The economic optimality of learning from marine protected areas

Klaas Hartmann[1]   Lance Bode[2]   Paul Armsworth[3]

Abstract

Marine Protected Areas (MPAs) are an emerging tool for managing marine resources. Many of the benefits associated with MPAs have been widely investigated and the field is an active area of research in theoretical ecology. One benefit of MPAs that has remained largely overlooked is their value as a tool for learning about the population dynamics of a fishery. We investigate the economic optimality of implementing an MPA, purely for the purpose of obtaining more informative data about a fish population, thereby allowing a better management strategy. A stochastic dynamic programming framework for finding optimal management strategies in this scenario is developed. A simple example is investigated using this framework, with the results illustrating that in some situations the knowledge gained from MPAs can be sufficient to make their creation economically optimal. This establishes an additional benefit of MPAs that should be considered further by fishery managers.

Contents

1   Introduction

The fish in the world's oceans were long considered an inexhaustible resource [9]. Until the advent of the more modern fishing equipment and techniques of the nineteenth century, this was for the most part an adequate assumption. However, today, 22% of the world's fisheries are over exploited or depleted, and a further 44% are fully to heavily exploited [6, 5]. Currently, more than 1 billion people worldwide rely on fish for over 30% of their animal protein intake [5], this fact emphasises the importance of sustainable fisheries management.

Efficient fisheries management requires a quantitative understanding of the underlying population dynamics. Unfortunately, due to the natural variability of the system and the difficulty in observing the stock, it is difficult to acquire such an understanding. Generally there will be a number of different possible models, each of which suggests a different management strategy, some of which may be contradictory. For example, the optimal target stock size for one model may be predicted by an alternative model to cause a population collapse. A compromise must therefore be made and the resulting management strategy is unlikely to be truly optimal.

Knowledge about the population dynamics can be improved by monitoring the change in the population density from year to year. However, the new data are usually obtained from a small range of population densities [17], often coinciding with that for which historical data are already available. Consequently little may be learnt.

One method for obtaining informative data is to purposefully manipulate the population density, either by fishing down the stock or letting it build up. This is known as active adaptive management, and was initially investigated by Walters [16], who found that in some situations it is economically optimal to temporarily reduce fishing in order to learn about the population dynamics. In this case the short term decrease in yield is outweighed by the long term economic gain, due to the improved management method that becomes possible.

Fishery managers will generally encounter major political difficulties in reducing fishing (even for the short term), especially as the justification is in essence to conduct an experiment on the fishery. On the other hand, fishing down a stock will provoke criticism from conservation groups, and it may be difficult to rebuild the stock after information has been gained (as suggested by the number of over exploited fisheries worldwide). An alternative option, which we investigate here, is to establish a marine protected area (MPA; an area where fishing is forbidden) in which higher population densities can be observed. MPAs are an emerging tool in fisheries management [10], with many

widely discussed benefits [12, 8, 14, 2, 1, for example].

In this article we present a stochastic dynamic programming (SDP) framework for calculating the economic benefit of creating an MPA exclusively for the purpose of learning about the population dynamics of a fishery. By applying this framework to a simple scenario, we show that in some instances the value of the information obtained from an MPA may be sufficient to make its creation economically optimal. This establishes an additional benefit of MPAs.

2   The management problem

In the scenario that we consider, the fishery is managed by setting a catch quota, C_t, in each year, t. Finding the optimal catch quota for the current year is the problem the fishery manager is faced with. We assume that, based on historical data, n possible models for the population dynamics have been proposed, and that each model i has some probability $\sigma_t(i)$ of being the most appropriate for the fishery. These n models are difference equations and have the form

$$\mathbf{E}(N_{t+1}) = g_i(N_t - C_t), \quad 0 < i \le n.$$

Here the expectation $\mathbf{E}(\cdot)$ is over process noise; any form of noise can be considered. However, we concentrate on normally distributed multiplicative noise such that

$$N_{t+1} = \epsilon g_i(N_t - C_t), \quad \epsilon \sim N(1, s^2), \quad 0 < i \le n. \tag{1}$$

The probability associated with each of the possible models is updated annually by the manager using Bayes' Law. Updating probabilities while taking both process noise and observational error into account is a complicated process, as outlined by Meyer and Millar [11]. It is infeasible to handle this complexity in an SDP framework, and hence the assumption of a perfectly

observable population is made. With this assumption and Bayes' law, the model probabilities are updated at a time $t+1$, using the observed population density at that time (N_{t+1}):

$$\sigma_{t+1}(i) = \frac{\sigma_t(i)p(N_{t+1} \mid N_t, i)}{\sum_{i=1}^{n} \sigma_t(i)p(N_{t+1} \mid N_t, \sigma_t)}, \tag{2}$$

where $p(N_{t+1} \mid N_t, i)$ is the probability of model i producing the observed population density at time $t+1$, given the previous population density (which is simply the probability of obtaining $\epsilon = N_{t+1}/g_i(N_t - C_t)$).

Fishery managers should be concerned not only with the catch obtained this year but also with the catch obtainable in future years—for example, if all fish are caught today, then there will be no future catches.

The ideal catch quota to set for the current year, C_T^*, maximises the expected long term time discounted yield from the fishery, assuming optimal choices are made in the future:

$$Y_T^* = \max_{C_T} \left[C_T + \mathbf{E}\left(\sum_{t=T+1}^{\infty} C_t^* \delta^t \right) \right], \quad 0 < \delta < 1. \tag{3}$$

Here δ is the time discounting factor, which represents the fact that future earnings are worth less than present earnings; generally this will be between 0.9 and 1.[1]

3   Single population stochastic dynamic program framework

We now outline an SDP framework similar to that by Walters [16], which can be used to solve the management problem we have introduced; we refer to

---

[1]See the work by Clark [4] for more details regarding time discounting.

this framework as the Single Population Stochastic Dynamic Programming (SP-SDP) framework.

At the present time, t, a fisheries manager wants to set a catch quota that maximises Equation (3). This optimal catch will depend on the state of the system which we summarise in a state vector $\boldsymbol{x}_t$. The state is uniquely defined by the population density and probabilities associated with the n possible models, such that $\boldsymbol{x}_t = [N_t, \sigma_t(1), \ldots, \sigma_t(n-1)]'$ (the nth probability is redundant as they sum to unity). SDP is based on the observation that Equation (3) can be written recursively:

$$Y^*(\boldsymbol{x}_t) = \max_{C_t} \left\{ C_t + \delta \mathbf{E}\left[Y^*(\boldsymbol{x}_{t+1})\right] \right\}, \tag{4}$$

where the expectation is over both the possible models and process noise.

The future state of the system $(\boldsymbol{x}_{t+1})$ depends on the current system state, the population dynamics and the decision made by the manager. Importantly, the manager's decision affects both the future population size and the information gained about the population dynamics (as some population sizes are more informative than others). Given their uncertain knowledge of the population dynamics, the manager needs to consider the outcome of a management decision under each possible model, and to weight these outcomes according to each model's probability:

$$Y^*(\boldsymbol{x}_t) = \max_{C_t} \left\{ C_t + \delta \sum_{i=1}^{n} \sigma_t(i) \mathbf{E}\left[Y^*(\boldsymbol{x}_{t+1} \mid i)\right] \right\}, \tag{5}$$

where the expectation is now only over process noise. $Y_{t+1}^*(\boldsymbol{x}_{t+1} \mid i)$ is the time discounted yield expected if optimal management decisions are made by the manager (given their uncertainty about the population dynamics) and model i is the correct model. Note that although we are conditioning on the true model being model i, this knowledge remains 'hidden' from the manager; it simply dictates how the population density will change and consequently what will be learnt about the population dynamics.

The best decision a manager can make, given their uncertainty, is that which maximises C_t^* in Equation (5). However, the actual value of this decision depends on which model is correct and is therefore the previously introduced $Y_t^*(\boldsymbol{x}_t \mid i)$. This can also be expressed recursively as

$$Y^*(\boldsymbol{x}_t \mid i) = C_t^* + \delta \mathbf{E}\left[Y^*(\boldsymbol{x}_{t+1} \mid i)\right] . \tag{6}$$

A SDP solution to this problem is found by beginning with arbitrary values $Y^*(\boldsymbol{x}_\tau \mid i)$ at some future time, τ, and working backwards. For each backwards step an optimal catch is found by maximising Equation (5) and the corresponding value of the optimal catch $Y^*(\boldsymbol{x}_t \mid i)$ is calculated from Equation (6) for each of the n possible models. After repeated backwards steps the optimal solutions and their corresponding values generally become independent of time, thus providing a solution to the problem.

In order to implement this procedure we require three further components which we now outline in some detail. Firstly, Equations (5) and (6) need to be solved for each possible state; as the state space is continuous it is necessary to discretise these in a sensible manner. The state vector $\boldsymbol{x}_t$ contains the population density and the probability associated with each model, $\boldsymbol{x}_t = [N_t, \sigma_t(1), \ldots, \sigma_t(n-1)]'$, and is thus n-dimensional as one probability is superfluous. Each of the dimensions in $\boldsymbol{x}_t$ is allowed to take on a discrete number of values; for the population density these range from 0 to the maximum attainable, and for the probabilities these range from 0 to 1. Denoting the number of possible values for the population density and probabilities by d_N and d_σ respectively, we note that there are $D = d_N d_\sigma^{n-1}$ possible states of the system.

Secondly, some arbitrary terminal values are required: a sensible choice will result in faster convergence to the solution. We assign the value of the remaining stock to each state $(Y^*(\boldsymbol{x}_\tau \mid i) = N_\tau)$; in effect this is the value of each state if the remaining fish were harvested.

Lastly, we need to calculate the probability of the system transitioning

from one state to the next. These probabilities depend on the current population density, the catch taken, the true model and process noise. In our model the catch is taken before the other processes take place, and hence we need only obtain probabilities for each post-catch population density and model. The system generally will not transition nicely from one discretised state at time t to another at $t+1$. We therefore consider a discrete set of values from the process noise distribution (we denote the number of possible values by d_ϵ) and interpolate values of $Y^*(\boldsymbol{x}_{t+1} \mid i)$ for the state resulting from each process noise value.

4   Adding a marine protected area

We now consider the situation where a marine protected area (MPA; an area in which no fishing is allowed) is introduced to permit the population dynamics to be observed at higher densities. The motivation for this is that an MPA may be easier to implement and that a higher population density increase may be obtainable in an MPA than in the whole population for the same loss in yield. To make this situation comparable to one without an MPA, we assume that the data from which the population dynamics are updated come from a fisheries-independent survey; during the existence of the MPA this survey monitors only the MPA population.

The population density inside and outside of the MPA must be modelled separately and migration between the two needs to be taken into account. The population dynamic models in Equation (1) are assumed to apply equally to both populations (the catch term is omitted for the protected population), but migration between the two populations takes place before the population densities are incremented by the model. This migration step is crucial: without it, the same increase in population density would be observed in an MPA regardless of its size.

We denote the population densities in and outside of the reserve prior to

migration by $\hat{R}_t$ and $\hat{N}_t$, and subsequent to migration by R_t and N_t. The form of migration derived in Appendix A for an MPA occupying a proportion, α, of the total area is

$$
\begin{aligned}
R_t &= \alpha(\hat{R}_t + \hat{N}_t) + ((1-\alpha)\hat{R}_t - \alpha\hat{N}_t)\, e^{-\eta B/(\alpha(1-\alpha))}\,, \\
N_t &= (1-\alpha)(\hat{R}_t + \hat{N}_t) + (\alpha\hat{N}_t - (1-\alpha)\hat{R}_t)\, e^{-\eta B/(\alpha(1-\alpha))}\,,
\end{aligned}
$$

where B is the length of the boundary of the MPA and η is a migration coefficient that gives the rate of movement over the boundary. The strength of migration will depend on many factors including the layout and size of the reserve and the species of fish being considered.

To evaluate the economic benefit of implementing an MPA, we consider a scenario where the MPA exists for some period of time, L, before it is reopened to fishing. During the existence of the MPA higher population densities are observed within the MPA, and consequently information about the population dynamics is gained more rapidly than is possible in the main fishery. The increased rate of learning comes at a cost as the MPA will result in lower catch rates[2]. However, this may be balanced by the improved management strategy that this additional information provides. Whether this is sufficient to make creation of the MPA economically optimal is the question we investigate.

Temporary MPAs, as considered here, will provide little additional benefit once a good understanding of the population dynamics has been attained. However, we do not suggest that real MPAs should be implemented temporarily, as they provide many additional benefits not considered here, which may ensure their long-term benefit, economic or otherwise.

---

[2]It has been argued that in some situations MPAs may result in an increased yield [12, 7, 8].

5   The full stochastic dynamic programming framework

The previously described SP-SDP framework applies after the MPA is reopened to fishing. SDP can also be used during the MPA period, although the states and transitions will be different, as there are two populations with migration. Here we present an SDP framework that gives the value of the system at the start of the MPA period, which we refer to as the MPA-SDP framework. By comparing the value of the system both with and without an MPA, it is possible to determine the economic benefit of implementing an MPA.

The state of a system with an MPA includes an additional dimension—the population density inside the MPA. The new state vector is $\boldsymbol{y}_t$, which now consists of $n + 1$ elements. At the start of the MPA period the population density is homogenous, and hence a straightforward mapping exists from the states $\boldsymbol{x}_0$ to $\boldsymbol{y}_0$: the population densities inside and outside of the MPA are simply equal to the original population density. At the end of the reserve period (time L) the reserve is re-opened to fishing. We assume a homogenous mixing of fish; this assumption is justifiable as economic theory suggests that the re-opened MPA would be selectively targeted until the population densities were again equal. Under this assumption the overall population density of the fishery is simply $\alpha R_L + (1 - \alpha)N_L$.

Since SDP works backwards, the period of time without an MPA must be considered first $(t > L)$. During this period the SP-SDP framework applies and can be used to obtain values for $Y_L^*(\boldsymbol{x}_L \mid i)$. Using the relationship between $\boldsymbol{x}_L$ and $\boldsymbol{y}_L$ just discussed, we can therefore interpolate the values of $Y_L^*(\boldsymbol{y}_L \mid i)$.

Using $Y_L^*(\boldsymbol{y}_L \mid i)$ and updated transitional probabilities for the two-population model, SDP can be applied for the period during which the MPA exists. The terminal values are now meaningful. Hence the procedure should not be repeated until convergence is reached; this corresponds to a perma-

nent MPA. Instead, L backwards steps are performed—this corresponds to the duration of the MPA. The values $(Y^*(\boldsymbol{y}_0 \mid i))$ obtained can be compared with those obtained from the SP-SDP framework to evaluate the economic optimality of implementing an MPA of a particular size and duration.

6   Simple application

We now consider a simple example where there are two possible models for the fishery. The two models we consider are a Ricker and a Beverton–Holt model (described in further detail in Appendix B). Both are common models for population dynamics and are derived from different biological principles.

The parameters for the two models have been selected so that the models provide similar dynamics at a population size between the optimal population size corresponding to each model, as illustrated in Figure 1. This figure also shows the expected yield as a function of the target stock size (the stock remaining after the catch is taken each year). A simple management strategy would select a catch each year so that a population density corresponding to the peak of the expected yield curve (around 0.55) is maintained. The trouble with such a strategy is that data obtained in the future will do little to differentiate between these two models as they are very similar at this population density. We expect that changing the population density to learn more about the population dynamics may be advantageous.

The MPA-SDP framework was applied to this example using a discretisation of $[d_N, d_\sigma, d_\epsilon] = [50, 40, 20]$; the resulting calculation took approximately 10 seconds on a current computer. Figure 2 shows the increase in value due to the introduction of an MPA occupying 5% of the fishery for five time steps. This increase is relative to the expected value if the manager does not take into account the effect that their choices have on learning. The increase in value is highest where initial uncertainty is high, as it is here that there is the most to be gained. The benefit of the MPA is reduced for high initial pop-
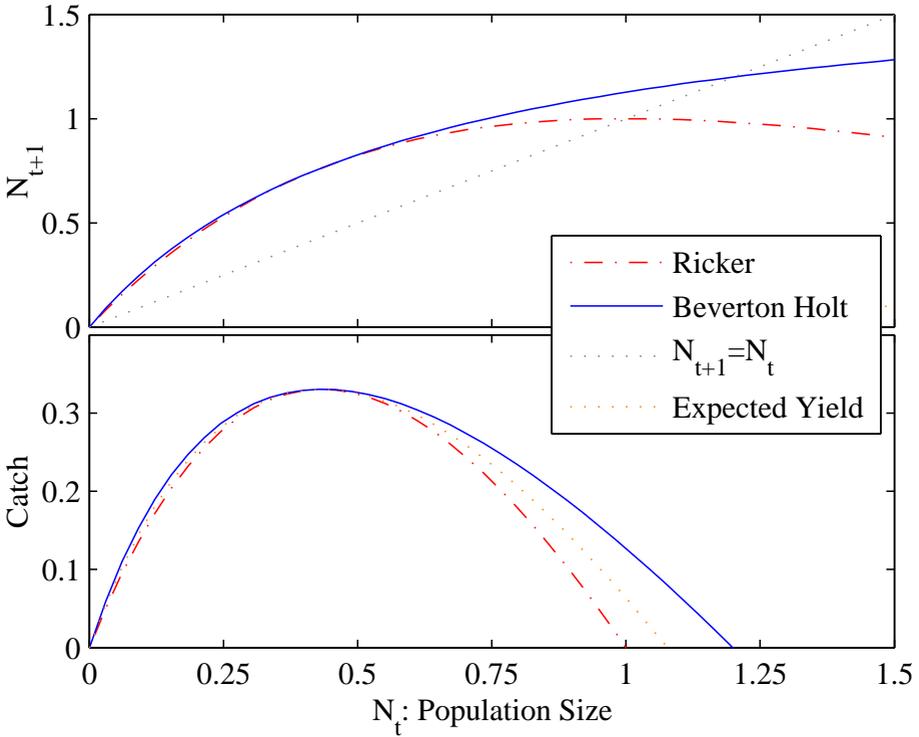
FIGURE 1: The simple scenario to which the MPA-SDP framework was applied. The two plausible models are a Ricker model (parameters $[r, K] = [1, 1]$) and a Beverton Holt model (parameters $[r, K] = [2, 1.9981]$). The black dotted line shows the expected yearly yield as a function of the target stock size if equal probability is placed on both models. A simple management strategy would set a target stock corresponding to the maximum of this curve (around 0.55). Data obtained from this region, however, will not help resolve the true model.
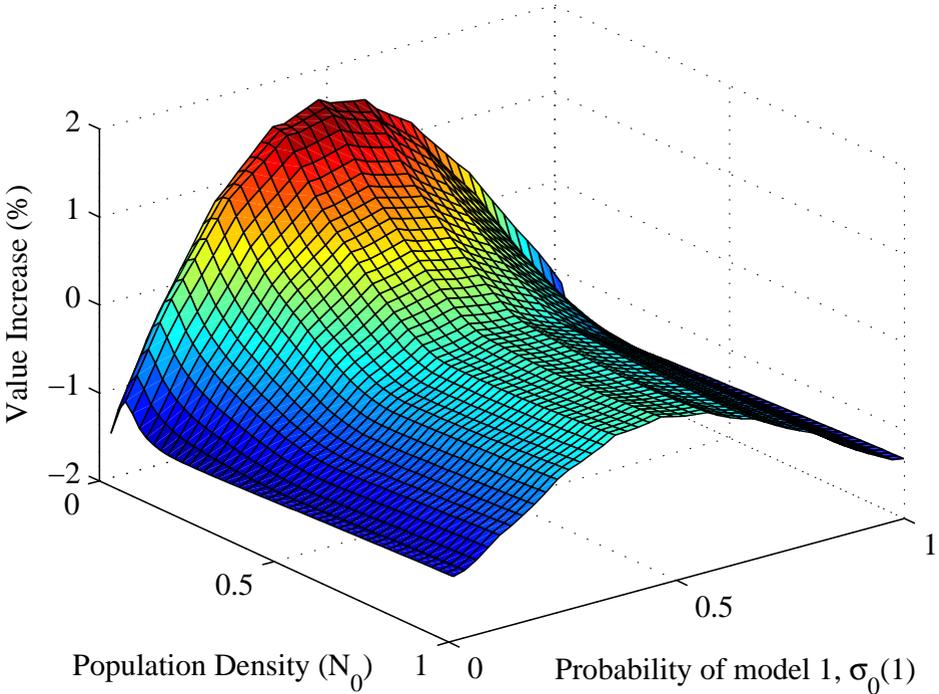
FIGURE 2: The increased value due to the introduction of an MPA occupying
5% of the fishery for five time steps with a migration coefficient of 0.01.
This increase is relative to the value obtained if a fishery manager does not
take into account the effect that their choice has on the information gained.
The highest increase (around 1.5%) is observed when the highest uncertainty
about the true model exists and the initial population density is low.

ulation densities, as it prevents the initial harvesting of a significant portion of the excess population.

Figure 3 shows the increase relative to the expected value if the manager is freely permitted to vary the population density in order to learn about the population dynamics. In this situation an MPA is significantly less beneficial. In order to obtain an increase in value it is necessary to have a very small migration factor, such that tiny MPAs yield high population increases. Note that the benefit of the MPA *increases* with the initial population density. Because of the small size of the MPA, its effect on the initial harvest of the excess is negligible and is outweighed by the rapid information gain occurring from the start, due to the elevated population density in the MPA.

The increases observed in these figures seem relatively small (the largest increase is about 1.5%). However, they must be considered relative to the maximum possible increase. The maximum possible increase from a simple management strategy is easily found for our application by considering Figure 1. A simple management strategy corresponds to the peak of the expected yield curve in the lower panel. If perfect information were to become available, the new expected yield would correspond to one of the model peaks with equal probability. The difference between these yields is the maximum increase attainable by gaining new information and is commonly referred to as the Expected Value of Perfect Information [17]; in our example the expected value of perfect information is an increase of only 8%.

7   Concluding comments

We have presented a framework for calculating the economic value of creating a reserve for the purpose of learning about the population dynamics of a fishery. Using this framework it is possible to illustrate a new and previously unconsidered economic value of MPAs.
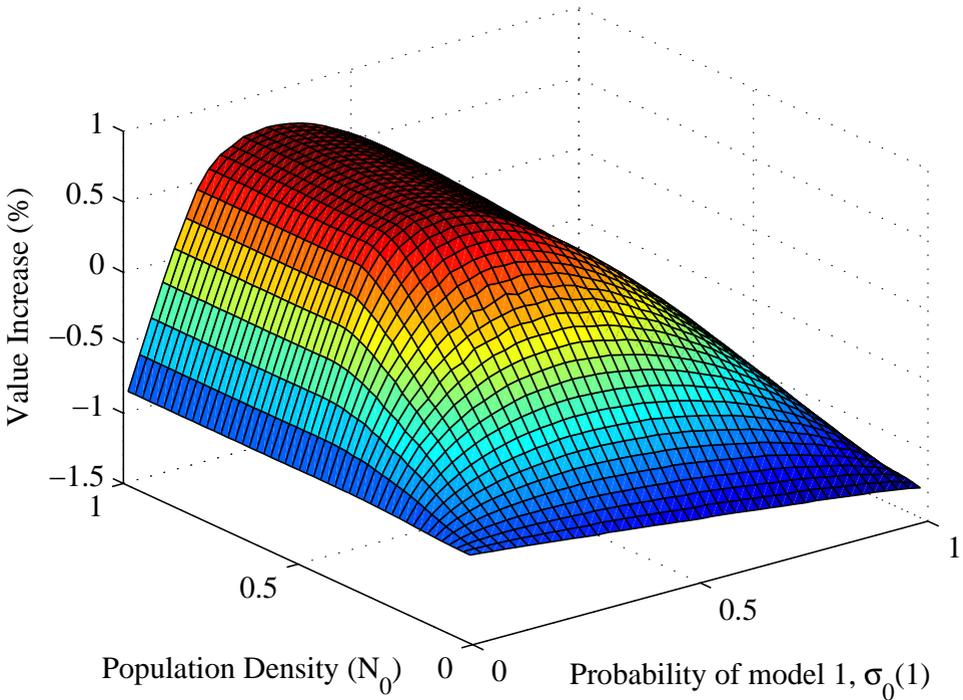
FIGURE 3: The increased value due to the introduction of an MPA occupying 1% of the fishery for ten time steps with a migration coefficient of 0.0001. This increase is relative to the value obtained if a fishery manager is permitted to manipulate the population density to gain informative data. Note that for clarity the population density axis has been reversed relative to Figure 2. The highest increase (around 0.75%) is observed when the highest uncertainty about the true model exists and the initial population density is high.

This framework poses challenging computational issues, particularly as the number of possible models increases. If the uncertainty occurs in a parameter of a model such that it too must be discretised, this may be particularly problematic. Careful consideration needs to be given to the discretisation of the parameters describing the state of the system; this is crucial for retaining acceptable computational performance.

There are several important simplifications inherent in the scenario we have outlined. One of the most important is that the population densities have been considered to be perfectly observable, but clearly this is not the case. Unfortunately, calculating the probabilities assigned to various models becomes complicated if both process noise and observational error are taken into account; developing an SDP framework that does this would be a formidable challenge. Furthermore, the population in the MPA may be more difficult to observe reliably due its smaller size and migration (an additional source of process noise which needs to be factored out).

The other major concern is the way in which the population dynamics are split for the two areas in the MPA-SDP. In reality, the population dynamics in any two areas will differ to some extent, and hence extrapolating results from the MPA to the whole fishery may result in incorrect management, depending on the variability of the population dynamics. This is particularly problematic as real MPAs are often selected for special attributes and may not be representative of the whole fishery.

Despite these limitations our model illustrates that it can be economically optimal to create a temporary MPA, solely for the purpose of learning about population dynamics. The economic benefit of creating an MPA is particularly pronounced if the species have a low migration factor or if active adaptive management (manipulating the population density to obtain informative data) is not possible. This provides an additional economic benefit of MPAs which on its own is sufficient to make an MPA economically optimal.

A   Migration

Migration can be considered a result of random movement, acting as diffusion, which will result in net movement from high to low population densities. Walters [15] implemented a continuous numerical scheme for modelling migration; a much simpler estimate will be presented here. Migration is considered over a period where the other processes that affect the population density are negligible. This may be the case, for example, with a species for which the juvenile fish relocate to find a new habitat over a short period of time and there is little adult movement.

In a small period of time, only individuals in the proximity of the boundary between the reserve and the exploited population are likely to cross the boundary. The population dynamics due to migration during the migratory period in the reserve are

$$dR = \left(\frac{wB\gamma}{1 - \alpha} N - \frac{wB\gamma}{\alpha} R \right) dt \,,$$

where for convenience R and N are the population sizes (not densities), respectively inside and outside the MPA. The first term gives the migration into the MPA and the second term the migration out of it. B is the length of the boundary and w is the distance from the boundary within which fish have to be in order to cross it. Effectively, therefore, $wBN/(1 - \alpha)$ is the number of fish that could cross the boundary; this is multiplied by γ—the rate at which an individual fish crosses the boundary—yielding migration into the MPA. The second term represents migration out of the MPA and has a similar interpretation. Making use of T as the total population size, this simplifies to

$$dR = \frac{\eta B}{\alpha(1 - \alpha)}(T\alpha - R) \, dt \,,$$

where $\eta = w\gamma$, a quantity we will refer to as the migration coefficient. With the assumption that the remaining population dynamics, birth and death,

are negligible during the period of migration, T will be constant and the above is easily solved to yield the population sizes after a migratory step of length one:

$$\begin{aligned} R &= \alpha T + (\hat{R} - \alpha T)\, e^{-\eta B/(\alpha(1-\alpha))}\,, \\ N &= (1-\alpha)T + (\hat{N} - (1-\alpha)T)\, e^{-\eta B/(\alpha(1-\alpha))}\,. \end{aligned}$$

Here $\hat{R}$ and $\hat{N}$ are the population sizes at the beginning of the migratory step and R and N are the population sizes at the end of this process. It is easy to show that for this model the MPA shape that minimizes migration is a circle; this gives the lowest boundary to area ratio. For a circular MPA of radius r the boundary length is $B = 2\pi r$. If the length unit is set such that the total habitat has area 1 we obtain $r = \sqrt{\alpha/\pi}$ and thus $B = 2\sqrt{\alpha\pi}$.

The migration coefficient η needs to be based on biological estimates of w and γ. Alternatively, if an MPA of size $\hat{\alpha}$ is believed to show negligible increase in density, then a bound on η is

$$\eta \geq \frac{\log(1/\epsilon)\sqrt{\hat{\alpha}}(1-\hat{\alpha})}{2\sqrt{\pi}}\,. \tag{7}$$

This ensures that for reserves less than $\hat{\alpha}$ in size or greater than $1 - \hat{\alpha}$, the factor $e^{-\eta B/(\alpha(1-\alpha))}$ will be less than ϵ.

This is a fairly simple approximation of the migration process. However, it exhibits all the qualitative features that are expected, and as such it is appropriate for the investigation conducted here.

B   Fishery models

The two common population models utilised in our example are derived from different ecological considerations [3, 13]. They both have two parameters:

one parameter, K, is the carrying capacity—the population density in the absence of fishing—and a parameter r, which is the growth rate at low densities.

The first of these is the Ricker model:

$$\mathbf{E}(N_{t+1}) = N_t e^{r(1-N_t/K)}.$$

The second model is a Beverton–Holt model:

$$\mathbf{E}(N_{t+1}) = \frac{r N_t}{1 + (r-1)N_t/K}.$$

The Ricker model displays over compensatory density dependence, in that for large enough population sizes an increase in population size in one year results in a decrease in total recruitment the following year (note the eventual down-turn in the dot-dashed curve in the upper panel of Figure 1). In contrast, the Beverton–Holt model displays compensatory density dependence where an increase in population size in one year results in a smaller increase in recruitment the following year. The two models lead to different predictions about maximum sustainable yield and have differing consequences for population dynamics.

The two models are derived under different assumptions about competition and predation. For example, over compensatory dynamics can arise if adults that recruited in one year cannibalise young that recruit the next year. Alternatively, the over compensatory Ricker model can be derived under an assumption of 'scramble' competition among recruits. Here, many individuals share a critical, limiting resource such as food. The more individuals that are present, the less resource each obtains. If individuals do not receive enough resource, they face increased mortality. The end result of which is that a large recruitment pulse may be counter productive and could lead to large scale recruitment failure. In contrast, compensatory dynamics would occur if the competition were of a 'contest' nature. In contest competition

a limited number of individuals are able to secure all of the resource they require, but additional individuals face acute resource scarcity, which gives the asymptotic behaviour shown in the solid curve in the upper panel of Figure 1.

Acknowledgement: We thank the two anonymous referees for their comments.

References

[1] Gary W. Allison, Steven D. Gaines, Jane Lubchenco, and Hugh Possingham. Ensuring persistence of marine reserves: Catastrophes require adopting an insurance factor. *Ecological Applications*, 13(1) Supplement:S8–S24, 2003. doi:10.1890/1051-0761(2003)013[0008:EPOMRC]2.0.CO;2. C310

[2] Paul R. Armsworth and Joan E. Roughgarden. The economic value of ecological stability. *PNAS*, 100:7147–7151, 2003. doi:10.1073/pnas.0832226100. C310

[3] Michael Begon, Colin Townsend, and John L. Harper. *Ecology: From Individuals to Ecosystems*. Blackwell Publishing, 2005. C324

[4] Colin W. Clark. *Mathematical Bioeconomics*. Wiley–Interscience, 1990. C311

[5] Fisheries and Agriculture Organization of the United Nations. *The State of the World Fisheries and Aquaculture*. FAO, Rome, 2004. ftp://ftp.fao.org/docrep/fao/007/y5600e/y5600e00.pdf. C308

[6] S. Garcia and C. Newton. In E. Pikitch, D. D. Hubert, and M. Sissenwine, editors, *Global Trends in Fisheries Management*, pages 3–27. American Fisheries Society, 1997. C308

[7] Fiona R. Gell and Callum M. Roberts. Benefits beyond boundaries: The fishery effects of marine reserves. *TRENDS in Ecology and Evolution*, 18:448–455, 2003. doi:10.1016/S0169-5347(03)00189-7. C315

[8] Fiona R. Gell and Callum M. Roberts. The fishery effects of marine reserves and fishery closures. Technical report, WWF, 2003. http://www.worldwildlife.org/oceans/pdfs/fishery_effects.pdf. C310, C315

[9] P. A. Larkin. Fisheries management – an essay for ecologists. *Ann. Rev. Ecol. Syst.*, 9:57–73, 1978. C308

[10] Jane Lubchenco, Stephen R. Palumbi, Steven D. Gaines, and Sandy Andelman. Plugging a hole in the ocean: The emerging science of marine reserves. *Ecological Applications*, 13(1) Supplement:S3–S7, 2003. doi:10.1890/1051-0761(2003)013[0003:PAHITO]2.0.CO;2. C309

[11] Renate Meyer and Russell B. Millar. Bayesian stock assessment using a state-space implementation of the delay difference model. *Can. J. Fish. Aquat. Sci.*, 56:37–52, 1999. doi:10.1139/cjfas-56-1-37. C310

[12] Michael G. Neubert. Marine reserves and optimal harvesting. *Ecology Letters*, 6:843–849, 2003. doi:10.1046/j.1461-0248.2003.00493.x. C310, C315

[13] Terrance J. Quinn and Richard B. Deriso. *Quantitative Fish Dynamics*. Oxford University Press, 1999. C324

[14] Callum M. Roberts and Julie P. Hawkins. Fully-protected marine reserves: A guide. Technical report, WWF, 2000. http://assets.panda.org/downloads/marinereservescolor.pdf. C310

[15] Carl Walters, V. Christensen, and D. Pauly. Structuring dynamic models of exploited ecosystems from trophic mass-balance assessments. *Reviews in Fish Biology and Fisheries*, 7:139–172, 1997. C323

[16] Carl J. Walters. Optimum escapements in the face of alternative recruitment hypotheses. *Can. J. Fish. Aquat. Sci*, 38:678–689, 1981. C309, C311

[17] Carl J. Walters. *Adaptive Management of Renewable Resources*. 1986. C309, C320

Author addresses

1. **Klaas Hartmann**, James Cook University, Townsville, Australia. Currently at University of Canterbury, Christchurch, New Zealand. `mailto:k.hartmann@math.canterbury.ac.nz`

2. **Lance Bode**, James Cook University, Townsville, Australia. `mailto:lance.bode@jcu.edu.au`

3. **Paul Armsworth**, The University of Sheffield, Sheffield, United Kingdom. `mailto:p.armsworth@sheffield.ac.uk`