VI FORECASTING AND DECISION FOR MULTIPLE ECOSYSTEMS SUBJECT TO REGIME SHIFTS

Introduction

All lakes have some essential similarities. The interaction of the phosphorus and carbon cycles is a key process for the functioning of nearly all lake ecosystems. Size-structured predation governs food web dynamics in nearly all lakes. Such repeatable patterns of biogeochemistry, food web organization, and other processes mean that information from one lake is often transferable to a different lake. Because of these fundamental similarities among lakes, lakes are to some extent substitutable or exchangeable for purposes of scientific inference. These similarities are the basis of comparative limnology (Cole et al. 1991, Rigler and Peters 1995). Other types of ecosystems, such as small watersheds or islands, may also be replicates in some respects. This degree of similarity among ecosystems provides ecosystem managers with potential approaches to the dilemma discussed in Chapter V.

Chapter V considered the possibility of anticipating regime shifts in advance for a single ecosystem, by monitoring key indicators and updating a dynamic model. The probability of measuring the threshold was low, unless a regime shift occurred. It proved difficult or impossible to learn the threshold for the regime shift without crossing it and transforming the ecosystem. If the regime shift is costly and difficult to reverse, then a cautious policy that minimizes risk of regime shift may be preferable to a learning policy that attempts to locate the threshold. Thus, in unique ecosystems subject to regime shifts, there is a dilemma between precaution and experimentation. Social-economic pressures to exploit the ecosystem may oppose precaution, but in the absence of information about the location of the threshold, one does not know which policies may cause a regime shift.

Prospects for measuring the threshold might improve, however, if many ecosystems are available for study and information is transferable among ecosystems. If a large number of similar ecosystems are available, it may be possible to use a few of them for experimentation, to gain information about the threshold and avoid unwanted regimes in the majority of ecosystems. That possibility is explored in this chapter. As in Chapter V, I will use an ecosystem model combined with models for learning and management to ask whether and how experimental learning might improve prospects for managing ecosystems subject to regime shifts.

A Landscape of Lake Fisheries

Consider a landscape with many lakes and fisheries subject to exploitation and management (Fig. 41). This general situation occurs in many regions of the world. For example, the walleye fishery of northern Wisconsin includes hundreds of lakes, which are managed using models calibrated with regular studies of a subset of 125 lakes (Beard et al. 2003a,b). Marine fisheries have also been assessed using the assumption that information on different stocks is exchangeable (Liermann and Hilborn 1997).

This chapter considers a set of similar lakes with exploited fisheries (Fig. 41). The dynamics within each lake ecosystem resemble those described in Chapter II for depensation in fish populations (Fig. 8). A piscivorous fish population is subject to harvest (Fig. 41). Juvenile piscivores are consumed by other species of fishes (forage fish), which are preyed upon by adult piscivores. As discussed in Chapter II, predator-prey interactions of this type can exhibit cultivation-depensation dynamics, in which recruitment of the piscivore depends on the capacity of adult piscivores to control the forage fish population (Walters and Kitchell 2001).

The piscivore stock collapses if the adult population falls below a certain critical level (Fig. 41). This phenomenon, called critical depensation or the Allee effect (Allee 1931, Begon et al. 1986), is marked by negative population growth rates when the population drops below the threshold for critical depensation. Critical depensation has important implications for resource management and conservation because it implies that population sizes well above zero may be nonviable (Dennis 1989, Stephens and Sutherland 1999). If this is true, then sustainable policies must maintain population levels well above the threshold, to ensure that stochastic events do not take the population below the threshold.

In the model used in this chapter, the critical depensation threshold depends on the availability of nearshore habitat for juvenile piscivores. Predatory interactions (shaded box in Fig. 41) occur primarily in nearshore habitat, such as rooted aquatic plants or trees fallen into the lake. If adequate habitat exists for juvenile piscivores, they suffer little mortality from predation by forage fishes. As habitat availability declines, the mortality of the juvenile piscivores increases. This phenomenon is discussed in Chapter 2 and illustrated in Fig. 8. The role of habitat in fish recruitment is important, because shoreland development practices are altering fish habitat and growth in some lakes (Christensen et al. 1996, Olson et al. 1998, Schindler et al. 2000). The documented effects of habitat on fish growth suggest potentially important impacts on size-structured predation and fish community change.

In aquatic ecosystems, size-selective interspecific predation creates many situations in which critical depensation is possible (Chapter II; de Roos and Persson 2002, Walters and Kitchell 2001). The location of the threshold may depend on life history characteristics of the fishes as well as characteristics of the habitat. Factors which may drive a population across the threshold include fishing, invasion of exotic species or weather during key times of the fishes' life history (Chapter II).

Although critical depensation is potentially widespread in aquatic food webs, and is profoundly important for management, it has proven difficult to quantify thresholds in advance of population collapses, or even to detect depensation statistically (Liermann and Hilborn 1997). This difficulty has led to controversy in both fishery management and conservation. Presence or absence of depensation has implications for the structure of ecosystem models and for harvest or conservation derived from them (Pascual et al. 1997, Runge and Johnson 2002). The problem, then, is one of choosing which model is most appropriate for decisions (Peterson et al. 2003). Ecosystem manipulations should be helpful in determining which model is most appropriate (Walters 1986, Petraitis and Latham 1999, Carpenter 2002).

I assume that the manager's task is to sustain the fisheries by regulating harvest rates. A program of monitoring and modeling is in place to guide the process of regulating harvest. Experiments using a subset of the lakes can be conducted to estimate the depensation threshold, and this information can be used to manage all the lakes of the landscape (Fig. 41). To assess the value of the information obtained from the experiments, it is necessary to measure the performance of the management system. I will use the expected future harvest over infinite time to measure management performance. Sport fisheries usually have substantial economic benefits which are largely related to expenditures on equipment and travel (Postel and Carpenter 1997). It seems reasonable to assume that expected future harvest over infinite time is an appropriate indicator of the total net economic benefits of managing the fishery.

Model

The model was designed to be as simple as possible, while retaining the necessary features demanded by the goals of this chapter (Fig. 42). These features include an ecosystem subject to regime shift, the possibility of different types of experiments to gain information about the threshold, plausible statistical models with different implications for management, and an explicit process for choosing harvest policies. The model first appeared in Carpenter (2002).

There is one state variable, a harvested fish population subject to critical depensation. The depensatory process depends on the quality of habitat, represented by the number of fallen trees in the littoral zone of the lake. The habitat is assumed to change slowly relative to the fish population. I assume that both the fish population and the number of fallen trees can be manipulated experimentally. Such experiments are presently underway in lakes of northern Wisconsin http://biocomplexity.limnology.wisc.edu/).

In this chapter, I will consider two contrasting experimental designs, one which yields little information and another which is highly informative. Data from experimental manipulations are fitted to two statistical models which (1) provide reasonably good fits to simulated data but (2) yield different estimates of the threshold and have sharply different implications for policy choice.

Population dynamics of the adult piscivores follow

$$A_{t+1} = A_t \exp(G_t + N_t) - C_t \qquad (10)$$

where A_t is adult stock in year t, C_t is harvest in year t, G_t is a function for population dynamics, and N_t is a function for stochastic shocks.

Population dynamics G_t are given by

$$G_t = k + f A_t (1 - m A_t - \{c (h/W_t)^q / [(h/W_t)^q + A_t^q] \}) \qquad (11)$$

Parameters include minimum population growth rate k, fecundity f, density-dependent mortality m, and three parameters for mortality due to interspecific predation on juvenile piscivores, c, h and q. W_t is the amount of habitat available for juvenile piscivores in year t. Maximum mortality due to interspecific predation is c, h determines the level of W at which interspecific predation occurs at half the maximum rate, and q determines the slope of the interspecific mortality curve near this half-maximal point. W_t is the time series of habitat. The interspecific predation term $\{c (h/W_t)^q / [(h/W_t)^q + A_t^q] \}$ is plotted against adult stock for two different levels of habitat in Fig. 8B. Interspecific predation on juvenile piscivores declines with adult population size, because the adults suppress the forage fishes that prey on the juveniles. Also, as the amount of habitat for juvenile piscivores (W) increases, predation on juvenile piscivores declines.

Process noise N_t represents stochastic effects on population dynamics due to factors such as weather. This term is given by

$$N_t = z - s^2/2, \text{ where } z \sim N(0,s^2) \qquad (12)$$

The process variance is s^2. The noise term in equation 12 is adjusted by subtracting $s^2/2$ so that the mean value of e^N is zero (Hilborn and Mangel 1997).

In the following, I refer to the true system dynamics as the dynamics given by equations 10 to 12. The deterministic portion of the true dynamics is calculated by setting $N_t = 0$.

For certain parameter values a threshold value of A can exist, below which the population collapses to zero, and above which the population reaches a stable point. This can be seen at equilibrium by rearranging the deterministic part of the model (Equations 10 and 12) as

$$\exp \{f\, A(m\, A + \{c\, (h/W)^q / [\, (h/W)^q + A^q] \}\,)\, \} = [A / (A + C)] \exp (k + f\, A) \qquad (13)$$

The left side of Equation 13 is the total rate of natural (non-harvest) mortality. The right side is the total birth rate adjusted by the proportion of the population that is not harvested. In this model, habitat is assumed to change more slowly than the fish population. This is the case in lakes where habitat is provided by trees fallen into the lake (Christensen et al. 1996), because tree population dynamics are generally slower than those of fishes. In that case it makes sense to calculate steady-state fish populations (see below) for a given level of habitat, as shown in equation 13. These are not true steady states because they change slowly over time as habitat changes (Rinaldi and Scheffer 2000).

When the left side and right side of equation 13 are plotted on the same axes (as in Fig. 43 described below), the intersections correspond to equilibria. In many cases, parameters lead to two intersections. In these cases, the lower equilibrium is an unstable threshold (the depensation threshold) and the upper one is stable. The threshold level is inversely related to W. Discussions of similar models are presented by Clark (1990), Carpenter (2001, 2002) and Walters and Kitchell (2001).

The management challenge is to sustain the fishery without crossing the threshold to collapse, given the complications of stochasticity, possibly unknown parameters which must be estimated from data, and slow change in W which may be poorly known.

The manager uses a subset of the lakes for ecosystem experiments to fit models for the piscivore population dynamics. The manager is aware that both population size and habitat may have important effects on the dynamics. Two types of experiments are considered, one in which contrasts are weak and another in which contrasts are strong. The experiments and their outcomes are described in the next section.

I assume that the manager does not know the true model, but instead fits approximate models to observed time series of A and W. Many approximate models are conceivable, but only two particularly interesting ones are analyzed here. In subsequent text, I refer to these as the estimated models. The first of these, referred to as the simple model, is a lag-1 autoregressive model with density dependence:

$$Y_t = b_0 + b_1 A_{t-1} + b_2 A_{t-1}^2 + \varepsilon_t \qquad (14)$$

The second estimated model, referred to as the habitat model, is the simple model augmented with a term that accounts for habitat effects

$$Y_t = b_0 + b_1 A_{t-1} + b_2 A_{t-1}^2 + (b_3 / W_{t-1}) + \varepsilon_t \qquad (15)$$

In extensive simulations , the simple habitat term, b_3 / W_{t-1}, produced reasonably good fits to data, leaving little pattern in the residuals. Thus the habitat term represents an estimator of habitat effects that might well be obtained from empirical regression analyses of data generated by equations 10 to 12. The b_i are parameters to be estimated from data and the errors ε_t are assumed to be normally distributed with mean 0 and unknown variance σ^2. The response series Y is

$$Y_t = \log [(A_t + C_{t-1}) / A_{t-1}] \qquad (16)$$

Predicted one-step-ahead distributions of $A_{t+1} = A_t\exp(Y_t) - C_t$ are obtained using the posterior distribution of Y_t. This posterior is calculated by analyzing equations 14 and 15 using Bayesian multiple linear regression with a prior distribution that is uniform on the b_i and $\log(\sigma)$ (Appendix; Gelman et al. 1995).

For the purposes of this chapter, I assume that management follows an optimal control policy. That is, the manager estimates the population size that leads to the maximum expected yield over infinite time, and adjust harvest to stabilize the population at that optimal level. Harvest policies of this type have many difficulties, both in theory and practice (Walters 1986, Hilborn and Walters 1992, Ludwig et al. 2001). In this chapter they are used as a benchmark to show how the difficulties of estimating thresholds can lead to extirpation of valuable populations. The analyses presented here depict only one of the difficulties of optimal control policies for living resources. For a broader view of the problems with optimal management of natural resources, and some alternatives, see Gunderson and Holling (2002).

If management follows an optimal control policy, the goal is assumed to be maximization of utility V derived from harvest (Clark 1990), defined here as

$$V = \Sigma \delta^t C_t \qquad (17)$$

In words, the performance of the fishery, V, is assumed proportional to the total future discounted harvest for all time. The summation is over time from the present to infinity, δ is the discount factor, and C_t is catch in year t. The economic discount factor δ ($0 < \delta < 1$) measures the manager's preference for current versus future harvest. For example, if $\delta = 0.98$ then 100 fish this year are considered to be as valuable as 98 fish next year, or $100 \times 0.98^2 \approx 96$ fish in two years. Discount factors are controversial in natural resource management (Heal 1997). For the purposes of this book, expression 17 is offered as an arbitrary benchmark for performance of the management system that is similar to the performance measures used by many other authors. Note that the economic discount factor used in this chapter is different in meaning from the statistical discount factor used in Bayesian nonlinear dynamic regression (Appendix).

The optimal policy can be found by maximizing the following expression with respect to harvest:

$$V(A_t) = C_t + \delta E[V(A_{t+1})] \qquad (18)$$

where E is the mathematical expectation operator, which corresponds to the mean over the posterior distribution of $V(A_{t+1})$ (Clark 1990). This expected value depends on the policy choice C_t. The optimal policy is found by establishing a target population size y and moving to that level as quickly as possible (Clark 1990). Define u(A,y) as the harvest that will move the population to a target level y in one time step

$$u(A,y) = A \exp(M_i) - y \qquad (19)$$

In equation 19, M_i is calculated from one of the fitted models (Equations 14 or 15), or from the true model (Equations 10 to 12). In the latter case, M_i corresponds to G from equation 11. When one of the fitted models is used, M_i corresponds to Y from equation 14 or 15. Because harvest cannot be negative, we set

$$C_t = u(A_t,y) \text{ if } u(A_t,y) \geq 0 \tag{20A}$$

$$C_t = 0 \text{ if } u(A_t,y) < 0 \tag{20B}$$

Each model yields a distribution of future values of A_t. The distribution of $\log(A_t)$ is normal in the case of the true model, and Student-t in the case of the fitted models (Gelman et al. 1995). In computing policies we must account for the probability of obtaining each possible value of A_{t+1} given a particular value of A_t. This was done by computing probabilities on a discrete mesh, then weighting each possible value of A_{t+1} by its probability, as in Carpenter et al. (1999b) Appendix C. I used a mesh of 50 grid cells uniformly distributed over the interval $[\log(\hat{A}) - 4\ s_{\log(\hat{A})},\ \log(\hat{A}) + 4\ s_{\log(\hat{A})}]$ where $\log(\hat{A})$ is the expected value of $\log(A)$ and $s_{\log(\hat{A})}$ is the standard deviation for the prediction of $\log(A)$. Given these definitions, equation 18 can be computed for any given value of y. The value of y that maximizes V is computed by Brent's method (Press et al. 1989) over the interval $[0,A_t]$.

When the experimental data do not clearly favor one model over another, model uncertainty may have a significant influence on policy calculations for resource management (Walters 1986, Peterson et al. 2003). Model uncertainty refers the probability that a given model is true. It is different from parameter uncertainty, which relates to the probability distribution of parameters, and from the error term of the model, which relates to the probability distribution of model residuals. Like parameter uncertainty and the model error term, model uncertainty affects the uncertainty of predictions.

To account for the effect of model uncertainty on predictions, I calculated a third estimate of optimal population size by Bayesian model averaging, or BMA (Raftery et al. 1997, Fernández et al. 2001). BMA is a weighting method, in which the impact of each model on the policy is weighted according to the credibility of the model. The credibility is based on the model's fit to data. According to this procedure, the Bayesian model average utility, V_{BMA}, for a given y value is calculated over both fitted models (Raftery et al. 1997, Fernández et al. 2001) as:

$$V_{BMA}(y) = p\ V_D(y) + (1-p)\ V_H(y) \qquad (21)$$

V_D is the utility under the simple density-dependent model (Equation 14) given y, V_H is the utility under the habitat model (Equation 15) given y, and p is the posterior probability of the simple density-dependent model. The posterior probability p is computed by the method of Fernández et al. (2001, using her equations 2.12 and 4.1). In this exercise, there are only two models so the posterior probability of the habitat model is 1-p. The optimal population size is computed by finding the target level y which maximizes $E[V_{BMA}(y)]$. As described above for the individual models, this expectation was computed over a discrete mesh, assuming a Student-t distribution for the predictive distribution of the two estimated models. The value of y that maximizes $E[V_{BMA}(y)]$ is computed by Brent's method (Press et al. 1989) over the interval $[0,A_t]$.

Results

The birth and death rates of the fish population for three different habitat conditions were calculated using Equation 13 (Fig. 43). In all cases the birth and death rate curves cross at two values of A. The right intersection point is stable: at slightly larger values of A, deaths exceed births and drive the population down to the intersection point, while at slightly smaller values of A, births exceed deaths and drive the population up to the intersection point. This stable value of A is the same for all three habitat conditions. The left intersection point, the threshold for collapse of the fish stock, is unstable. At slightly larger values of A, births exceed deaths and drive the population toward the right intersection point. At slightly lower values of A, deaths exceed births and the population is driven extinct. This threshold occurs at increasingly large population sizes as the habitat deteriorates.

Two simulated experiments were calculated to fit the approximate models of equations 14 and 15 (Table 6). In each experiment, fish stock size and habitat condition are manipulated in a number of lakes, and then the fish dynamics are observed for two years. The “strong contrast” experiment is an example of data which might derive from an aggressive program of experimental management. It substantially reduces the model uncertainty. The “weak contrast” experiment represents data which poorly discriminate the models. Note, however, that even this rather uninformative experiment assumes that 6 manipulated lakes are studied. Also, in these simulations I have assumed that observation error is negligible due to intensive sampling. Thus even the "weak contrast" data in this exercise are substantially better than those available for most lake management decisions.

Posterior probabilities of the two fitted models were calculated by the method of Fernandez et al. (2001) (Table 6). The weak-contrast experiment does not discriminate the two models; each has a posterior probability near 0.5. The strong-contrast experiment detects the habitat effect. Consequently, given the strong-contrast experiment there is a high posterior probability for the habitat model and a low posterior probability for the simple model.

The experiments differed sharply in their estimation of the threshold (Table 7). With strong-contrast data, the estimated threshold was within about 10 fish / ha of the true threshold. The estimated threshold decreased as habitat quality increased, consistent with the true ecosystem dynamics. Standard deviations of threshold estimates were 11 to14 fish ha^{-1}, reflecting the magnitude of uncertainty about the location of the threshold. With weak-contrast data, the estimate of the threshold was much more uncertain. The mean estimate was negative in all cases. There was no clear trend of the threshold with changing habitat conditions. In practice, this would be interpreted as evidence for only one positive intersection of the birth and death curves, the stable one (Fig. 43). The analyst would conclude, incorrectly, that depensatory collapse was not possible for this fish stock. However, the standard deviations are quite large, indicating considerable uncertainty about the location of the threshold.

I compared the performance of optimal management using weak-contrast data and strong-contrast data for lakes with the three habitat conditions shown in Fig. 43. In each case, I will compare optimal population targets as a function of fish population size (Figs. 44 to 46). These targets are computed under 4 models: the fitted simple model, the fitted habitat model, the Bayesian average of these two fitted models, and the true model. In each case, it is interesting to assess the performance of the models near the threshold for fish population collapse.

When the fish habitat is poor, the threshold for collapse occurs when the fish population declines to about 50 fish ha^{-1} (Fig. 43A). Note that in this case the habitat level used for the policy calculations was a treatment level in the experiments (Table 6). Given the true system dynamics, the policy calculations show correctly that the population must drop to zero if the initial stock size is less than about 50 fish ha^{-1} (Fig. 44D). With weak-contrast data, under which the simple model is credible (Table 6), the simple model indicates that positive populations can be maintained with initial stock sizes as low as about 42 fish ha^{-1} (Fig. 44A). With strong-contrast data, the simple model indicates that positive populations can be maintained from initial stock sizes as low as nearly 20 fish ha^{-1}, but the credibility of the simple model is very low if the contrast in the data is strong. The habitat model is close to the correct threshold with weak-contrast data (Fig. 44B). With strong-contrast data, the habitat model incorrectly indicates that positive fish populations can be maintained when the initial stock is as low as about 42 fish ha^{-1}. The Bayesian model average gives similar results for both weak-and strong-contrast data (Fig. 44C). It indicates incorrectly that positive stocks can be maintained if the initial stock is as low as about 42 fish ha^{-1}.

When the fish habitat is moderate, the threshold for collapse occurs at a fish population of about 22 fish ha^{-1} (Fig. 43B). In this case, the habitat level used for the policy calculations is not included in either experiment, but is close to one of the experimental treatments used in the strong-contrast data (Table 6). Given the true population dynamics, the policy calculation is close to the true threshold (Fig. 45D); the small error is attributable to the mesh interval used for initial stock sizes. The simple model is close to the true threshold given strong contrast data (Fig. 45A), but of course this model has low credibility under the strong contrast data. With weak contrast data, the simple model indicates, incorrectly, that the threshold is about 40 fish ha^{-1}. The habitat model indicates that the threshold is about 30 fish ha^{-1} for both weak- and strong-contrast data (Fig. 45B). The Bayesian model average indicates that the threshold is between about 30 and 40 fish ha^{-1}, depending on the data set (Fig. 45C).

When the fish habitat condition is high, the threshold for collapse occurs at a fish population of about 17 fish ha^{-1} (Fig. 43C). In this case, the habitat level used for the policy calculations is included in the strong-contrast data set, but must be extrapolated in the weak-contrast data set (Table 6). Given the true population dynamics, the shift in policy occurs below the true threshold (Fig. 46D). This is a conservative policy, in the sense that a population below the threshold is underharvested and may recover due to a fortunate stochastic shock. The simple model overestimates the true threshold under both data sets (Fig. 46A). The habitat model accurately estimates the threshold for the strong-contrast data set, and overestimates the threshold for the weak contrast data set (Fig. 46B). The Bayesian model average also gets the threshold about right given the strong-contrast data, and overestimates the threshold given the weak-contrast data (Fig. 46C).

To assess the overall performance of the estimated models, the difference in optimal stock sizes between the true model and the Bayesian model average was calculated (Fig. 47). This difference is zero when the estimation is perfect. Note that, in every case, the deviation between true and estimated models is greatest near the threshold. True threshold values were about 50 fish ha^{-1} for poor habitat, 22 fish ha^{-1} for moderate habitat, and 17 fish ha^{-1} for good habitat (Fig. 43). With weak-contrast data, the performance is significantly worse. The estimated model routinely underestimates optimal population size, an error that will lead to overharvest.

Discussion

The model is a highly simplified representation of fish management. The simplifications were designed to investigate the effects of ecosystem experimentation on inference about thresholds and optimal harvest policies based on these inferences. The exercise shows that thresholds are difficult to measure without strong experimental manipulations that yield high-contrast data. Even when such data are available, uncertainty is high near the threshold and policy choices are strongly affected by this uncertainty.

There are several assumptions that have strong effects on these findings. First, the model assumes that factors governing fish dynamics are the same across lakes, so the model and parameter estimates apply to all the lakes on the landscape. If it is not appropriate to describe the lakes using a common model, then the management problem is more similar to the case described in Chapter V, where each ecosystem is unique. This issues could be explored using a hierarchical Bayesian model (Gelman et al. 1995), in which the parameters for the threshold of a particular lake were viewed as a sample from a meta-distribution of parameters for the set of lakes on the landscape.

The model assumes that the threshold is constant over time. In reality, the threshold could move slowly over time due to changes in habitat, food web composition, or other ecological factors that were not included in the model. If the threshold moves over time, then repeated experiments will be necessary to track the changes. In general, this will increase uncertainty for model predictions and policy choices at any given time.

It is assumed that the habitat and harvest can be manipulated independently in entire lakes. Such replicated, multi-factor ecosystem experiments are rarely possible. Most ecosystem experiments have involved one or two manipulated ecosystems and one or two reference (or control) ecosystems (Carpenter et al. 1995, Carpenter 1998).

When an ecosystem experiment reveals something important, the experiment tends to be repeated by different research groups in different regions (Carpenter 1998). This type of replication leads to especially robust inferences, because it samples across regional differences in ecosystems as well as differences in methodology and practice among research groups. However, this type of replication takes a long time. It leads to inferences more slowly than the experiments assumed in this model. Slow inference could lead to much greater uncertainty about the location of the threshold, especially if the threshold is moving over time.

Some of the experimental manipulations assumed in the model are drastic ones that involve massive changes to habitat, deliberate overharvest of fish stocks, or severe reductions in harvest of fish stocks. While these extreme treatments are the most important ones for identifying the threshold, they are the most difficult ones to conduct. Extreme manipulations often face political barriers (Walters 1997, Carpenter 1998). When extreme manipulations are omitted, the situation is similar to the “weak contrast” situation shown in this chapter: high uncertainty, poorly-characterized thresholds, underestimate of optimal stock sizes, and overharvest with risk of inadvertent collapse. Ironically, a manager who avoids extreme manipulations will promote overharvest and risk collapse of fish stocks in most of the lakes on the landscape. While the extreme manipulations will damage some ecosystems, they lead to better policies in many other ecosystems. For the manager who is responsible for a landscape of lakes, there are significant advantages in investing in a few extreme experimental treatments.

Data and inference about thresholds

This exercise shows that analyses of thresholds demand excellent data. Given strong-contrast data, it may be possible to locate the general neighborhood of a threshold and take appropriate management action.

Even with excellent data, however, it is very difficult to locate the threshold with precision. This modeling exercise employed simplifying assumptions that made it easier to locate the threshold. Despite these assumptions, there was considerable variability in estimates of the threshold. Uncertainty about the threshold resulted from the process variance of the fish dynamics and uncertainty about which model was appropriate. In a more realistic situation, there would be additional sources of uncertainty, including observation error and dynamics of other slowly-changing variables (in addition to W) that could cause the threshold to shift over time. Also, I have assumed that whole-ecosystem observations for similar lakes are available for fitting the models. In reality, it might be necessary to estimate some parameters using data from small-scale experiments. In this case there are additional uncertainties due to extrapolating parameters from one scale to another. Therefore, this chapter has presented a rather optimistic scenario for estimating thresholds. In most practical situations, the uncertainties will be substantially larger than portrayed here.

In most cases, the estimated models underestimate the size of the population that can be sustained in the neighborhood of the threshold. Put another way, the estimated models tend to indicate that the threshold occurs at a population size larger than the true threshold. The impact of this overestimate on decisions would depend on the goals of the manager. A manager who was purely interested in optimal yield would assume that the population should be driven extinct from starting levels below the apparent threshold, but above the true the threshold. This error would lead to extirpation of stocks that are in fact potentially sustainable. A manager who sought to avoid thresholds, however, would behave more conservatively. Such a manager would harvest cautiously at stock sizes well above the true threshold, because of the fitted models' overestimation of the true threshold. These very different outcomes depend entirely on the managers' preference for maximizing expected present value of harvest versus avoiding the threshold for collapse.

Effect of the threshold on policy choice

Near the threshold, there is a sharp change in the optimal control policy. Above the threshold, the policy is to sustain a positive fish stock. Below the threshold, the policy is to harvest all remaining fish and drive the stock to zero. When the model for fish dynamics is estimated from data, the estimated location of the threshold has a powerful effect on policy choice. In any particular situation, the decision depends on the initial stock size, the choice of model for fish dynamics, and the estimates of parameters for that model, including parameter variances and the model variance.

The “use them or lose them” feature of the optimal control solutions occurs in many analyses of this type (Clark 1990, Ludwig et al. 2003). It is a consequence of the arbitrary decision to maximize net future discounted harvest. If this goal was replaced by a different goal, such as to avoid the threshold, different policies would result (Limburg et al. 2002). If the goal is to avoid collapse, one would sharply reduce harvest at stock sizes well above the threshold. At lower stock sizes, one would try to rebuild the population, for example by stocking or creating connections to refuges.

Implications for experimental management

If the goal of management is to avoid the threshold, ecosystem experiments can significantly improve the information for choosing management targets. The experiments also show the cost of crossing thresholds in direct and powerful ways. Thus ecosystem experiments combined with precautionary management appear to be a promising approach for modular ecosystems like lakes.

A landscape of similar ecosystems offers interesting possibilities for actively adaptive experimental management. By using a subset of the available ecosystems for strong-contrast experiments, it is possible to gain substantial information about regimes and thresholds. However, it is risky to use this information for optimal control near the threshold. If the goal of management is sustaining the ecosystem, then a more precautionary approach is necessary. In this case, the precautionary choices are informed by relatively precise information about the location of the threshold. The precautionary manager should recognize that the location of the threshold is uncertain and that the threshold might be moving. The experimental results are best used to design rough guidelines and simple rules for maintaining the system in desired regimes.

Strong-contrast data provide better information about the location of the threshold, and thereby lead to policies that are more likely to avoid the threshold while coming close to maximizing the long-term yield of the fishery. With weak-contrast data, the performance of the management system is significantly worse. One can reach the erroneous conclusion that there is no threshold. At best the estimate of the threshold will have high variance. The resulting policies aim for stock sizes that are too low. Consequently, stocks are overharvested, leading to frequent collapses as thresholds are crossed. These unwanted stock collapses are a direct result of weak-contrast data.

Extreme manipulations are necessary to create strong-contrast data. These extreme manipulations will damage some ecosystems, and for this reason they may be opposed by the public and by some managers. However, strong-contrast data are a good investment when it is possible to manipulate a few ecosystems to obtain information that is useful for managing many ecosystems. Weak-contrast data can lead to far more damage as ecosystems are over-exploited due to poor knowledge of thresholds.

Summary

Actively adaptive experimental management has great potential for landscapes with a multiplicity of similar ecosystems. While lakes are the focus of this chapter, analogous approaches may be possible for other types of modular ecosystems, such as small watersheds or islands. Experimentation on a subset of the ecosystems can reveal which regimes are possible and provide rough estimates of the thresholds. Such information is valuable for maintaining desired regimes.

Management actions may cross thresholds inadvertently when thresholds are poorly known. Results of this chapter show that poor knowledge of the threshold leads to overly aggressive harvest policies that increase the risk of crossing the threshold. This may occur in many situations where policies push ecosystem state near to unknown thresholds. It may be possible to devise utility functions that lead to less risky choices when a threshold is more uncertain (Heal and Kriström 2002).

Successful experiments involve strong manipulations that cross thresholds in some ecosystems. Such experiments yield data with strong contrasts between experimental treatments. While locations of thresholds will be uncertain after even the best experiments, experimental estimates of thresholds lead to better policy choices within the framework described in this chapter. Prospects for maintaining a desired regime are improved when experimental estimates of thresholds are employed in precautionary policies that seek to avoid thresholds.

To provide useful information about thresholds, thresholds must be crossed in some experimental ecosystems. This means that some ecosystems will be damaged by the manipulations. This damage must be weighed against the improved performance of management for the other ecosystems of the landscape. If extreme manipulations are not performed, the weak-contrast data that result lead to a number of policy problems related to high uncertainty. Within the framework of this chapter, weak-contrast data lead to erroneous conclusions about absence of a threshold, or large errors in estimating the location of the threshold. The consequence is overexploitation with high risk of inadvertently crossing the threshold to an unwanted regime. Avoidance of extreme manipulations has severe consequences, leading to ecosystem damage that far exceeds the potential impact of the experiments.

Even under the best conditions, estimates of threshold location for real ecosystems will have substantial errors. Given the complexity of ecosystems, and the challenges of experimenting with them, it seems unlikely that precise quantitative measurements of any important threshold will be attainable. Nevertheless, ecosystem experiments that cross thresholds lead to rough estimates of threshold location that could be extremely useful in management. Such experiments reveal the levels of exploitation that can be maintained without a costly regime shift. Policies that seek to avoid collapse by maintaining ecosystems far from thresholds, and well within desirable domains, may succeed. If management seeks to avoid the threshold, this implies a different utility function with a heavy penalty for transcending the threshold (Heal and Kriström 2002). Thus, the combination of experimentation on a subset of ecosystems with precautionary management appears to be a promising management approach for modular ecosystems such as lakes, small watersheds or islands.

Tables

Table 6. Description of the simulated experiments.

Contrast	Number Of Lakes	Initial Fish Stocks (A)	Habitat Conditions (W)	Posterior probability for Simple Model	Posterior probability for Habitat Model
Weak	6	10, 55, 100	0.4, 1.4	0.49	0.51
Strong	25	10, 32, 55, 77, 100	0.4, 0.75, 1.1, 1.45, 1.8	0.01	0.99

Table 7. Estimated threshold (fish ha^{-1}) for an unharvested fish population given the weak-contrast and strong-contrast data sets described in Table 6. Means and standard deviations (in parentheses) are presented. The threshold was estimated using the simple model (Equation 14) and the habitat model (Equation 15) at three habitat levels, assuming C = 0. Means and standard deviations of the threshold were calculated by parametric bootstrapping from the multivariate Student-t distribution of estimated parameters and the Student-t distribution of residuals (Efron and Tibshirani 1993).

	Simple Model	Habitat Model	Habitat Model	Habitat Model
Contrast	All W values	W = 0.4	W = 0.8	W = 1.8
Weak	-58.1 (1530)	-19.4 (277)	-14.3 (250)	-18.9 (249)
Strong	26.2 (14.1)	40.3 (11.6)	27.1 (11.8)	18.2 (11.3)

Figures

Figure 41. Overview of the system analyzed in this chapter. Many similar lake fisheries on a landscape are subject to management. A subset of the lakes is available for experimentation to gain information about fish population dynamics and possible thresholds. (Original)

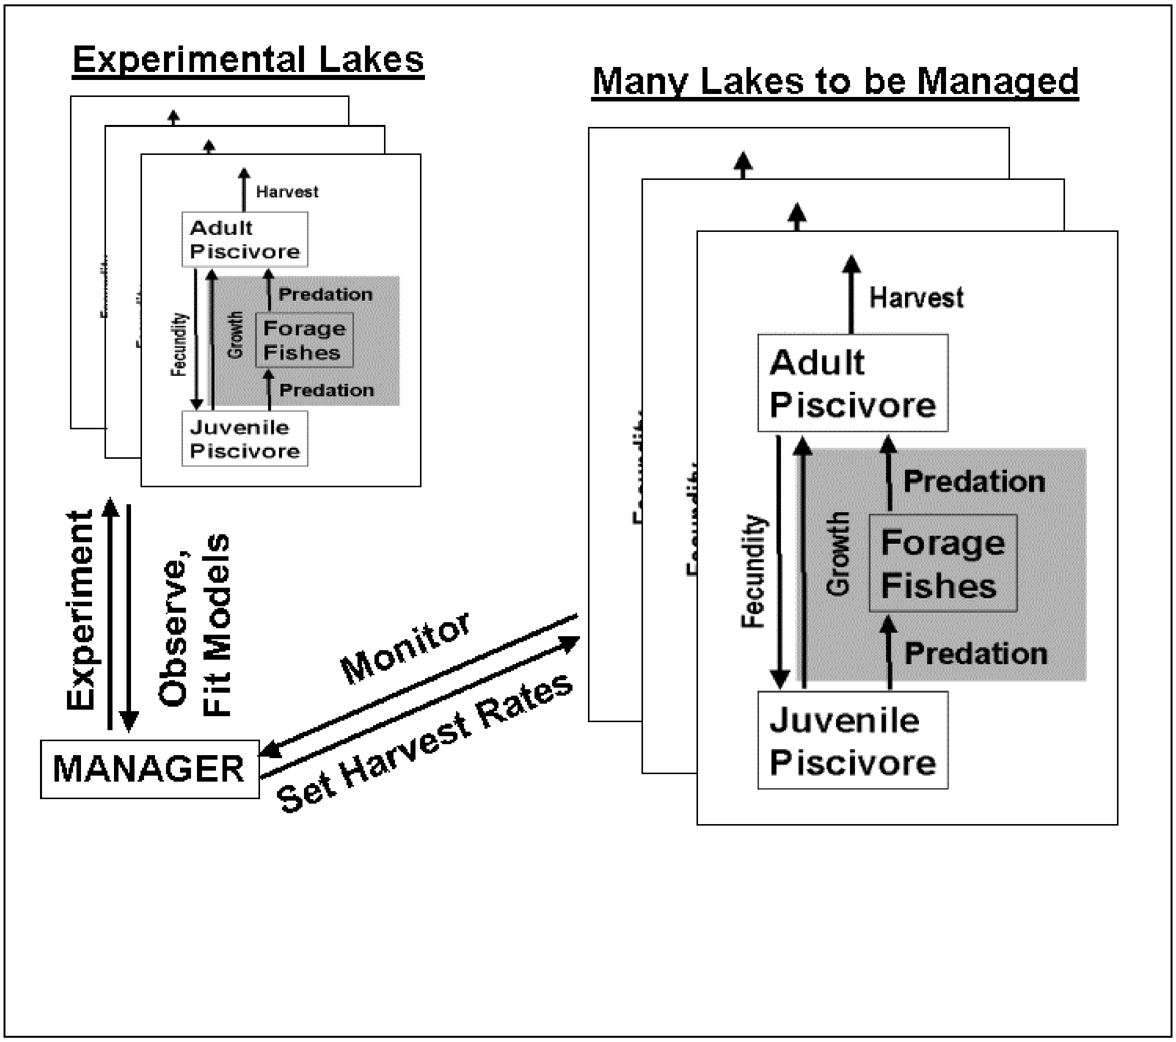

Figure 42. Flow chart of model calculations presented in this chapter. (Original)

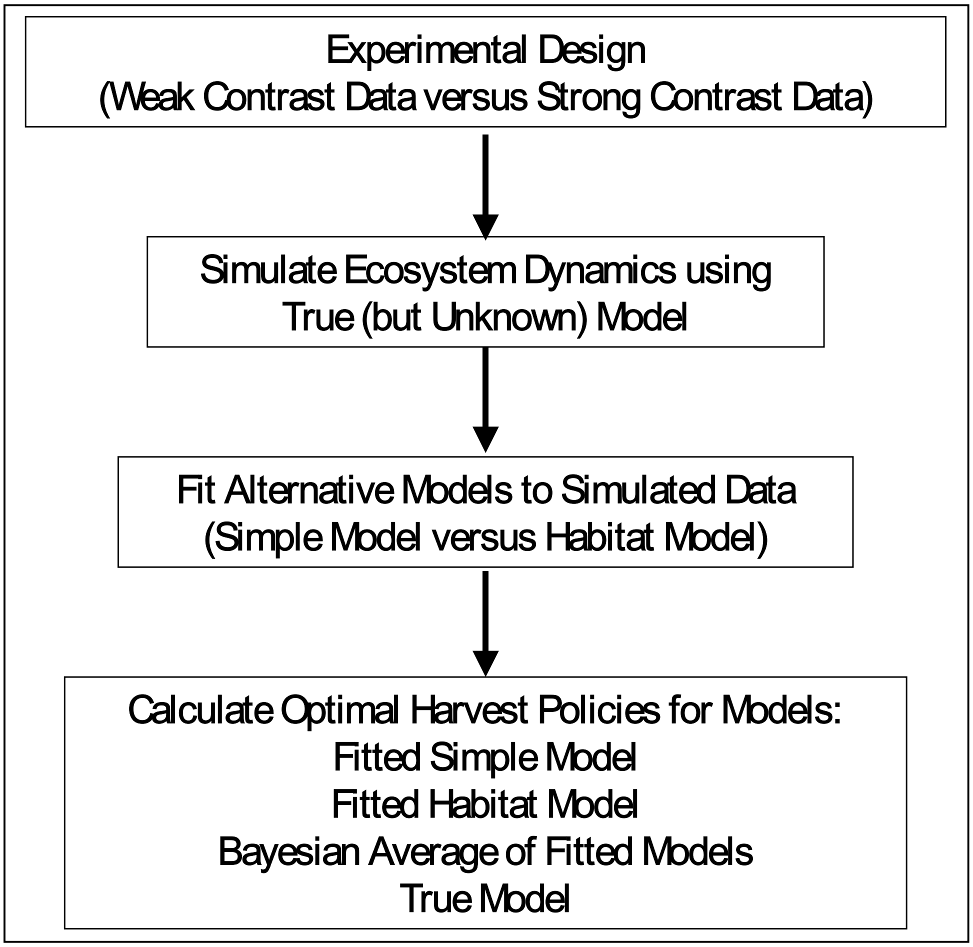

Figure 43. Vital rates from the deterministic part of the fish population model (births solid line, deaths dashed line; both in fish ha^{-1} y^{-1}) versus population size (fish / ha). A. Poor habitat (W = 0.4). B. Moderate habitat (W = 0.8). C. Abundant habitat (W = 1.8). Values of other parameters are C = 4, q = 4, k = 0.1, f = 0.01, m = 0.01, c = 10, h = 10. (Original)

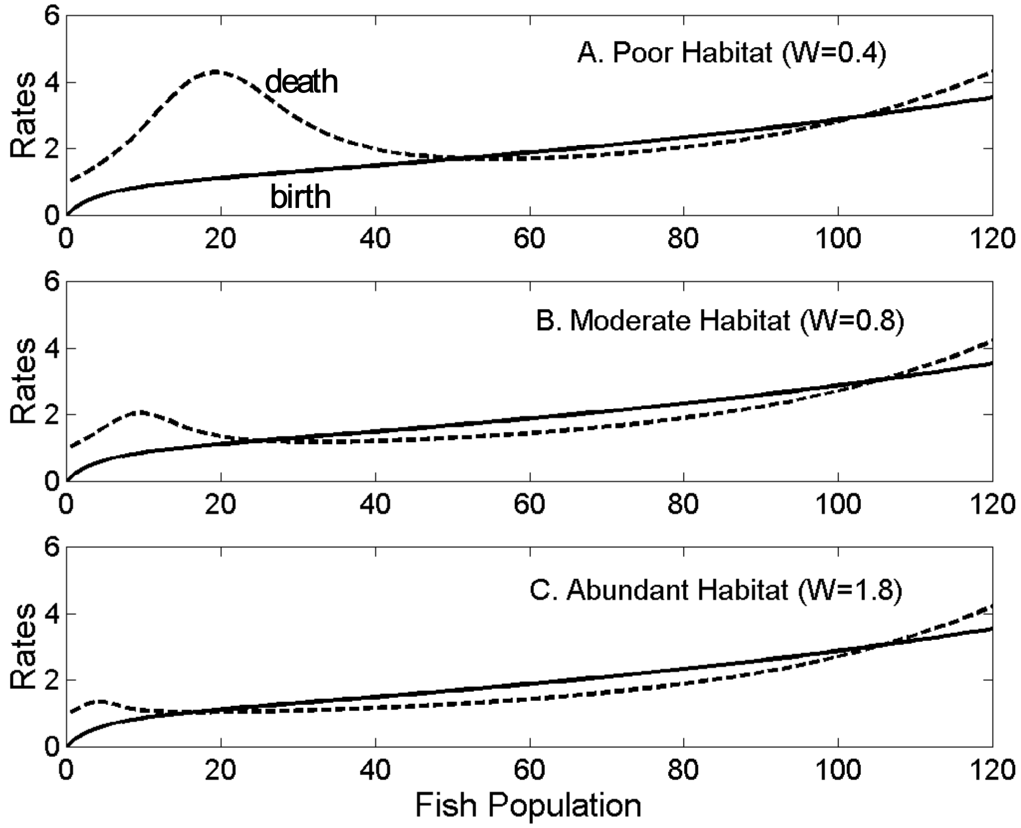

Figure 44. Apparently optimal population size (fish / ha) versus initial fish population (fish / ha) in simulations with poor habitat (W = 0.4). In panels A-C, optimal population sizes are calculated by fitting models to data with weak contrast (dashed line) or strong contrast (solid line). A. Simple model with density-dependence. B. Habitat model, with density dependence and a habitat effect. C. Bayesian model average of the simple model and the habitat model. D. Optimal population size calculated from true system dynamics. Values of parameters as in Fig. 1. Additional parameter values are s = 0.1, $\delta = 0.98$. (Original)

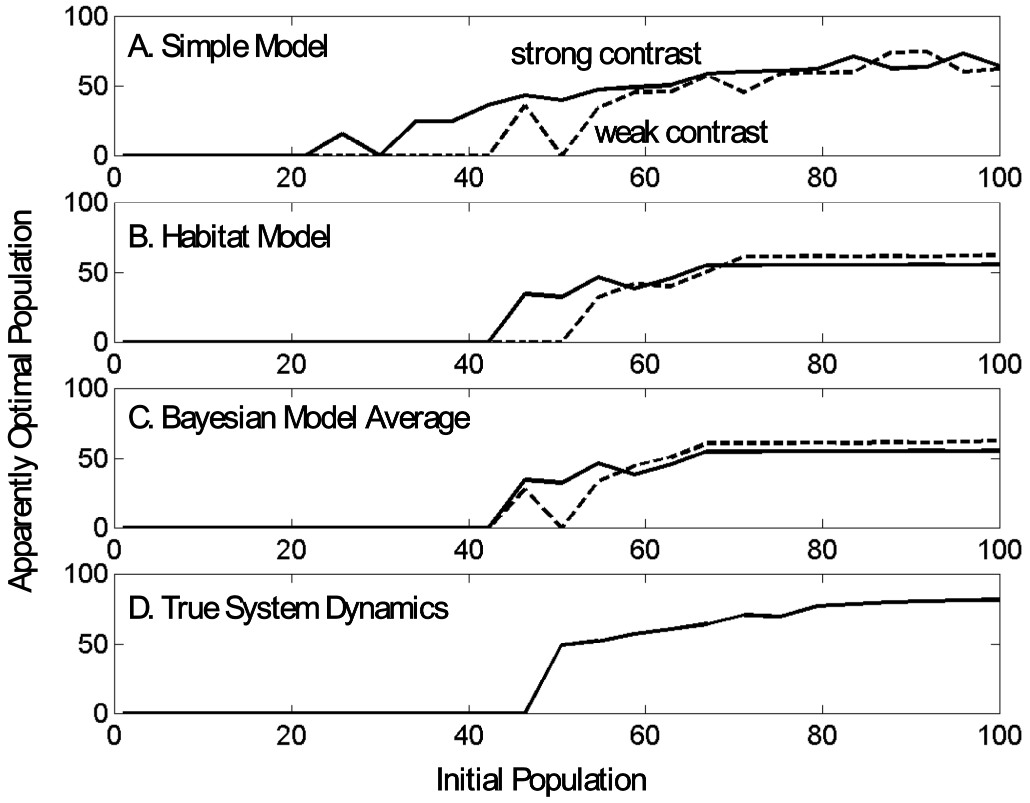

Figure 45. Apparently optimal population size (fish / ha) versus initial fish population (fish / ha) in simulations with moderate habitat (W = 0.8). In panels A-C, optimal population sizes are calculated by fitting models to data with weak contrast (dashed line) or strong contrast (solid line). A. Simple model with density-dependence. B. Habitat model, with density dependence and a habitat effect. C. Bayesian model average of the simple model and the habitat model. D. Optimal population size calculated from true system dynamics. Values of parameters as in Fig. 1. Additional parameter values are s = 0.1, δ = 0.98. (Original)

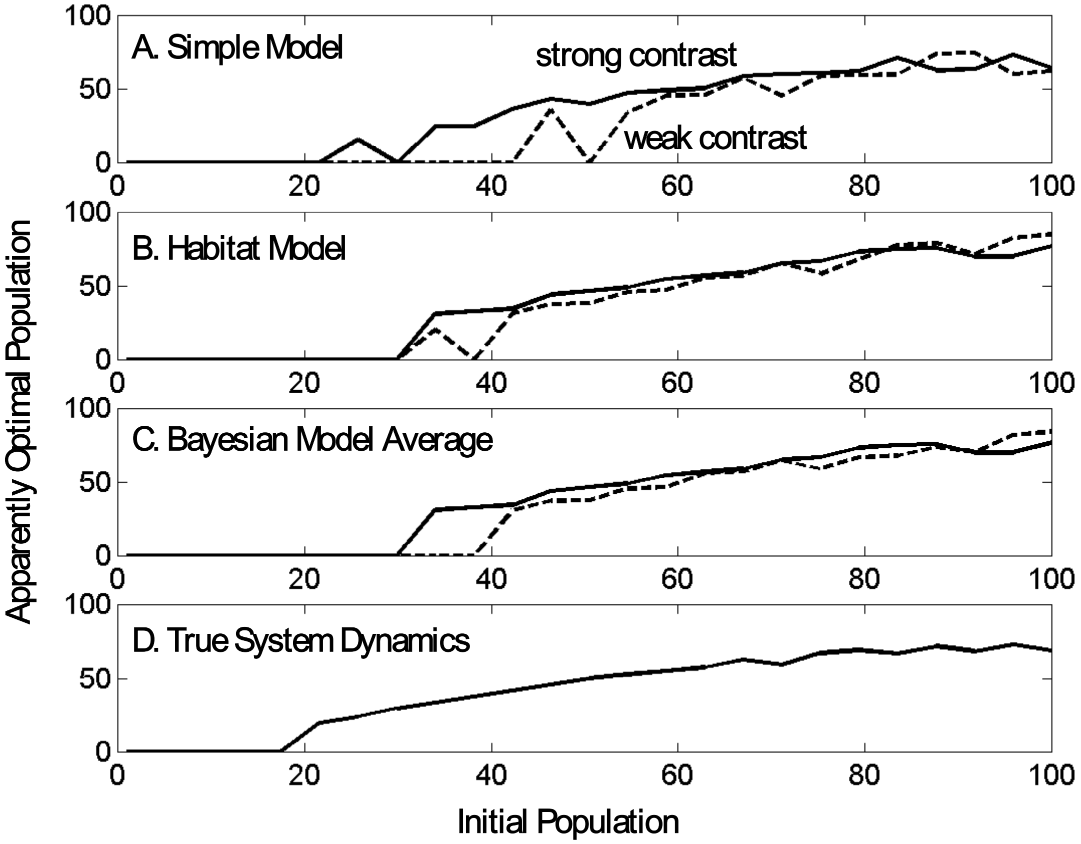

Figure 46. Apparently optimal population size (fish / ha) versus initial fish population (fish / ha) in simulations with abundant habitat (W = 1.8). In panels A-C, optimal population sizes are calculated by fitting models to data with weak contrast (dashed line) or strong contrast (solid line). A. Simple model with density-dependence. B. Habitat model, with density dependence and a habitat effect. C. Bayesian model average of the simple model and the habitat model. D. Optimal population size calculated from true system dynamics. Values of parameters as in Fig. 1. Additional parameter values are s = 0.1, δ = 0.98. (Original)

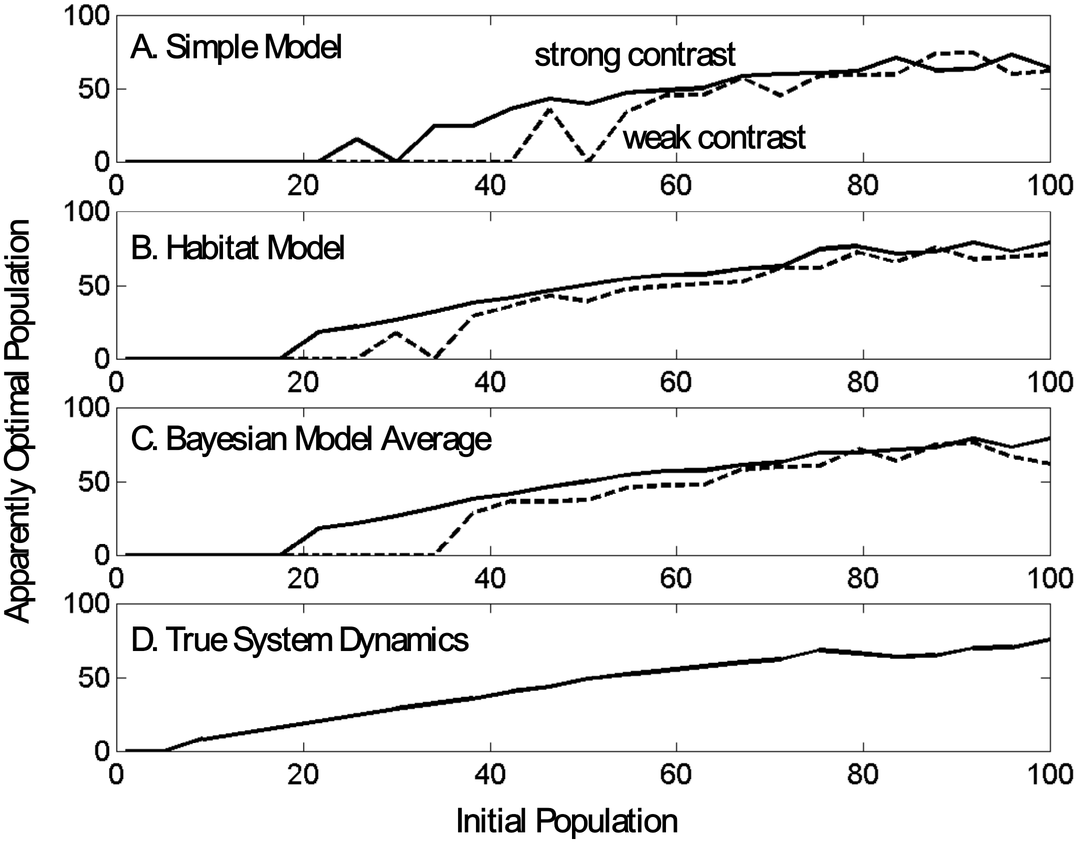

Figure 47. Difference between estimated optimal population size and the optimal population size (fish / ha) computed from the true system dynamics versus the initial population (fish / ha). The y-axis value is calculated as the estimated optimal population size from Bayesian model averaging minus the optimal population size if the true system dynamics were known. Solid lines show results with strong-contrast data set, and dashed lines show results with weak-contrast data set. A. Poor habitat (W = 0.4). B. Moderate habitat (W = 0.8). C. Abundant habitat (W = 1.8). All other parameters as in Figs. 2 – 4. (Original)