CCAMLR Science, Vol. 20 (2013): 119–137

A POTENTIAL FEEDBACK APPROACH TO ECOSYSTEM-BASED MANAGEMENT: MODEL PREDICTIVE CONTROL OF THE ANTARCTIC KRILL FISHERY

S.L. Hill✉
British Antarctic Survey
Natural Environment Research Council
High Cross, Madingley Road
Cambridge, CB3 0ET
United Kingdom
Email – sih@bas.ac.uk

M. Cannon
Department of Engineering Science
University of Oxford
Parks Road
Oxford, OX1 3PJ
United Kingdom

Abstract

CCAMLR aims to develop a feedback approach to aid ecosystem-based management (EBM) of Antarctic krill fisheries. It is important to assess whether a feedback approach is likely to achieve the multiple objectives that EBM implies in the complex and uncertain conditions typical of Antarctic marine ecosystems. This study used Model Predictive Control (MPC) to achieve objectives for a harvested species, its predators and the fishery, in a simulation model that incorporates uncertainty and spatial and trophic complexity. The approach adjusted spatially resolved annual catch limits in response to estimates of the state of the system. It suggests that feedback management is both feasible and a more effective way to achieve multiple objectives than fixed catch limits, which are currently used to manage Antarctic krill fisheries. The study demonstrates that optimisation-based approaches such as MPC are computationally capable of dealing with EBM-type problems. They are also useful for assessing the feasibility of candidate management policies or objectives, and characterising the trade-offs that they imply. This study characterises the trade-off between catch levels and the risk of harvested species biomass falling to unacceptable levels.

Introduction

CCAMLR aims to develop a feedback management approach for the Antarctic krill fishery in the Scotia Sea and southern Drake Passage (FAO Statistical Subareas 48.1 to 48.3) (SC-CAMLR, 2011), which has a potential catch limit equivalent to 7% of current global fisheries landings (Hill, 2013a). CCAMLR is committed to managing the fishery according to conservation principles that articulate the key goals of ecosystem-based management (EBM) (Grant et al., 2013). Feedback management must therefore overcome important challenges, including: (i) the need to meet multiple objectives for krill, its predators and the fishery (Grant et al., 2013); (ii) the complex spatial structure and significant spatial scale of the area requiring management (Murphy et al., 2012); (iii) the high levels of spatial and temporal variability in the ecosystem and substantial uncertainty in current understanding of its structure and operation (Nicol and Siegel, 2000; Hill et al., 2006; Hill, 2013a); and (iv) ongoing ecosystem change, which has multiple putative drivers, including past and current harvesting and climate change (Murphy et al., 2012).

CCAMLR's Working Group on Ecosystem Monitoring and Management (WG-EMM) defined feedback management as an approach that 'will use decision rules to adjust selected activities (distribution and level of krill catch and/or research) in response to the state of monitored indicators' (SC-CAMLR, 2011). The harvest strategies and decision rules used in modern fisheries management derive from control theory (Cunningham and

Maguire, 2002), which is concerned with controlling the states of dynamic systems. WG-EMM's broad definition of feedback management encompasses various control strategies, including both adaptive and robust control. Adaptive control aims to reduce uncertainty by estimating both the state of the system and the parameters that describe its dynamics (Bitmead et al., 1990). It is analogous to active adaptive management of human activities affecting living resources (Walters, 1986). Robust control aims to achieve specified objectives for the state of systems with characteristics such as uncertain dynamics, unmeasurable states and significant observation error (Zhou and Doyle, 1997; Rawlings and Mayne, 2009).

The management procedure approach used by fisheries scientists illustrates many of the features of control strategies (Figure 1). According to Rademeyer et al. (2007), a management procedure (MP) consists of data that provides information on resource status and productivity, and a set of harvest control rules (HCRs) that adjusts catch or effort controls in response to this information (see also Kell et al., 2006). MPs that are implemented to control fisheries in the real world are known as operational management procedures, whereas management procedure evaluation (MPE) is a simulation method that assesses the likely performance of MPs by using them to control operating models representing the controlled system and monitoring process. In both cases there is a controlled system (the harvested stock or ecosystem, or the operating model), a process for estimating its state (direct empirical measurements, which might be used with an assessment model that integrates available data), and a controller (the HCRs) which attempts to align the state of the system with specified objectives by adjusting catch or effort.

MPE generally attempts to find an HCR which will make adjustments without needing to be refined itself. In control theory it is also possible to continuously refine the control law (which is analogous to an HCR in this context) in response to the current state of the system. This refinement occurs within the feedback loop and is a characteristic of the control method known as Model Predictive Control (MPC) (Figure 2). Recent advances in algorithms and computing resources mean that it is now feasible, within MPC, to solve previously intractable robust control problems featuring high levels of complexity and uncertainty (Mayne et al., 2000; Grüne and Pannek, 2011). Complexity and uncertainty are defining characteristics of EBM (Hill et al., 2007; Link et al., 2012) and of the ecosystem in which the Antarctic krill fishery operates (Hill et al., 2006; Murphy et al., 2012; Hill, 2013a). There might therefore be some advantage in considering this type of continuously updating HCR.

This study uses MPC to develop an HCR designed to achieve multiple objectives within an illustrative operating model. The operating model incorporates important features of the ecosystem in which the Antarctic krill fishery operates. These include trophic and spatial structure, stochastic temporal variability, harvesting and movement of the harvested species between areas. The management objectives concern catch stability and the biomass of the harvested species and its predators. The state estimates used to derive the HCR include estimation error.

The aims of this study are: to assess whether a feedback approach can feasibly achieve the multiple objectives that EBM implies in the complex and uncertain conditions typical of Antarctic marine ecosystems; to compare the performance of a feedback approach with a fixed catch limit approach such as that currently used to manage Antarctic krill fisheries; to illustrate the elements of feedback control strategies, and the issues that must be addressed in their development; and to characterise some of the major trade-offs between objectives. There are radical differences between ecosystems and man-made systems in terms of objectives, uncertainty and complexity. Acknowledging these differences, this contribution discusses the ways in which control theoretic approaches, developed initially for man-made systems, can aid the development of management approaches for ecosystems and, in particular, how control theory can contribute further to the management procedure approach.

Methods

Overview

The current study develops an illustrative MPC strategy, and uses a simulation approach to assess its capabilities and evaluate trade-offs between objectives. Figure 2 illustrates the elements of an MPC strategy. The current study mainly addresses the objective function and the control law, using illustrative representations of the remaining elements to provide the necessary context. The control strategy is applied to a fished ecosystem which, in this study, is represented by an operating model. This model has three areas that are connected by transport of the harvested species between them. The fishery and a predator species target the harvested species in two of these areas. The area-specific carrying capacity for these two areas, and the influx of the harvested species from the third area, varies stochastically over time. This is a bottom-up model in which harvested species biomass affects predator biomass, but not vice versa. The operating model and all processes within the feedback loop have a nominal time-step of one year.

The controlled system (here the operating model) provides output in the form of measurable state variables such as biomass and, trivially, catch. An estimator uses information from the limited suite of measurable variables to infer the state of the wider suite of relevant state variables. This study does not explicitly model an estimator but it represents uncertainties due to observation and inference by adding error to the state estimates that are passed to the controller.

EBM-like objectives specify the reference, or desired, state of the controlled system. Some of these objectives specify limit reference points (LRPs) (Caddy and Mahon, 1995) for the predator and target species. These define the boundary between biomass levels that are considered desirable and those that are considered undesirable. These objectives are, in fact, soft limit reference points (SLRPs) (Hill, 2013b) because they also specify the probability with which biomass must remain in the desirable range. The study also considers another important fisheries objective, which is limiting interannual catch variability.

If the stochastic temporal variability in the operating model was set to zero, the model would converge towards steady-state conditions. The control strategy could be characterised as an attempt to bring the system as close as possible to the hypothetical steady-state conditions which achieve the specified objectives, subject to additional considerations such as the requirement to minimise catch variability. These steady-state conditions can be used to identify a set of target reference points (TRPs – states that management objectives are focused on achieving) (Caddy and Mahon, 1995) based on trade-offs between catch levels and the risk of biomasses falling below SLRPs.

The controller includes a prediction model, here a linear approximation of the operating model, parameterised with the uncertain state estimates. Its main function is to predict the future state of the system when supplied with a sequence of annual area-specific catches.

At each time-step, the controller identifies the sequence of future catches most likely to achieve the objectives, given the available uncertain estimates of current state. It identifies this optimal sequence using an objective function which minimises the predicted deviation of multiple state variables (provided by the prediction model), in multiple future years, from their TRPs. It also penalises interannual variability in biomass or catch. The importance of any particular state variable or penalty in this minimisation process can be controlled through weighting. The area-specific catches for the first year in the optimal sequence define the revised catch limits in the operating model. The sequence is revised at each time-step in response to new state estimates.

The following sections provide more detail. Tables 1 and 2 and Equation 8 specify the quantitative inputs used in this study: the parameters used in the operating model and the controller, and the weights used in the objective function. The 'State estimation' section also assesses which variables it is necessary to measure in order to infer the state of the operating model.

Operating model

The controlled system in this study is referred to as an operating model to distinguish it from the prediction and uncertainty models that are also components of MPC. However, it does not strictly conform to the definition of an operating model in Rademeyer et al. (2007) because the evaluated control strategy is illustrative.

The operating model is based on the simplified formulation of Constable's (2001) multi-species model in Hill et al. (2006). It is modified to include spatial structure. For simplicity, the linear fishing mortality term is removed from the equations of Constable (2001), and harvesting and movement are represented with additive loss or gain terms.

Table 1: Parameters used to implement the operating model.

Parameter	Area 0	Area 1	Area 2
Mortality rate (harvested species)		$M_1 = 0.2$	$M_2 = 0.2$
Mortality rate (predators)		$M'_1 = 0.05$	$M'_2 = 0.05$
Intrinsic growth rate (harvested species)		$r_1 = 0.4$	$r_2 = 0.5$
Intrinsic growth rate (predators)		$r'_1 = 0.1$	$r'_2 = 0.1$
Carrying capacity (harvested species): mean		$\bar{K}_1 = 10^4$	$\bar{K}_2 = 10^4$
standard deviation		$\sigma_{K,1} = 0.27\times10^4$	$\sigma_{K,2} = 0.18\times10^4$
Stable biomass factor (predators)		$\gamma_1 = 10^4$	$\gamma_2 = 10^4$
Degree of density dependence (predators)		$\zeta_1 = 2$	$\zeta_2 = 2$
Initial biomass (harvested species)		$B_{1,0}= 6\,500$	$B_{2,0}= 6\,500$
Initial biomass (predators)		$P_{1,0} = 45$	$P_{2,0} = 45$
Biomass (harvested species): mean	$\bar{B}_0 = 10^4$		
standard deviation	$\sigma_{B_0} = 0.2\times10^4$		
Transport coefficients	$T_{10} = 0.1$	$T_{11} = 0.1$	$T_{22} = 0.1$
Spectral bandwidths use to generate autocorrelated temporal variability	$\alpha_0 = 0.8$	$\alpha_1 = 0.8$	$\alpha_2 = 0.8$

Table 2: Controller and constraint parameters used in the simulations.

Parameter	Value
MPC prediction horizon	$N = 20$
Input bounds	$U_{min} = [400\ \ 400]^T$ $U_{max} = [1\,000\ \ 1\,000]^T$ $\Delta U_{max} = [200\ \ 200]^T$
State bounds	$X_{min} = [1\,300\ \ 1\,300\ \ 10\ \ 10]^T$ $X_{max} = [10^4\ \ 10^4\ \ 100\ \ 100]^T$
Constraint probability	$p_X = 0.9$
Mean steady state harvested species biomass	$(B_{1,ss}, B_{2,ss}) = (5\,200, 4\,000)$
Mean steady state predator biomass	$(P_{1,ss}, P_{2,ss}) = (37.2, 28.6)$
Mean steady state harvest	$(H_{1,ss}, H_{2,ss}) = (594, 344)$
Linear feedback gain	$L = \begin{bmatrix} 0.73 & 0 & 2.80 & 0 & 0.09 & 0 & 0.08 & -0.89 & 0 \\ 0.06 & 0.81 & -0.09 & 2.71 & 0 & 0.07 & 0 & 0 & -0.90 \end{bmatrix}$
Asymptotic value in terms of the predicted cost	$l_{ss} = 1.26 \times 10^5$

The implemented model has three areas. The first area simply provides a source of the harvested species for the remaining two areas. The biomass dynamics of the harvested species are explicitly represented in these two areas:

$$\begin{bmatrix} B_{1,k+1} \\ B_{2,k+1} \end{bmatrix} = \begin{bmatrix} e^{-M_1} & 0 \\ 0 & e^{-M_2} \end{bmatrix} \begin{bmatrix} B_{1,k} \\ B_{2,k} \end{bmatrix} + \begin{bmatrix} r_1 B_{1,k}\left(1 - B_{1,k} / K_{1,k}\right) \\ r_2 B_{2,k}\left(1 - B_{2,k} / K_{2,k}\right) \end{bmatrix} + \begin{bmatrix} -T_{11} & 0 \\ T_{11} & -T_{22} \end{bmatrix} \begin{bmatrix} B_{1,k} \\ B_{2,k} \end{bmatrix} + \begin{bmatrix} T_{10} \\ 0 \end{bmatrix} B_{0,k} - \begin{bmatrix} H_{1,k} \\ H_{2,k} \end{bmatrix} \qquad (1)$$

where $B_{i,k}$ is biomass in area i at time k; M_i is the area-specific non-harvesting mortality rate; r_i is the area-specific intrinsic biomass growth rate; $K_{i,k}$ is the area-specific maximum biomass (carrying capacity) at time k; $H_{i,k}$ is the catch biomass from area i at time k. $T_{ij} \geq 0$ is the proportion of harvested species biomass in area $j \neq i$ that moves into area i in one time-step, whereas $T_{ii} \geq 0$ is the proportion of harvested species biomass in area i that moves out of area i in one time-step. Harvested species biomass at time k in the source area is a random variable drawn from a normal distribution with mean $\bar{B}_0$ and variance $\sigma^2_{B_0}$:

$$B_{0,k} \sim \mathcal{N}\left(\bar{B}_0, \sigma^2_{B_0}\right). \qquad (2)$$

The equation for the predator species is:

$$\begin{bmatrix} P_{1,k+1} \\ P_{2,k+1} \end{bmatrix} = \begin{bmatrix} e^{-M'_1} & 0 \\ 0 & e^{-M'_2} \end{bmatrix} \begin{bmatrix} P_{1,k} \\ P_{2.k} \end{bmatrix} + \begin{bmatrix} r'_1 P_{1,k}\left[1 - \gamma_1\left(P_{1,k} / B_{1,k}\right)^{\zeta_1}\right] \\ r'_2 P_{2,k}\left[1 - \gamma_2\left(P_{2,k} / B_{2,k}\right)^{\zeta_2}\right] \end{bmatrix} \qquad (3)$$

where $P_{i,k}$ is biomass in area i at time k; M'_i is the area-specific mortality rate; r'_i is the area-specific intrinsic biomass growth rate; γ_i is an area-specific scale factor in the relationship between harvested species biomass and the maximum stable predator biomass that it supports; and ζ_i is an area-specific shape parameter controlling the degree of density dependence in this relationship.

The operating model, therefore, describes logistic biomass growth for both the harvested species and its predators. Local carrying capacity constrains biomass growth in the harvested species, and prey availability constrains growth in the predator species. There are constant natural mortality terms for both the harvested species and its predators. Mortality for the harvested species is independent of predator biomass and the model is therefore a bottom-up model with no top-down effects. The catches from areas 1 and 2 are the control variables, i.e. those that the controller adjusts. Additional stochasticity is introduced into the model via the carrying capacities for the harvested species, which are modelled as normally distributed random variables in each area (Constable, 2001):

$$K_{i,k} \sim \mathcal{N}\left(\bar{K}_i, \sigma^2_{K_i}\right). \qquad (4)$$

The sequences of random variables (Equations 2 and 4) are autocorrelated (Figure 3). This was achieved by linearly filtering sequences of normally distributed uncorrelated random variables:

$$\begin{aligned} & B_{0,k+1} = \alpha_0 B_{0,k} + \left(1 - \alpha_0\right) B^u_{0,k}, \\ & B^u_{0,k} \sim \mathcal{N}\left(\bar{B}_0, \sigma^2{}_{B_0}\left(1 - \alpha_0^2\right) / \left(1 - \alpha_0\right)^2\right), \\ & k = 1, 2, \ldots \end{aligned} \qquad (5)$$

$$\begin{aligned} & K_{i,k+1} = \alpha_i K_{i,k} + \left(1 - \alpha_i\right) K^u_{i,k}, \\ & K^u_{i,k} \sim \mathcal{N}\left(\bar{K}_i, \sigma^2{}_{K_i}\left(1 - \alpha_i^2\right) / \left(1 - \alpha_i\right)^2\right), \\ & i = 1, 2; k = 1, 2, \ldots \end{aligned} \qquad (6)$$

where parameters α_0, α_1 and α_2 specify the bandwidths of the spectra of the sequences of state variables and therefore determine the degree of autocorrelation.

Table 1 gives the parameter values used to implement the model. These include area-specific differences in the biomass growth rate of the harvested species as might occur with spatial differences in water temperature and primary production (Atkinson et al., 2006). The movement parameters specify unidirectional flow of a proportion of the harvested species biomass from area 0 through area 1, then area 2 and finally out of the system. The mortality rates, growth rates and mean carrying capacities were taken from Constable (2001). The variances of random variables were chosen to be similar to the equivalent values in Constable (2001). However, because of the differences between the operating model and that of Constable (2001), these values had to be adjusted so that the chance of population collapse in any one area remains at a plausible level, which is necessarily non-zero due to the use of normally distributed parameters. For example, with the values in Table 1, the likelihood of the predator population falling to zero in either area 1 or 2 is less than 0.1% over an interval of 100 time-steps when there is no harvesting. The parameters controlling the degree of density dependence in predator populations were selected with reference to Constable (2001), and the values chosen for the scale factors ensure that the maximum stable predator biomass is around 1% of the corresponding harvested species biomass. Thus, the harvested species nominally represents Antarctic krill, and each modelled predator represents one of a number of predators of Antarctic krill in each area.

Objectives

CCAMLR's conservation principles and their interpretation by CCAMLR's Scientific Committee over the past three decades imply that management objectives are needed for the Antarctic krill stock, its predators and the fishery (Grant et al., 2013). CCAMLR has made progress in specifying objectives for the future state of the krill stock as represented in population projection models, but specific objectives for predators and the fishery remain to be defined (Hill, 2013b). This study therefore uses the following illustrative objectives:

(i) with a probability of at least 90%, the area-specific biomasses of harvested species must remain above 20% of their respective pre-exploitation levels,

(ii) with a probability of at least 90%, the area-specific biomasses of predators must remain above 60% of their respective pre-exploitation levels,

(iii) the interannual catch variability must be minimised, subject to the other objectives and constraints, and

(iv) the interannual variability in biomass must be minimised, subject to the other objectives and constraints.

Objective (i) is based on the decision rule for the krill stock suggested by Butterworth et al. (1991) and later included in the three-part decision rule used to set the precautionary catch limit (Constable et al., 2000). Objective (ii) is based on the depletion level used by Smith et al. (2011) to indicate a serious fisheries impact on a non-target trophic group. Objective (iii) acknowledges that catch stability is a key goal of fisheries management (Rademeyer et al., 2007; Rosenberg, 2009), and objective (iv) indicates the preference for relatively stable ecosystems implied by various statements of objectives (McLeod and Leslie, 2009), reference points (Smith et al., 2011) and policy (Penas, 2007).

Objectives (i) and (ii) are expressed in terms of LRPs, but they also specify the probability of the system entering an undesirable state. These reference points are therefore SLRPs, which, in the terminology of control theory, define probabilistic state constraints. In addition to these operational constraints, there are also the physical constraints that catch levels and carrying capacities must be non-negative.

Although the objectives do not specify any TRPs, they do imply such targets. These are the hypothetical steady-state conditions in which all other objectives and constraints are satisfied. There may be many potential TRPs for each state variable. Furthermore, the constant area-specific catch levels that would exist under these hypothetical steady-state conditions imply TRPs for catch. The choice of TRPs is a trade-off between objectives. For example, within the set of potential TRPs, higher TRPs for catch are likely to have lower counterpart TRPs for biomass. The TRPs can be selected using the feasibility analysis depicted in Figure 7 and discussed in 'Results'.

State estimation

In living resource management in general, and EBM in particular, accurate measurement of all, or indeed any, of the system's relevant state variables is unlikely to be practicable. The fields of stock assessment (e.g. Hilborn and Walters, 1992) and estimation theory (e.g. Luenberger, 1966; Kwakernaak and Sivan, 1972) have given considerable attention to the significant problem of estimating system state from the available data. For the purposes of the current study, it is assumed that noisy measurements of a limited set of state variables are available, from which to estimate the wider suite of relevant states.

The relatively simple operating model has seven state variables describing the area-specific biomass of the harvested species and its predators, and the carrying capacity for the harvested species ($B_{1,k}$, $B_{2,k}$, $P_{1,k}$, $P_{2,k}$, $K_{1,k}$, $K_{2,k}$ and $B_{0,k}$). The observability of a system is a measure of whether a limited set of measurable variables appropriately indicates the full suite of state variables, including unmeasurable states such as the time-specific carrying capacities (Kwakernaak and Sivan, 1972). Observability can therefore be used to identify which variables to measure. Assessment of the observability of the linearised dynamics of the operating model (Isidori, 1995) indicates that it is necessary to measure: (i) the predator biomass in areas 1 and 2; and (ii) the harvested species biomass in area 0. The remaining state variables can be inferred from measurements of these variables.

In the current study there was no explicit representation of this inference process and the controller was supplied with estimates of all seven state variables, which were adjusted to incorporate random error drawn from a lognormal distribution with a CV of 0.10. The additional error in estimates of $P_{1,k}$, $P_{2,k}$ and $B_{0,k}$ therefore represents observation error, while that in estimates of the remaining state variables represents uncertainties in the inference process. In addition to state variables, assessment and prediction models require estimates of other parameters to represent the system dynamics. The prediction model in the current study used the same biomass growth, mortality and transport parameters as the operating model.

Defining the optimal control problem

The optimal control problem is a mathematical formulation of the objectives for the controlled system. It is formulated as an objective function to be optimised subject to an accompanying set of constraints. As stated previously, the control strategy attempts to minimise deviations of the state (biomass and carrying capacity) and input (catch) variables from their respective TRPs, while the SLRPs define probabilistic constraints. The requirement to control the temporal variability in biomass can be addressed through minimisation of the expected value of a quadratic cost function (e.g. Kwakernaak and Sivan, 1972). This leads to the following definition of the optimal control problem:

$$\min_{\{u_0,u_1,u_2,K\}} \sum_{k=0}^{\infty} \mathbb{E}\left(\begin{array}{l} \left\| x_k - x^0 \right\|_Q^2 \\ + \left\| u_k - u^0 \right\|_R^2 + \left\| \Delta u_k \right\|_S^2 \end{array} \right)$$

subject to

$$U_{min} \le u_k \le U_{max},\ \Delta u_k \le \Delta U_{max},\ k = 0,1,2,K$$

$$\Pr\left(x_k \ge X_{min}\right) \ge p_X,\ k = 1,2,K \tag{7}$$

where $\mathbb{E}(.)$ denotes the expectation operator and $\|x\|_Q^2$ denotes the quadratic form $x^T Qx$. The vectors x_k and u_k respectively contain the area-specific state and input variables, and x^0 and u^0 are the vectors of their respective area-specific TRPs, while $\Delta u_k = u_k - u_{k-1}$ is the area-specific interannual catch variation.

The probabilistic constraint $\Pr\left(x_k \ge X_{min}\right) \ge p_X$ applies separately to each state variable, making it possible to specify the relevant SLRP using the LRP, X_{min}, and the probability, p_X. The constraints imposed by the limits on u_k and Δu_k specify an allowable range and interannual change in area-specific catch levels, thus imposing limits on catch variations and rates of change. Q, R and S are positive-definite matrices specifying the weights applied in the optimisation. Q and R specify the degree to which the deviations of states and inputs from their TRPs are penalised, while S penalises interannual catch variability. The specific matrices used in this development were:

$$Q = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 100 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 100 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$R = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix},\ S = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} \quad . \tag{8}$$

The choice of weights in Q reflects the relative biomasses of predators and the harvested species in the operating model; the weights in Q on the harvested species biomass in area 0 and the carrying capacities are set to zero since the controller cannot influence these states. The weights in R and S control catch variability by penalising its variance. Thus, the choice of weights Q, R and S determine the trade-off between objectives (iii) and (iv).

Controller design

The dynamics of the operating model are non-linear and stochastic to reflect the nature of real ecosystems. Consequently, identifying the exact optimal solution to the optimal control problem is likely to be computationally intractable. MPC can be used to approximate the problem and so obtain an approximate solution (Mayne et al., 2000; Couchman et al., 2006; Cannon et al., 2009). MPC is also appropriate to use with constraints of the type specified in the optimal control problem (Kouvaritakis et al., 2010; Cannon et al., 2011).

The MPC approach identifies an optimal sequence of input variables using a prediction model to project the future dynamics of the controlled system. The length of the sequence is known as the prediction horizon, N, and is usually chosen based on estimates of the response time of the system (defined as the time taken to return to within 1% of the steady state following a perturbation on the scale of those generated by the controller) when constraints are not active. The optimal input sequence at time k therefore depends on the current estimated state of the system. In the current context, MPC generates a sequence of annual area-specific catch levels, implements the first of these, and re-estimates subsequent catch levels in the next time-step.

The illustrative prediction model in the current study is a linear approximation of the operating model. This simplifies the optimisation process. It also avoids exact replication of the operating model in the prediction model, which reflects the impossibility of identifying an accurate model of ecosystem dynamics. In this study, the accuracy of the linear approximation was verified by simulation. If necessary, it is also possible to reduce the discrepancy between the dynamics of the operating model and its linear approximation by imposing additional constraints in the optimal control problem. MPC can also accommodate the more complex non-linear models that might be necessary to represent the dynamics of real ecosystems.

The operating model can be written in the general form:

$$x_{k+1} = f\left(x_k, u_k, w_k\right) \tag{9}$$

where the vectors of state, input and disturbance variables are denoted x_k, u_k and w_k respectively. The disturbance variables are the state variables which are subject to stochastic variability. If these three sets of variables consist of small perturbations (indicated by superscript δ) of the TRP values, x^0, u^0 and w^0, so that $x_k = x^0 + x_k^\delta$, $u_k = u^0 + u_k^\delta$ and $w_k = w^0 + w_k^\delta$, then the linear approximation to the system dynamics around x^0, u^0 and w^0 can be expressed:

$$x_{k+1}^\delta = Ax_k^\delta + Bu_k^\delta + Dw_k^\delta \tag{10}$$

$$A = \left.\frac{df}{dx}\right|_{\left(x^0,u^0,w^0\right)} \quad B = \left.\frac{df}{du}\right|_{\left(x^0,u^0,w^0\right)}$$
$$D = \left.\frac{df}{dw}\right|_{\left(x^0,u^0,w^0\right)} \quad . \tag{11}$$

The above notation refers to the value of the derivative under equilibrium conditions.

This system, computed with respect to equilibrium conditions implied by the TRPs compatible with the input parameters, provides the linear prediction model.

An optimisation method known as quadratic programming (QP) can be used to compute the optimal control law (sequence of area-specific catch limits). This requires the optimal control problem (Equation 7) to be expressed as a convex quadratic program which can be solved using standard QP algorithms (e.g. Mayne et al., 2000; Grune and Pannek, 2011) to give the feedback control law:

$$u_k = u^0 + L\left(x_k - x^0\right) + v_k,\ k = 0,1\ldots \tag{12}$$

Here L is the fixed feedback gain (i.e. how much output is fed back to the input, see Figure 2) with the property that, if v_k = 0 for all k, then Equation 12 solves the problem of minimising the objective function in Equation 7 when the constraints are removed. This unconstrained optimal control problem, which is known as the linear-quadratic regulation problem (Kwakernaak and Sivan, 1972), can be solved to determine L using standard computational tools (e.g. Matlab control toolbox). Thus v_k contains the variables over which the objective function is to be minimised, and the purpose of v_k is to optimally modify this linear feedback law to meet the specified constraints.

The variance of states and inputs in steady state is necessarily non-zero under any feedback control law, and this implies that the objective function defined in Equation 7 is necessarily infinite. Therefore, to allow the objective function to be optimised numerically, MPC was implemented by minimising a modified objective function:

$$\min_{\{u_0, u_1, u_2, K\}} \sum_{k=0}^{\infty} \mathbb{E} \begin{pmatrix} \left\| x_k - x^0 \right\|_Q^2 \\ + \left\| u_k - u^0 \right\|_R^2 + \left\| u_k^{'} \right\|_S^2 - l_{ss} \end{pmatrix} \quad (13)$$

where l_{ss} is the minimum achievable value of $\mathbb{E}\left(\left\| x_k - x^0 \right\|_Q^2 + \left\| u_k - u^0 \right\|_R^2 + \left\| u_k^{'} \right\|_S^2 \right)$ in steady state, which can be computed using the linearised dynamics around equilibrium. This modification ensures that the optimal value of the objective function is finite.

Simulations

The performance of the MPC controller was explored in 500 simulations of 60 time-steps, each implemented with a unique random sequence of values for the disturbance variables. For comparison, the mean steady-state catch levels from the MPC were used to implement a fixed harvest strategy. Also, the performance in terms of catch variability was contrasted with that of a linear feedback controller. This controller has an optimal solution to the objective function (Equation 7) when the constraints are removed (i.e. it is governed by the TRPs but not by the LRPs). These contrasting strategies were simulated using the same 500 sets of disturbance variables as the MPC strategy. Table 2 details the additional parameters used to implement these simulations, including solutions for the state values, gain L, and l_{ss}. Further solutions were obtained with different constraint probabilities for objective (i) in the range [0.5, 1] to explore the relationships between constraint probability and feasible values of state variables.

Results

CCAMLR currently manages the Antarctic krill fishery with a fixed catch limit. With sufficient information about the system, and the assumption that it has a constant underlying equilibrium, this catch limit could be set so that the average biomass for the harvested species in each area is equal to its TRP. Figure 4 shows the biomass trajectories from the operating model with such fixed catch limits in each area. This shows a wide dispersion of trajectories around the mean (target) state. Predator biomasses in areas 1 and 2 fell below the 60% LRP in 1% and 25% of simulations respectively. In contrast, application of the MPC strategy to simulations using the same sequences of random numbers results in a much narrower dispersion and only 0% and 3% respectively of trajectories falling below the 60% LRP (Figure 5). Overlaying 500 trajectories for the harvested species produces the effect of regular peaks and troughs. This is because each of the individual trajectories is regularly steered back towards the target state. The controller is better able to reduce biomass (by increasing the catch level) than it is able to increase biomass (by reducing the catch level), so the peaks appear sharper than the troughs.

The MPC strategy achieves tight control on the system by varying the catch levels in the two areas. Figure 6 compares the catch trajectories for MPC and for the linear feedback controller. MPC achieves lower catch variability than the linear feedback strategy (demonstrated by its lower variance) as a result of constraints on catch and catch variability. This lower catch variability is achieved at the expense of higher biomass variability. The fixed catch strategy has no catch variability but it also has the highest biomass variability.

Figure 7 shows the feasible sets of mean steady-state biomass and catch levels with different constraint probabilities. The feasible set is the set of values for which there is a viable solution to the control problem. The feasible set has as many dimensions as there are state variables with specified objectives. When the constraint probability for objective (i) >0.95 (i.e. it allows no more than 5% probability of biomass falling below the LRP), the feasible sets are vanishingly small. Thus, there are trade-offs between constraint probability and feasible TRPs, as well as between feasible TRPs for different variables (catch and biomass). The lower edges of the sets are due to the specified LRP while the upper edges are caused by the constraint that catch levels must be >0.

Figure 8 summarises the trade-off between the constraint probability specified in objective (i) and the mean steady-state catch level. The steepness of the curve at high probabilities implies a rapid decrease in feasible TRPs for catch with increasing constraint probability, and is consistent with the feasible set vanishing at high constraint probabilities.

Discussion

This study identified an MPC strategy that achieved the specified objectives in the operating model. A contrasting fixed catch limit with a slightly lower average catch was twice as likely to result in the depletion of modelled predators below 60% of their initial levels. The chosen objectives and constraints illustrate the sort of objectives that might be developed from CCAMLR's conservation principles, and the operating model incorporated spatial and trophic structure and uncertainty, which are important characteristics of the fished ecosystem. The results, therefore, support the intention within CCAMLR to replace the current fixed catch limit for Antarctic krill with a feedback approach. They also support the maintenance of a highly precautionary catch limit (such as the current 620 000 tonnes trigger level) until an appropriate feedback approach is implemented.

The development of a feedback approach is a substantial undertaking. Figures 1 and 2 illustrate the elements of such an approach if it were developed as an MP. MPs often include a formal state estimation process, known as the assessment model, to integrate available information and estimate the relevant state variables. This, in turn, requires regular measurements of some of these variables. Consideration of the observability (Kwakernaak and Sivan, 1972) of plausible models of ecosystem dynamics might help to prioritise variables for measurement. This identification, collection and integration of relevant information is a priority task in the development of feedback management. Ultimately such an approach must be evaluated in terms of its performance for relevant objectives (Rademeyer et al., 2007), so identifying these objectives, or at least the performance statistics that they require, is also a priority.

It is also necessary to develop candidate HCRs. These are often phrased as a response to a single state variable, usually a spatially aggregated estimate of harvested species biomass. For example, when biomass is above one LRP, catch is a fixed proportion of biomass; when biomass is below that LRP, the catch proportion declines with biomass; and when the biomass is below a second, lower LRP, catch is zero (Hilborn and Hilborn, 2012). Such HCRs do not require the controller components illustrated in Figure 2 and have the advantage that they can be understood and calculated by non-specialists.

Nonetheless, EBM implies objectives for multiple state variables, such as predators and prey in multiple subareas. For example, Watters et al. (2013) considered 36 separate state variables in their evaluation of candidate management measures for the krill fishery. It is unlikely that a single variable (or index that summarises multiple variables, e.g. de la Mare and Constable, 2000) will adequately indicate the state of the entire ecosystem in which the krill fishery operates. Furthermore, CCAMLR has defined 15 small-scale management units for the Scotia Sea and southern Drake Passage (Hewitt et al., 2004) and currently has catch limits for three larger subareas (Hill, 2013a). The potential spatial complexity of management and the complex interrelationships between state variables suggest that it is worth considering HCRs based on methods that can integrate this complex information. The current study suggests that MPC could, in principle, be used to develop such candidate HCRs.

The illustrative MPC-based HCR presented here would require further development and validation in order to be considered a candidate HCR. Such development might include a non-linear prediction model and a more detailed cost function, and should occur in concert with the development of a state estimator mentioned above. An immediate priority, which is feasible within a theoretical framework, is to assess the robustness of the approach to higher levels of uncertainty.

MPE is an appropriate framework for selecting an HCR by identifying the candidate HCR that provides the best trade-off between objectives (Rademeyer et al., 2007). The best candidate is identified through manual inspection of performance statistics, which integrates subjective, and potentially unstated, opinions. These could, of course, include preferences for parsimonious HCRs. A key issue in assessing any candidate MP is to ensure an appropriate representation of uncertainty in the state estimation process and any expectations about how system dynamics will respond to the HCR (Hill et al., 2007; Rademeyer et al., 2007; Link et al., 2012). These uncertainties can include, amongst other things, the effects of lags introduced by the time that elapses between measurements of state and the implementation of management measures.

CCAMLR's scientific working groups have established good practice to deal with uncertainty in the understanding of krill–predator–fishery–environment interactions (Link et al., 2012). This involves the use of two separate operating models, each of which is implemented using four separate parameterisations and used to produce multiple stochastic simulations (Hill et al., 2007; Plagányi and Butterworth, 2012; Watters et al., 2013). In the current study, the operating and prediction models shared common values for demographic and transport parameters. Such parameters might also be used in an assessment model. MPE evaluates the whole feedback process, including estimation, and should account for the effects of parameter uncertainty. One solution is to represent a range of plausible values for these parameters in the operating models, while maintaining the best estimates in the estimator and prediction model. The specific optimisation algorithm used to identify MPC-based HCRs will depend on the nature of the represented uncertainty and it might be appropriate to include robust min–max optimisation (e.g. Ben-Tal et al., 2009) to minimise the risk associated with the worst-case scenario. The performance of any MP in simulation does not guarantee the achievement of objectives in the real world, which is why monitoring and contingency plans are also important.

The approach described in this study offers opportunities for non-specialists to participate in the design of candidate HCRs. The selection of TRPs is based on trade-offs between objectives, as illustrated in Figure 7. The available trade-offs will be influenced by the weighting of objectives and state variables in Equation 8, such that there are likely to be more options for achieving more highly weighted objectives. Involving stakeholders in the selection of weights and TRPs therefore allows some expression of unstated or subjective preferences. It would be possible to develop multiple alternative HCRs based on different TRPs implying different trade-offs. The selection of TRPs is based on trade-offs in steady-state conditions, so it will still be necessary to use MPE to evaluate performance in dynamic and uncertain conditions.

In addition to the potential development of candidate HCRs, optimisation-based approaches can contribute to the development of MPs in other ways. Firstly, it is useful to assess whether it is likely to be feasible to maintain the state of an ecosystem within the range that satisfies the diverse objectives implied by EBM. Conversely, it is also useful to assess whether objectives are realistic. CCAMLR has already embarked on the substantial enterprise of developing a feedback approach for Antarctic krill and the current study provides some indication that this is feasible in principle. The study's methods could be adapted to use the operating models of Plagányi and Butterworth (2012) and Watters et al. (2013) to provide a further assessment of feasibility.

Secondly, optimisation-based approaches provide a straightforward, quantitative way to identify and evaluate trade-offs by assessing whether combinations of candidate objectives imply a feasible solution and simulating the outcomes. Figure 8 characterises an important trade-off. The quantitative information is uncertain because it applies to an operating model and not the real world, but the shape is also informative. Hill et al. (2007) suggested that, because of the limited processing capacity of the human mind (Miller, 1956), MPEs should assess on no more than about seven performance statistics. Nonetheless, Watters et al. (2013) provided statistics on 36 separate state variables because of a lack of prior information about which state variables were important to decision-makers. Any arbitrary aggregation of the state variables could have biased the evaluation (Hill, 2013b). Multiparametric optimisation methods (Fiacco, 1983) can deal with far more than seven pieces of information. In addition to deriving optimised HCRs, these methods can be used to assess how close pre-specified HCRs are to optimal. They can identify those state variables which have the strongest influence on the optimal solution, and assess how this solution varies with the values of different state variables. In other words, they can be used to identify those objectives which are in most need of modification, or the most problematic uncertainties.

Conclusions

This analysis demonstrates that robust control, implemented using an optimised HCR, is a feasible way of achieving multiple EBM-like objectives in a spatially and trophically structured model system when state estimates are uncertain. An optimised HCR is more likely than fixed catch levels to achieve these objectives. Robust control might therefore be a feasible approach to EBM of the krill fishery. However, the development of such an approach needs investment in the definition of objectives, the characterisation of uncertainty, monitoring of the appropriate state variables, and the development of operating models, assessment models and controllers. A more immediate application of the approach outlined in this study is to help develop objectives for the state of the system by evaluating the feasibility of objectives, and the trade-offs that they imply.

Acknowledgements

This paper is a contribution to the British Antarctic Survey's Natural Environment Research Council core-funded Ecosystems Programme.

References

Atkinson, A., R.S. Shreeve, A.G. Hirst, P. Rothery, G.A. Tarling, D.W. Pond, R.E. Korb, E.J. Murphy and J.L. Watkins. 2006. Natural growth rates in Antarctic krill (*Euphausia superba*): II. Predictive models based on food, temperature, body length, sex, and maturity stage. *Limnol. Oceanogr*., 51: 973–987.

Ben-Tal, A., L. El Ghaoui and A. Nemirovski. 2009. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, New Jersey.

Bitmead, R.R., M. Gevers and V. Wertz. 1990. *Adaptive Optimal Control: the Thinking Man's GPC*. Prentice-Hall, Upper Saddle River, New Jersey.

Butterworth, D.S., A.E. Punt and M. Basson. 1991. A simple approach for calculating the potential yield of krill from biomass survey results. In: *Selected Scientific Papers,* 1991 (*SC-CAMLR-SSP/8*). CCAMLR, Hobart, Australia: 207–217.

Caddy, J.F. and R. Mahon. 1995. Reference points for fisheries management. *FAO Fish. Tech. Pap.*, 347. FAO, Rome, Italy: 83 pp.

Cannon, M., B. Kouvaritakis and X. Wu. 2009. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. *IEEE Trans. Automatic Control*, 54 (7): 1626–1632.

Cannon, M., B. Kouvaritakis, S. Raković and Q. Cheng. 2011. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Trans. Automatic Control*, 56 (1): 194–200.

Constable, A.J. 2001. The ecosystem approach to managing fisheries: achieving conservation objectives for predators of fished species. *CCAMLR Science*, 8: 37–64.

Constable, A.J., W.K. de la Mare, D.J. Agnew, I. Everson and D. Miller. 2000. Managing fisheries to conserve the Antarctic marine ecosystem: practical implementation of the Convention on the Conservation of Antarctic Marine Living Resources (CCAMLR). *ICES J. Mar. Sci.*, 57 (3): 778–791.

Couchman, P.D., M. Cannon and B. Kouvaritakis. 2006. Stochastic MPC with inequality stability constraints. *Automatica*, 42 (12): 2169–2174.

Cunningham, S. and J.J. Maguire. 2002. Factors of unsustainability in fisheries. In: Gréboval, D. (Ed.). *Report and Documentation of the International Workshop on Factors Contributing to Unsustainability and Overexploitation in Fisheries*. Bangkok, Thailand, 4 to 8 February 2020. *FAO Fisheries Report*, No. 672. FAO, Rome: 55–92.

de la Mare, W.K. and A.J. Constable. 2000. Utilising data from ecosystem monitoring for managing fisheries: development of statistical summaries of indices arising from the CCAMLR Ecosystem Monitoring Program. *CCAMLR Science*, 7: 101–117.

Fiacco, A.V. 1983. *Introduction to Sensitivity and Stability: Analysis in Nonlinear Programming.* Academic Press, Waltham, Massachusetts.

Grant, S.M., S.L. Hill, P.N. Trathan and E.J. Murphy. 2013. Ecosystem services of the Southern Ocean: trade-offs in decision-making. *Ant. Sci*., doi: 10.1017/S0954102013000308.

Grüne, L. and J. Pannek. 2011. *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer-Verlag, Berlin: 360 pp.

Hewitt, R.P., G. Watters, P.N. Trathan, J.P. Croxall, M.E. Goebel, D. Ramm, K. Reid, W.Z. Trivelpiece and J.L. Watkins. 2004. Options for allocating the precautionary catch limit of krill among small-scale management units in the Scotia Sea. *CCAMLR Science*, 11: 81–97.

Hilborn, R. and U. Hilborn 2012. *Overfishing: What Everyone Needs to Know*. Oxford University Press, Oxford.

Hilborn, R. and C.J. Walters. 1992. *Quantitative Fisheries Stock Assessment: Choice, Dynamics and Uncertainty*. Chapman and Hall, New York: 570 pp.

Hill, S.L. 2013a. Prospects for a sustainable increase in the supply of long chain omega 3s: lessons from the Antarctic krill fishery. In: De Meester, F., R.R. Watson and S. Zibadi (Eds.). *Sustainable Long Chain Omega-3 Fatty Acids in Cardiovascular and Mental Health*. Humana Press, New York: 267–296.

Hill, S.L. 2013b. From strategic ambiguity to technical reference points in the Antarctic krill fishery: the worst journey in the world? *Environmental Conservation,* doi: 10.1017/S0376892913000088.

Hill, S.L., E.J. Murphy, K. Reid, P.N. Trathan and A.J. Constable. 2006. Modelling Southern Ocean ecosystems: krill, the food-web, and the impacts of fishing. *Biol. Rev. Camb. Philos. Soc.,* 81: 581–608.

Hill, S.L., G.M. Watters, A.E. Punt, M.K. McAllister, C. LeQuere and J. Turner. 2007. Model uncertainty in the ecosystem approach to fisheries. *Fish Fish*., 8: 315–333.

Isidori. A. 1995. *Nonlinear Control Systems* (3rd ed.). Springer-Verlag, Berlin.

Kell, L.T., J.A.A. De Oliveira, A.E. Punt, M.K. McAllister and S. Kuikka. 2006. Operational Management Procedures: an introduction to the use of Evaluation Frameworks. In: Motos, L. and D. Wilson (Eds.). *The Knowledge Base for Fisheries Management*. Elsevier Limited, Amsterdam: 379–407.

Kouvaritakis, B., M. Cannon, S. Raković and Q. Cheng. 2010. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46 (10): 1719–1724.

Kwakernaak, H. and R. Sivan. 1972. *Linear Optimal Control Systems*. Wiley-Interscience, Hoboken, New Jersey: 608 pp.

Link, J.S., T.F Ihde, C.J. Harvey, S.K. Gaichas, J.C. Field, J.K.T. Brodziak, H.M. Townsend and R.M. Peterman. 2012. Dealing with uncertainty in ecosystem models: the paradox of use for living marine resource management. *Prog. Oceanogr.*,102: 102–114.

Luenberger, D.G. 1966. Observers for multivariable systems. *IEEE Transactions on Automatic Control*, 11 (2): 190–197.

McLeod, K.L. and H.M. Leslie. 2009. Why ecosystem-based management? In: McLeod, K. and H. Leslie (Eds). *Ecosystem-Based Management for the Oceans*. Washington, DC, USA: Island Press: 3–12.

Mayne, D.Q., J.B. Rawlings, C.V. Rao and P.O.M. Scokaert. 2000. Constrained model predictive control: Stability and optimality. *Automatica*, 36: 789–814.

Miller, G.A. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *The Psychological Review*, 63: 81–97.

Murphy, E.J., J.L. Watkins, P.N. Trathan, K. Reid, M.P. Meredith, S.L. Hill, S.E. Thorpe, N.M. Johnston, A. Clarke, G.A. Tarling, M.A. Collins, J. Forcada, A. Atkinson, P. Ward, I.J. Staniland, D.W. Pond, R.A. Cavanagh, R.S. Shreeve, R.E. Korb, M.J. Whitehouse, P.G. Rodhouse, P. Enderlein, A.G. Hirst, A.R. Martin, D.R. Briggs, N.J. Cunningham and A.H. Fleming. 2012. Spatial and temporal operation of the Scotia Sea ecosystem. In: Rogers, A.D., N.M. Johnston, E.J. Murphy and A. Clarke (Eds). *Antarctic Ecosystems: An Extreme Environment in a Changing World*. Blackwell Publishing, Oxford: 160–212.

Nicol, S. and V. Siegel. 2000. Population parameters. In: Everson, I. (Ed.). *Krill: Biology, Ecology and Fisheries*. Blackwell Science, Oxford: 40–79.

Penas, E. 2007. The fishery conservation policy of the European Union after 2002: towards long-term sustainability. *ICES J. Mar. Sci*., 64 (4): 588–595.

Plagányi, É.E. and D.S. Butterworth. 2012. The Scotia Sea krill fishery and its possible impacts on dependent predators – modelling localized depletion of prey. *Ecol. Appl.*, 22 (3): 748–761.

Rademeyer, R.A., É.E. Plagányi and D.S. Butterworth. 2007. Tips and tricks in designing management procedures. *ICES J. Mar. Sci*., 64 (4): 618–625.

Rawlings, J.B. and D.Q. Mayne. 2009. *Model Predictive Control*. Nob Hill Publishing, Madison, Wisconsin: 533 pp.

Rosenberg, A.A. 2009. Changing U.S. ocean policy can set a new direction for marine resource management. *Ecology and Society*, 14 (2): 6.

SC-CAMLR. 2011. *Report of the Thirtieth Meeting of the Scientific Committee (SC-CAMLR-XXX)*. CCAMLR, Hobart, Australia: 454 pp.

Smith, A.D., C.J. Brown, C.M. Bulman, E.A. Fulton, P. Johnson, I.C. Kaplan, H. Lozano-Montes, S. Mackinson, M. Marzloff, L.J. Shannon, Y.-J. Shin and J. Tam. 2011. Impacts of fishing low-trophic level species on marine ecosystems. *Science*, 333: 1147–1150.

Walters, C. 1986. *Adaptive Management of Renewable Resources*. Macmillan Publishing, New York: 374 pp.

Watters, G.M., S.L. Hill, J. Hinke, J. Matthews and K. Reid. 2013. Decision making for ecosystem based management: evaluating options for a krill fishery with an ecosystem dynamics model. *Ecol. Appl.*, 23 (4): 710–725.

Zhou, K. and J.C. Doyle. 1997. *Essentials of Robust Control*. Prentice Hall, Upper Saddle River, New Jersey: 411 pp.

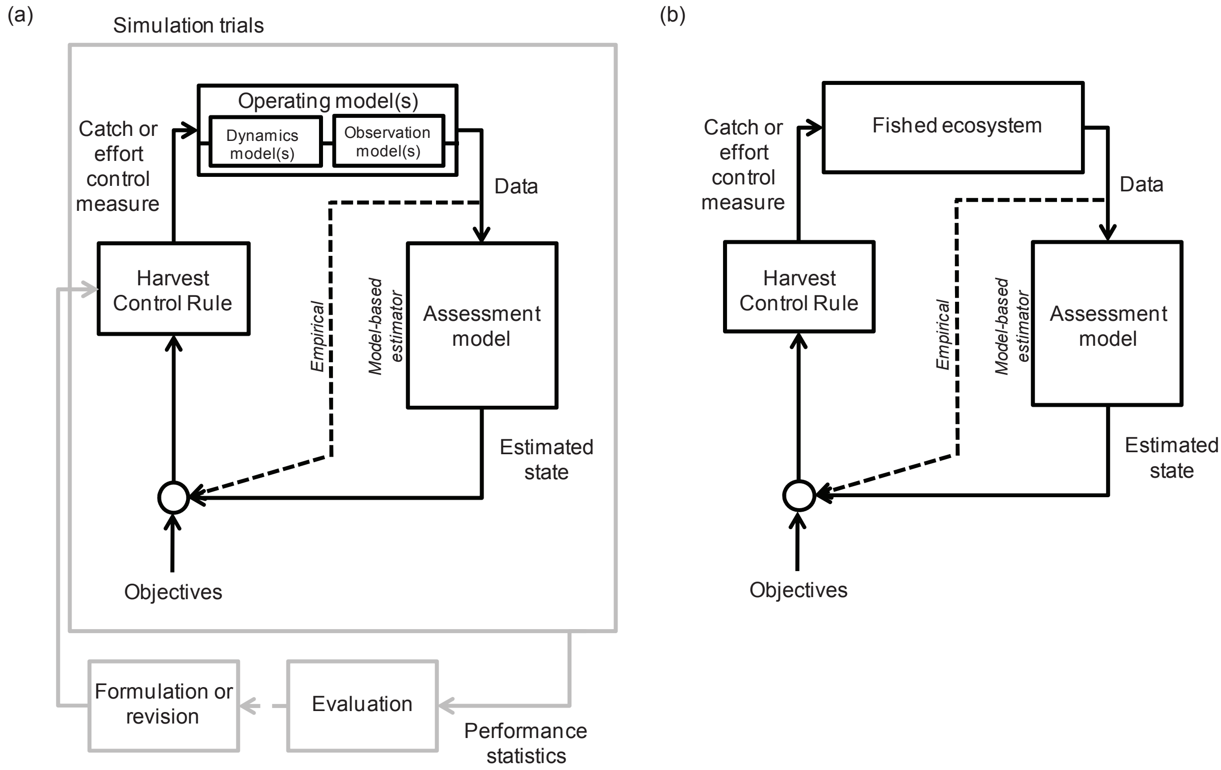

Figure 1: The management procedure approach showing: (a) management procedure evaluation and (b) an operational management procedure, based on the descriptions in Rademeyer et al. (2007).

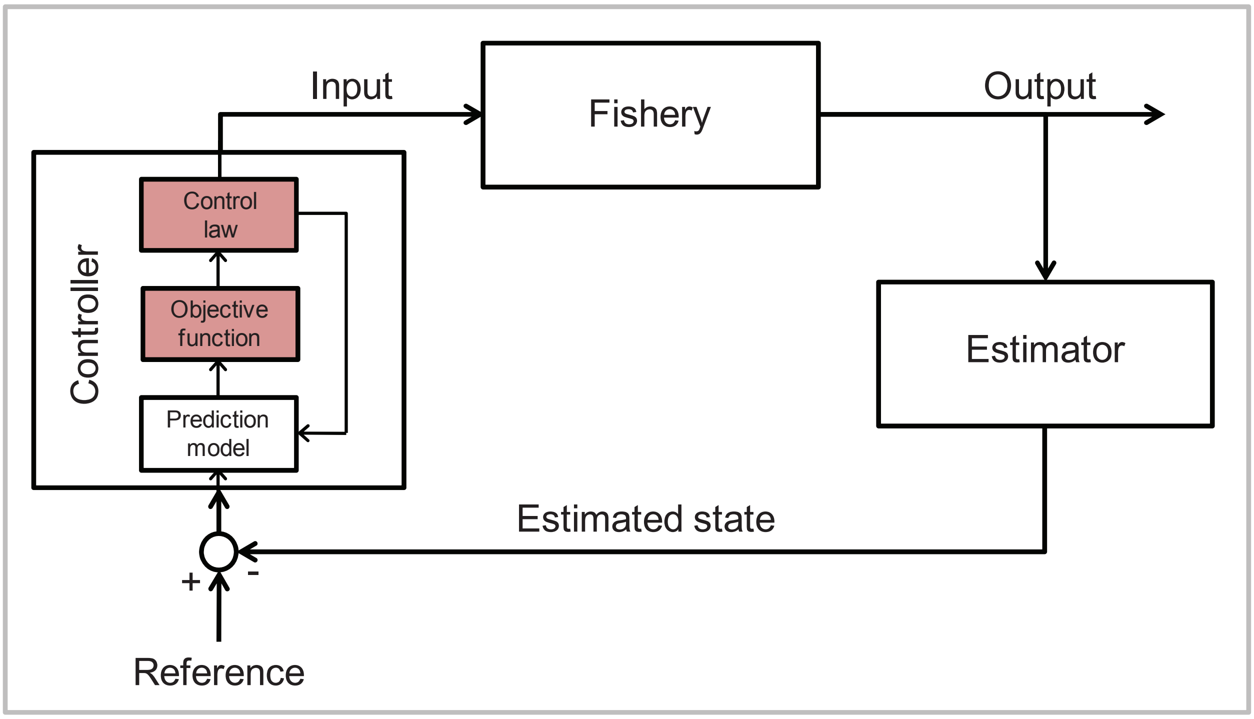

Figure 2: A Model Predictive Control strategy applied, hypothetically, to a fishery. Coloured elements indicate the methodological focus of the current study.

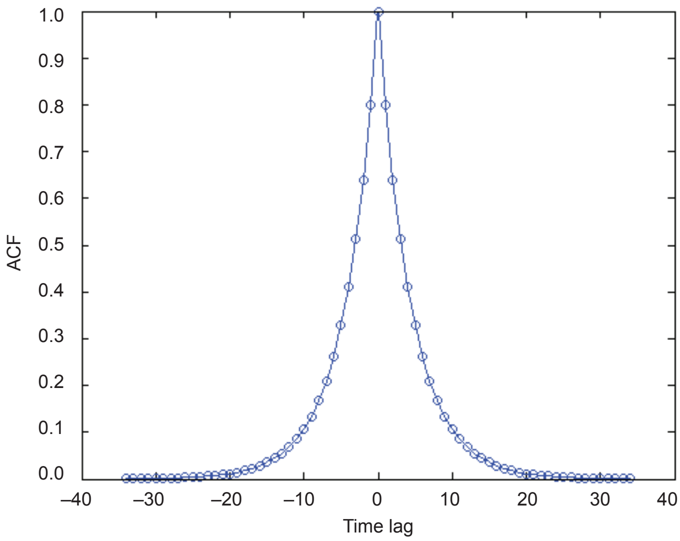

Figure 3: Autocorrelation function of the disturbance variables (harvested species biomass in area 0 and carrying capacities in areas 1 and 2) generated using the spectral bandwidths in Table 1.

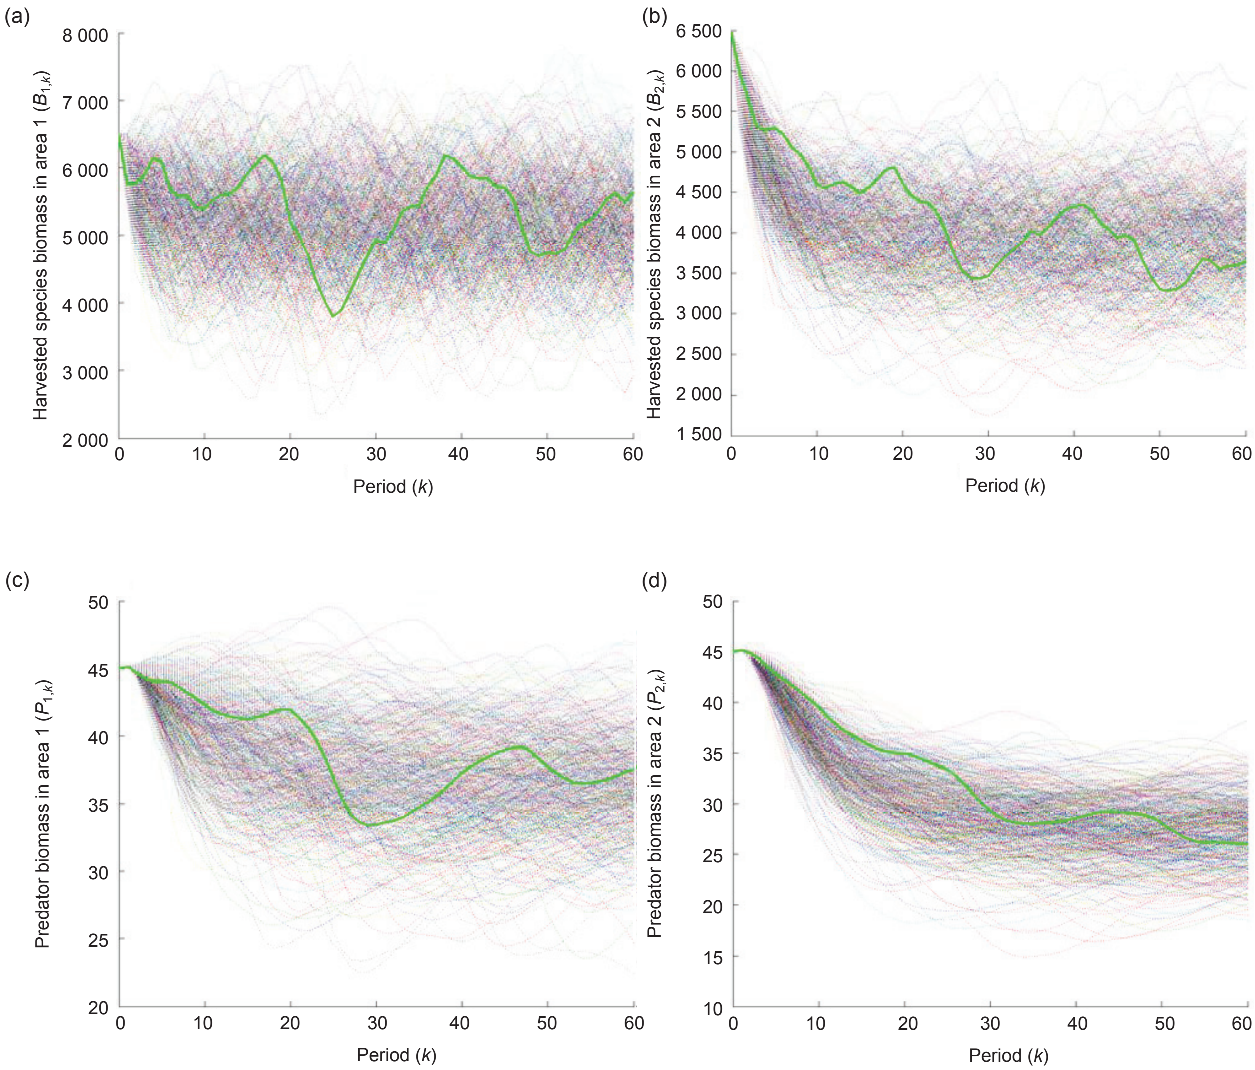

Figure 4: Biomass of the harvested species (a, b) and its predators (c, d) in two areas from 500 stochastic realisations of the operating model with fixed catch limits in each area. Solid green lines highlight a single randomly selected simulation.

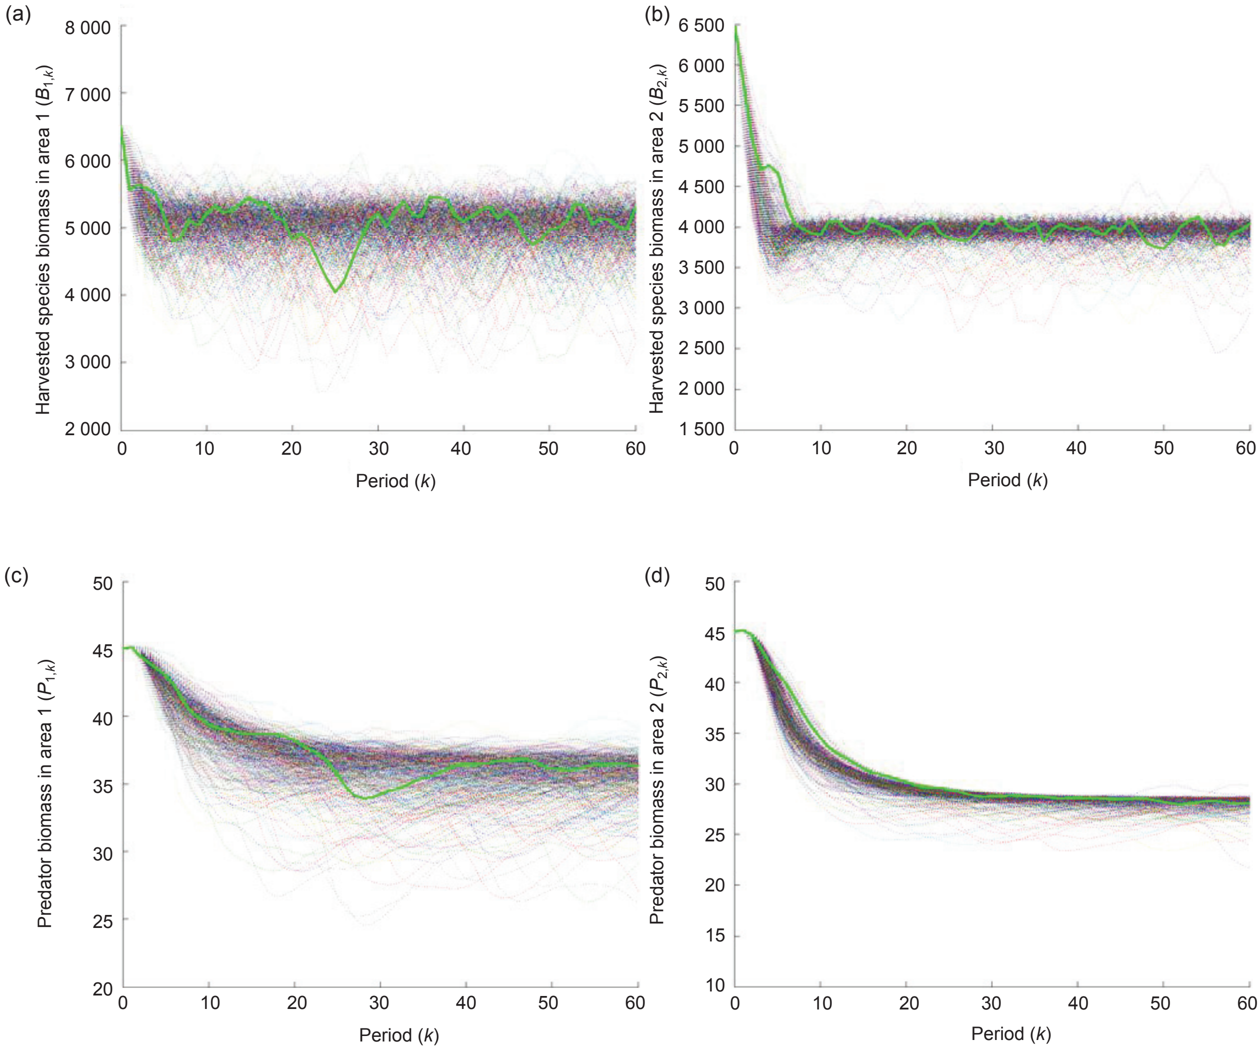

Figure 5: Biomass of the harvested species (a, b) and its predators (c, d) in two areas from 500 stochastic realisations of the operating model with an MPC strategy. Solid green lines highlight a single simulation using the same random number sequences as that highlighted in Figure 4.

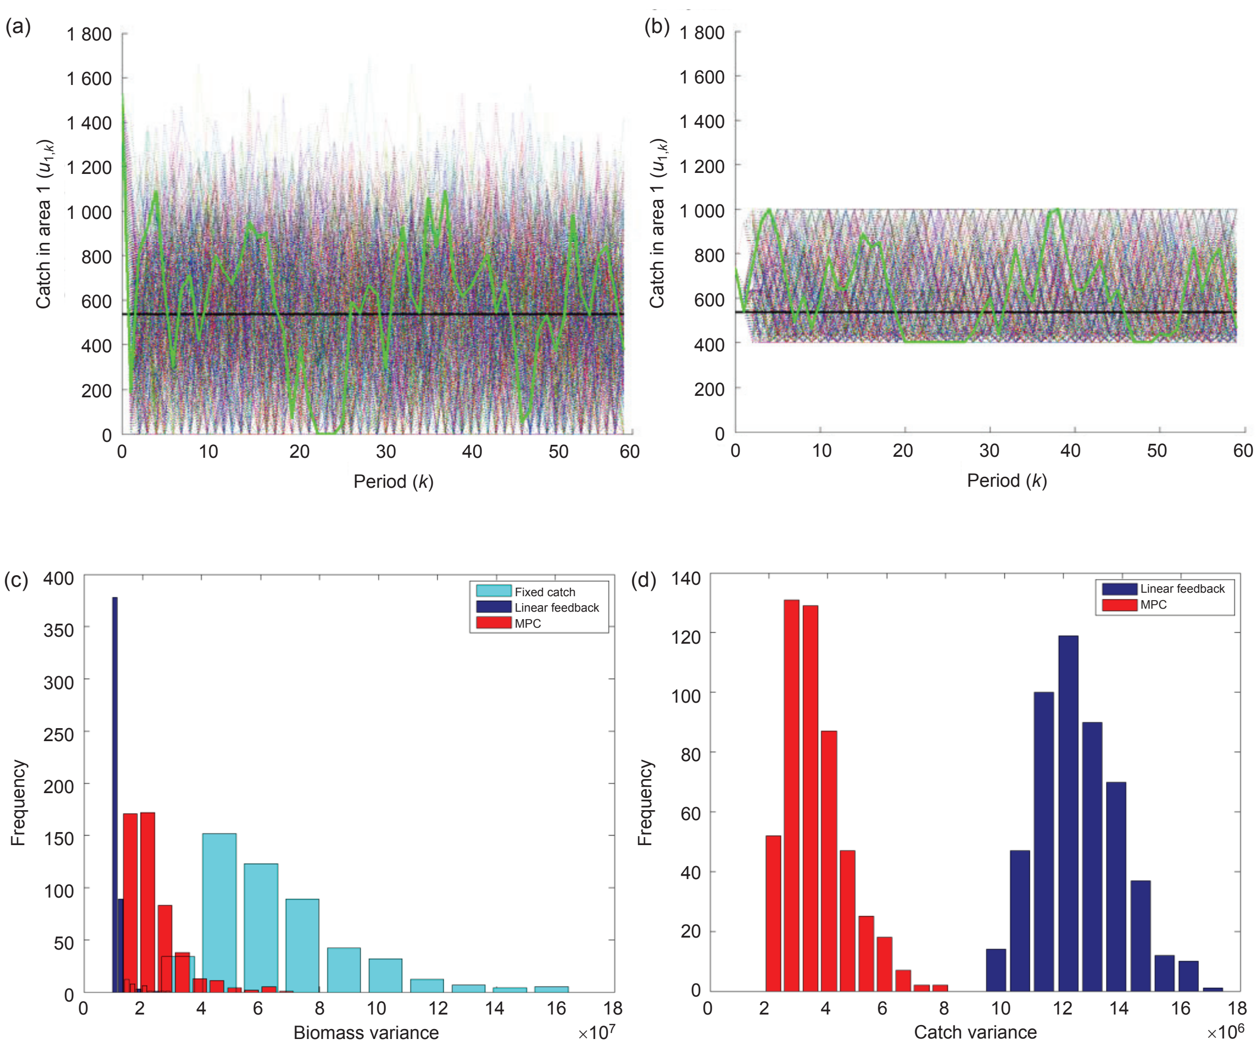

Figure 6: Catch in area 1 from 500 stochastic realisations of the operating model for linear feedback (a) and MPC (b). Solid black lines indicate the mean steady-state catch levels (which are also the TRPs). Distributions of variances of biomass (c) and catch (d) for 500 stochastic operating model realisations.

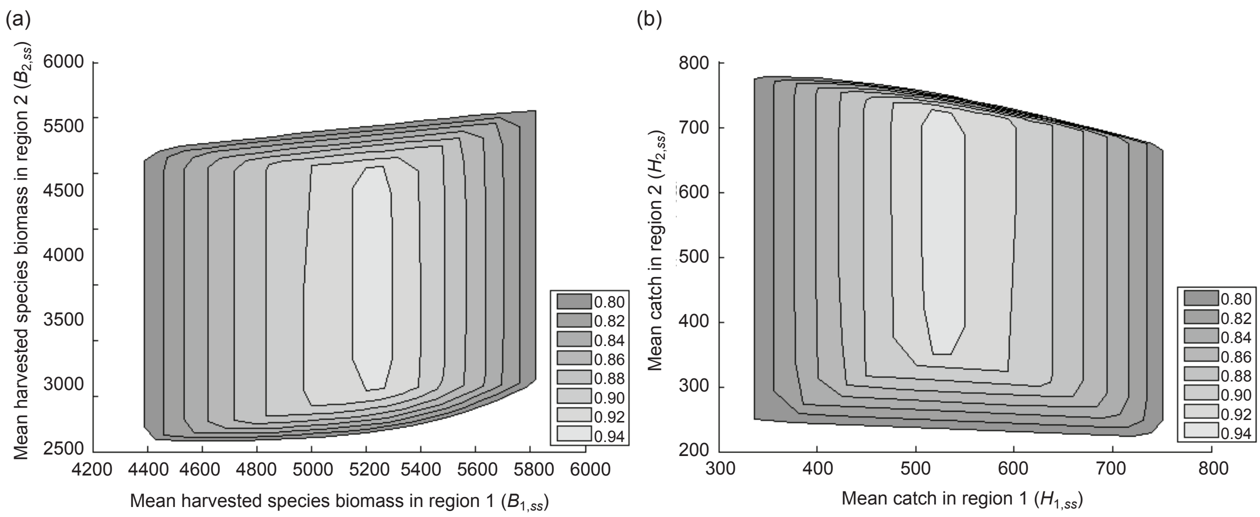

Figure 7: Feasible mean harvested species biomass (a) and mean catch level (b) in steady-state operation under MPC, versus the constraint probability for objective (i) illustrating the trade-offs between constraint probability, catch and biomass.

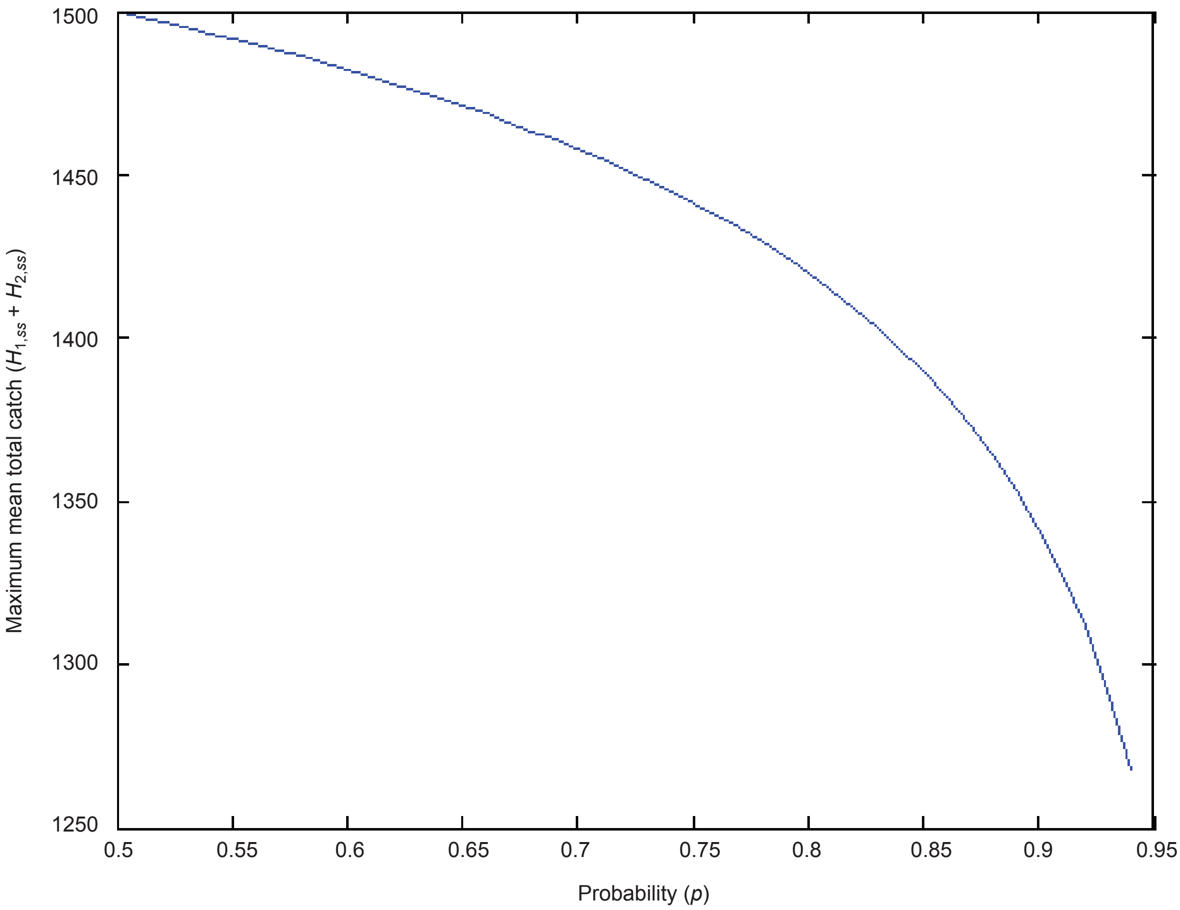

Figure 8: Maximum feasible catch level under steady-state operation with MPC, versus constraint probability for objective (i), illustrating the trade-off between constraint probability and catch.