Canadian Journal of
Fisheries and
Aquatic Sciences

Article

Optimal dynamic spatial closures can improve fishery yield and reduce fishing-induced habitat damage

Anna J. Poulton [a], **Suresh A. Sethi** [b], **Stephen P. Ellner** [a,c], and **T. Scott Smeltz** [b]

[a]Center for Applied Mathematics, Cornell University, Ithaca, NY 14853, USA; [b]Fisheries, Aquatic Science and Technology Laboratory, Alaska Pacific University, Anchorage, AK 99508, USA; [c]Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, NY 14853, USA

Corresponding author: **Anna J. Poulton** (email: ajp374@cornell.edu)

Abstract

Bottom-towed fishing gears produce significant amounts of seafood globally but can result in seafloor habitat damage. Spatial closures provide an important option for mitigating benthic impacts, but their performance as a fisheries management policy depends on numerous factors, including how fish respond to habitat quality changes. Spatial fisheries management has largely focused on marine protected areas with static locations, overlooking dynamic spatial closures that change through time. To investigate the performance of dynamic closures, we develop a spatial fishery model with fishing-induced habitat damage, where habitat quality can affect both fish productivity and movement. We find that dynamic spatial closures often achieve greater harvest and habitat protection than fixed marine protected areas or conventional nonspatial maximum sustainable yield management, especially under strong habitat–stock interactions. Determining optimal dynamic spatial closures may require considerable information, but we find that simple policies of fixed-schedule rotating closures also perform well. Dynamic spatial closures have received less attention as fisheries management tools, and our results demonstrate their potential value for addressing both harvest and habitat impacts from fishing.

Key words: dynamic spatial closures, ecosystem-based fisheries management, habitat damage mitigation, marine protected areas, optimal harvesting

1. Introduction

While fisheries management has historically focused on maximizing yield for single stocks, more attention is now given to maintaining the health of ecosystem components such as seafloor habitats (Pikitch et al. 2004). Along with providing valuable ecosystem services (Costanza et al. 1997; Barbier 2017), healthy seafloor habitat may increase juvenile survival and recruitment (Collie et al. 1997; Bradshaw et al. 2003; Scharf et al. 2006), enhance fish body condition (Hiddink et al. 2011), and improve fish population recovery (Auster 1998; Hutchings and Reynolds 2004). This recognition has resulted in laws and mandates to protect seafloor habitat: for instance, in the United States the Magnuson–Stevens Fishery Conservation and Management Act requires that "essential fish habitat" be considered in fishery management plans (NOAA National Marine Fisheries Service 2007). The Marine Strategy Framework Directive in the European Union similarly calls for maintaining seafloor integrity (European Commission 2008).

Fishing gear that contacts the seafloor can damage benthic habitat in a variety of ways, including reducing seafloor complexity, removing epiflora and epifauna, and disrupting biogeochemical processes (Peterson et al. 1987; Collie et al. 1997; Watling and Norse 1998; Watling et al. 2001; Sala et al. 2021). Mobile bottom gears that contact the seafloor (e.g. bottom trawls) are widely used, accounting for approximately a quarter of worldwide seafood catch (Cashion et al. 2018). For these reasons, bottom gears are a primary anthropogenic source of physical seafloor disturbance in commercially fished areas (Watling and Norse 1998; Foden et al. 2011; Ramirez-Llodra et al. 2011). Often habitat quality can recover over time, but the rate of recovery may differ greatly depending on the type of habitat (Rooper et al. 2011; Grabowski et al. 2014).

The three main management tools generally available to limit seafloor habitat damage from fishing are effort regulations, gear regulations, and spatial closures (National Research Council 2002; Smeltz et al. 2019). It is crucial to understand how and when these tools might be able to protect seafloor habitat while sustaining or improving harvests. Here we focus on spatial closures. While gear and effort regulations can be useful, they may require difficult socioeconomic tradeoffs and may not be feasible in fisheries with limited resources for monitoring and enforcement (National Research Council 2001, 2002; Russ 2002; Reithe et al. 2014; Cabral et al. 2019). In contrast, spatial closures may be logistically easier to manage once implemented, and have been used widely in developing nations (Myers et al. 2000; Reithe et al. 2014; Cabral et al. 2019). Not all spatial closures completely ban all types of fishing (Costello and Ballantine 2015), but here we focus on strictly "no-take" closures.

We classify spatial closures based on whether the placement of closures is fixed (marine protected areas, MPAs) or allowed to change over time (dynamic spatial closures, DSCs). Broadly, spatial fisheries management has focused on MPAs, with studies and real-world implementations of DSCs being less abundant. DSCs may range greatly in complexity, from simple fixed schedules of rotations to adaptive policies, and have often focused on the management of sedentary benthic species (Cohen and Foale 2013; Plagányi et al. 2015; O'Boyle et al. 2017). For example, the Australian Queensland sea cucumber fishery uses rotational closures where the fishing area is split into over 150 zones, and each zone is open one of every 3 years (Plagányi et al. 2015; Commonwealth of Australia 2021). The U.S. Atlantic sea scallop fishery uses adaptive rotations, closing areas with many small scallops and reopening areas when the scallops have grown large (New England Fishery Management Council 2003). DSCs have also been used for small-scale fisheries throughout the Indo-Pacific (where they are a traditional form of management); for instance, on Kakarotan Island (Sangihe-Talaud Archipelago, Indonesia), community leaders select one area of reef each year to be closed for several months (Cinner et al. 2005; Cohen and Foale 2013).

Beyond limiting seafloor habitat damage, spatial closures can result in increased stock density and biomass within the closure (Hilborn et al. 2004; Lester et al. 2009). Adults and/or larvae may "spill over" into fished areas, counteracting losses in harvest due to the closure, but the overall effect of spatial closures on yield depends on a number of factors (Hilborn et al. 2004; Hart 2006). Many models of no-take areas (e.g., Hastings and Botsford 1999; Hart 2006; Le Quesne and Codling 2009; Carvalho et al. 2019) have found that they cannot increase overall yield beyond what can be obtained with optimal nonspatial management. However, such models do not consider interactions between stocks and habitat quality, potentially missing an important driver of stock dynamics under fishing-induced habitat damage. Along with possible impacts to population growth, empirical results have suggested that fish might prefer being in areas with higher quality habitat, reducing their movement out of such areas (Claudet et al. 2010; Grüss et al. 2011). Both of these interactions may affect the performance of spatial closures in meeting harvest and habitat protection goals.

Our work extends past models for the optimization of spatial closures in fisheries with fishing-induced habitat damage in several significant ways (Moeller and Neubert 2013, 2015; Akpalu and Bitew 2014; Reithe et al. 2014; Nichols et al. 2018; Alqawasmeh and Lutscher 2019; Kelly et al. 2019; Falcó and Moeller 2022). First, we consider two potential habitat–stock interactions: the impact of habitat quality on fish population growth and on fish movement. While several previous models have incorporated the former effect, only Alqawasmeh and Lutscher (2019) have considered the latter, even though the movement behavior of fish may greatly impact the effectiveness of different spatial management policies. Second, we explicitly model the impacts of fishing on seafloor habitat quality and the recovery of habitat over time. Most of the previously mentioned works have not explicitly considered the dynamics of habitat quality (but see Kelly et al. 2019 and Nichols et al. 2018). Third, we model an open-access fishery whereby spatial closures may be the only viable management option available. This differs from most previous works, in which spatial management is examined in combination with (or arising from) harvest/effort controls (but see Reithe et al. 2014). Finally, we compare MPAs with DSCs to see when significant management benefits can be gained by letting spatial closures change over time. With the exception of Kelly et al. (2019), past works have focused exclusively on optimizing MPAs.

We develop an n-patch model of a fishery in which the stock, habitat quality, and fishing effort in each patch are explicitly modeled. The closure status of each patch may change on a yearly basis, allowing us to examine DSCs. Due to the complexity of our model, we utilize an ADP algorithm to find near-optimal policies of DSCs. These DSC policies, in which the manager decides which patches are open/closed each year, are then compared with optimal MPAs, i.e. fixed spatial closures. We assess the performance of DSCs and MPAs under both conventional management focused solely on maximizing yields, and the multiobjective case where both yield and habitat protection are recognized as management objectives.

The primary question we ask is: under what conditions can DSCs significantly outperform MPAs? We explore this question by varying the strength of the habitat–stock interactions, along with key habitat properties (recovery rate and sensitivity), for both a sedentary stock and a mobile stock. As an important performance benchmark, we also compare our DSC policies to nonspatial management at maximum sustained yield (MSY) harvest rates. Recognizing that optimal DSC policies may be hard to develop and implement in real life, we additionally consider the performance of simple heuristic policies of rotating spatial closures on fixed schedules. Finally, we briefly examine a stochastic version of our model in which recruitment is variable, reflecting a key source of population variability in real fisheries (Thorson et al. 2014).

2. Methods

2.1. Model framework

We model a fishery with n patches. Let $x_i(t)$, $h_i(t)$, and $f_i(t)$ be the stock biomass, habitat quality, and fishing effort in patch i at time t, respectively. Thus, at any given time the state of the fishery can be described by the vector $\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{h}(t), \mathbf{f}(t))$, where $\mathbf{x}(t) = (x_1(t), \ldots, x_n(t))$, $\mathbf{h}(t) = (h_1(t), \ldots, h_n(t))$, and $\mathbf{f}(t) = (f_1(t), \ldots, f_n(t))$. We will let $\mathcal{S}$ denote the state space, the set of all potential states of our model. Note that since we only require $x_i(t), f_i(t) \geqslant 0$ and $0 < h_i(t) \leq h_{\max}$, the state space is continuous and large ($3n$-dimensional). Note that our model has been rescaled such that all the state variables are unitless.

Our model is not based on a specific real fishery or fish species, but is rather meant to represent the general case of a fishery with fishing-induced habitat damage that is managed only with spatial closures. However, as our equations use specific functional forms, we must make some assumptions about the fishery and the life history of the fish. Notably, we assume that the early life stages of fish contribute to a common juvenile pool, from which they may recruit into any of the patches. From there, our model allows for adults that are either sedentary or mobile. We also assume specific forms for the effects of habitat on the stock, although we vary the strength of these to represent a wide range of possible interactions. While these restrictions are important to note, our model is still relevant for a wide range of species. Model state dynamics and assumptions are described in detail in Section 2.2.

A summary of our modeling framework is shown in Fig. 1. At the beginning of each year, the manager makes a decision about which patches should be open to fishing, and which should be closed to fishing (in which case no fishing effort or take is permitted within the patch). Note that throughout this work we use d to denote the year, while t is used for continuous time. The action taken by the manager at the beginning of year d is denoted by $\mathbf{a}(d) = (a_1(d), \dots, a_n(d))$, with $a_i(d) = 1$ indicating that patch i is open, and $a_i(d) = 0$ indicating that patch i is closed, for the duration of year d (in other words, $a_i(t) = 1$ for all $t \in [d, d+1)$ if $a_i(d) = 1$). We also assume that due to stakeholder concerns, the manager is not allowed to close all the patches at once. Letting $\mathcal{A}$ denote the set of all actions the manager can take in a single year; we have that $|\mathcal{A}| = 2^n - 1$.

The manager's goal can include both maximizing harvest and protecting habitat. The manager receives Reward(d) at the end of year d, which represents how well they met these two objectives. More specifically, Reward(d) is a weighted sum of the total harvest and the total habitat quality throughout the fishery over the course of year d. The weights can be adjusted to control how much the manager cares about maximizing harvest versus protecting seafloor habitat. The ultimate goal of the manager is to maximize the sum of the discounted rewards over an infinite time horizon. When the closure status of each patch is allowed to change on a yearly basis, the resulting policies are called DSCs. A key focus of our work will be comparing these DSC policies with MPAs, in which the closure status of each patch is not allowed to change over time ($a_i(d) = a_i$ for all d).

2.2. State dynamics

Recall that $x_i(t)$, $h_i(t)$, and $f_i(t)$ are the stock, habitat quality, and effort in patch i at time t, respectively. The dynamics of our fishery model are a system of ordinary differential equations (ODEs) having the form

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = \mathrm{Growth}_i(t) + \mathrm{Movement}_i(t) - \mathrm{Harvest}_i(t) \tag{1}$$

$$\frac{\mathrm{d}h_i}{\mathrm{d}t} = \mathrm{Recovery}_i(t) - \mathrm{Damage}_i(t) \tag{2}$$

$$\frac{\mathrm{d}f_i}{\mathrm{d}t} = \mathrm{Entry}_i(t) - \mathrm{Exit}_i(t) + \mathrm{Relocation}_i(t) \tag{3}$$

In this section, we first discus the two potential effects of habitat on the stock that we consider. Subsequently, we discuss each of the state processes (Growth, Movement, etc.) in eqs. 1–3, paying particular attention to where and how each habitat–stock effect is incorporated. Parameter values will be discussed later in Section 2.5.

2.2.1. Habitat–stock effects

We assume that habitat quality may potentially impact both the population growth and the movement behavior of the stock.

Empirical evidence has suggested several possible benefits of higher habitat quality to fish population growth. Here we specifically focus on how healthy seafloor habitat may increase the survival of early life stages (and subsequently, increase recruitment). We refer to this as the "habitat–stock survival effect". To represent this effect formally, we denote the scaled effect of patch i (with habitat quality $h_i(t)$) on survival by

$$A_i^{\mathrm{s}}(t) = 1 + 2\phi_{\mathrm{s}}(h_i(t) - 0.5) \tag{4}$$

Here ϕ_{s} denotes the strength of the habitat–stock survival effect, with $0 \le \phi_{\mathrm{s}} \le 1$. A value of $\phi_{\mathrm{s}} = 0$ indicates that early-stage survival is independent of habitat quality, while $\phi_{\mathrm{s}} = 1$ indicates a strong positive effect of habitat quality on early-stage survival. Note that with $h_{\max} = 1$ (which is used throughout this work), we have that $0 < h_i(t) \le 1$, and subsequently $1 - \phi_{\mathrm{s}} < A_i^{\mathrm{s}}(t) \le 1 + \phi_{\mathrm{s}}$.

Empirical evidence has also suggested that fish might have a preference for areas with greater habitat quality. We call this the "habitat–stock movement effect". In our model, we will assume that higher habitat quality may both (1) increase the preference recruits have for settling in a patch, and (2) decrease the movement rate of mobile adults from a patch. As recruits immediately join the adult population after settling in a patch, we assume that they have the same preference for habitat as adults. To represent the habitat–stock movement effect, we let the scaled preference that fish have for patch i (with habitat quality $h_i(t)$) be denoted by

$$A_i^{\mathrm{m}}(t) = 1 + 2\phi_{\mathrm{m}}(h_i(t) - 0.5) \tag{5}$$

Here ϕ_{m} denotes the strength of the habitat–stock movement effect, with $0 \le \phi_{\mathrm{m}} \le 1$. A value of $\phi_{\mathrm{m}} = 0$ indicates that patch preference is independent of habitat quality, while $\phi_{\mathrm{m}} = 1$ indicates a strong preference for patches with greater habitat quality. Note that with $h_{\max} = 1$ (which is used throughout this work), we have that $0 < h_i(t) \le 1$, and subsequently $1 - \phi_{\mathrm{m}} < A_i^{\mathrm{m}}(t) \le 1 + \phi_{\mathrm{m}}$.

While several past models considering optimal spatial closures have assumed that habitat impacts fish population growth, only Alqawasmeh and Lutscher (2019) consider that habitat might affect fish movement behavior. Our approach differs slightly from Alqawasmeh and Lutscher (2019), who do not explicitly model habitat quality, but rather assume that fishing activity may instantly affect fish movement behavior (by increasing the movement rate outside of protected areas

Fig. 1. Summary of the steps in our modeling framework.

Model Framework Summary: Starting at the beginning of year d,

1. The manager selects an action $\mathbf{a}(d)$, representing which patches are open/closed for the duration of year d
2. The dynamics of the fishery are simulated for times $t \in [d, d+1)$
3. The manager receives Reward(d), representing the dual goals of maximizing harvest and maximizing habitat quality
4. Advance to year $d+1$ and return to step 1

and/or bias toward protected areas). However, both our approaches maintain the same key assumption that fish may prefer less disturbed areas.

2.2.2. Stock dynamics

The growth equation represents the instantaneous change to the stock in patch i through recruitment and natural mortality.

$$\text{Growth}_i(t) = \underbrace{\left(\frac{A_i^m(t)}{\sum_i A_i^m(t)}\right)}_{\substack{\text{recruit spatial}\\ \text{distribution}}} \underbrace{\frac{r\sum_i x_i(t)}{1+\frac{\kappa\sum_i x_i(t)}{\sum_i A_i^s(t)}}}_{\substack{\text{total}\\ \text{recruitment}}} - \underbrace{\mu x_i(t)}_{\substack{\text{natural}\\ \text{mortality}}} \tag{6}$$

We will assume that (1) recruitment happens throughout the year (and is not affected by seasonality), and (2) the early life stages disperse widely and function as a common juvenile pool. The "total recruitment" term represents the total instantaneous recruitment for the entire fishery coming from this common pool. The *potential* total instantaneous recruitment is given by $r\sum_i x_i(t)$, where r is the stock growth rate. *Actual* total instantaneous recruitment is reduced through density-dependent effects, represented by the denominator of the "total recruitment" term, where κ is a scaling parameter. As we assumed a common juvenile pool model, to represent the habitat–stock survival effect, we let the survival of early life stages depend on the status of the entire fishery aggregated across all spatial patches (represented by $\sum_i A_i^s(t)$; see eq. 4). As $\sum_i A_i^s(t)$ increases, survival (and, subsequently, recruitment) increases.

The total instantaneous recruitment is then distributed between patches by the "recruit distribution" term. To incorporate the habitat–stock movement effect for recruits, we assume that recruits settle in patches in proportion to their scaled patch preferences ($A_i^{\mathrm{m}}(t)$; see eq. 5), upon which they enter the adult population. Finally, we assume that the natural mortality of the stock occurs at a constant rate μ.

The movement equation represents the net movement of the stock (i.e., adults) in/out of patch i. The first term gives the instantaneous movement out of patch i, while the second gives the total instantaneous movement into patch i.

$$\text{Movement}_i(t) = \alpha\left(\underbrace{Y_{i,i}\frac{x_i(t)}{A_i^m(t)}}_{\substack{\text{movement out}\\ \text{of patch i}}} + \sum_{\substack{j=1\\ j\neq i}}^{n} \underbrace{Y_{i,j}\frac{x_j(t)}{A_j^m(t)}}_{\substack{\text{movement to i}\\ \text{from other patches}}}\right) \tag{7}$$

Here α is the stock movement rate parameter, and $Y_{i,j}$ are movement weights. To represent the habitat–stock movement effect for a mobile stock ($\alpha > 0$), we assume that adults do not have any information about the habitat quality in other patches, and make the decision to leave a patch based only on its scaled patch preference ($A_i^m(t)$; see eq. 5). Specifically, we let the rate of adult movement out of a patch be inversely proportional to $A_i^{\mathrm{m}}(t)$, meaning that mobile adults are less likely to leave a patch with a higher $A_i^{\mathrm{m}}(t)$. On the other hand, when modeling a sedentary stock, we let $\alpha = 0$, in which case no movement of adults occurs (note, though, that recruits are still assumed to be mobile; see eq. 6).

The weights $Y_{i,j}$ determine the pattern in which mobile adults move between patches (representing the geometry of patches, directional biases, etc.). For $i \neq j$, $Y_{i,j}$ gives the proportion of fish leaving patch j that enter patch i. On the other hand, $Y_{i,i}$ represents the movement out of patch i. In our analysis, we use the following matrix of $Y_{i,j}$ values:

$$\mathbf{Y} = \begin{bmatrix} -0.5 & 0.5 & 0 & 0 & 0 \\ 0.5 & -1 & 0.5 & 0 & 0 \\ 0 & 0.5 & -1 & 0.5 & 0 \\ 0 & 0 & 0.5 & -1 & 0.5 \\ 0 & 0 & 0 & 0.5 & -0.5 \end{bmatrix} \tag{8}$$

This represents a fishery with $n = 5$ patches in a line, fish movement only allowed between neighboring patches, and no directional bias. Essentially, of the fish leaving each patch, half will end up in each neighboring patch (reflecting our assumption that fish don't know the habitat quality of a patch until entering that patch). This matrix of weights also assumes a reflecting outer boundary (half the fish that attempt to leave an end patch instead hit the habitat boundary and do not leave the patch).

Finally, the harvest equation represents the harvest of the stock in patch i by the fishing effort in patch i. For this we assume a simple type I functional response (the Schaefer catch equation):

$$\text{Harvest}_i\,(t) = qa_i\,(t)\,x_i\,(t)\,f_i\,(t) \tag{9}$$

where q is the catchability parameter. Multiplication by the management action $a_i(t)$ is included to emphasize that no harvest occurs in a closed patch (when $a_i(t) = 0$).

2.2.3. Habitat dynamics

We assume that damaged habitat is able to recover in quality over time, which we represent with the recovery equation. Similar to Smeltz et al. (2019), we assume that habitat recovery is of an exponential asymptotic form,

$$\text{Recovery}_i\,(t) = \beta\,(h_{\max} - h_i\,(t)) \tag{10}$$

where β is the habitat recovery rate and $h_{\max}$ is the maximum habitat quality that can occur in a given patch.

The damage equation represents damage to each patch by the fishing effort in that patch. As with the harvest term for stock, we assume a type I functional response for habitat damage caused by fishing activity. This is the same functional form as used in past works that explicitly model habitat and habitat damage (Nichols et al. 2018; Kelly et al. 2019).

$$\text{Damage}_i\,(t) = \gamma\,a_i\,(t)\,h_i\,(t)\,f_i\,(t) \tag{11}$$

The habitat impact parameter γ represents how damaging the fishing gear is to habitat (or equivalently, how sensitive the habitat is to damage). A value of $\gamma = 0$ implies that no habitat damage occurs. Multiplication by the management action $a_i(t)$ is again included to emphasize that no damage occurs in a closed patch (when $a_i(t) = 0$).

2.2.4. Effort dynamics

We assume that fishing effort follows open-access dynamics in open patches, and that no effort is permitted within closed patches (the closures are strictly "no-take"). This means that effort will enter an open patch from outside the fishery when it is profitable to do so, and leave when it is not. Let $C_i(t) = qa_i(t)x_i(t)$ be the instantaneous catch per unit effort (CPUE) in patch i, p be the price per unit harvest, and c be the cost per unit effort (we multiply CPUE by the management action to ensure there is no catch allowed in a closed patch, where $a_i(t) = 0$). Then, the instantaneous average profit in patch i can be written as $P_i(t) = pC_i(t) - c$.

The classic open-access model (Smith 1969; Sanchirico and Wilen 2001) would assume that effort enters and exits an open patch from outside the fishery according to the equation $\mathrm{d}f_i/\mathrm{d}t = P_i\,(t)\,f_i\,(t)$. However, this form poses issues when it comes to modeling patch closures and reopenings, as if there is no effort in a patch ($f_i(t) = 0$), then $\mathrm{d}f_i/\mathrm{d}t = 0$ always. This means that when a closed patch is reopened, no effort would move into it from outside the fishery, even if the potential CPUE was very high. We rectify this issue by making slight changes to the classic open-access equation while maintaining the key aspects of its dynamics. Notably, we want our model to satisfy the property that effort will enter an open patch from outside the fishery when $P_i(t) > 0$, and exit when $P_i(t) < 0$. We thus define the entry and exit terms for effort in our model by

$$\text{Entry}_i\,(t) = \lambda_{\text{ent}} P_i^+\,(t) \tag{12}$$

$$\text{Exit}_i\,(t) = \lambda_{\text{exit}} P_i^-\,(t)\,f_i\,(t) \tag{13}$$

where $P^+ = \max\,(P,\ 0), P^- = -\min\,(P,\ 0)$, and $\lambda_{\text{ent}},\ \lambda_{\text{exit}}$ represent respective rates for entry and exit. We assume that when effort entry occurs, it is independent of the amount of effort currently present in the patch. This allows effort to enter into recently reopened patches from outside the fishery. In the case of effort exiting the fishery, our form matches the classic open-access model exactly.

Additionally, we allow effort already in the fishery to quickly move between patches. We assume that fishers will move between patches in response to $C_i(t)$ (CPUE), and define the relocation term for effort as

$$\text{Relocation}_i\,(t) = \lambda_{\text{rel}} \sum_{j \neq i} \left[C_{i,j}^+\,(t)\,f_j\,(t) - C_{i,j}^-\,(t)\,f_i\,(t) \right] \tag{14}$$

where $C_{i,j}(t) = C_i(t) - C_j(t)$ gives the difference in CPUE between patch i and patch j, λ_{rel} represents the rate of effort movement between patches, and $C^+ = \max\,(C,\ 0), C^- = -\min\,(C,\ 0)$. Essentially, if $C_i(t) > C_j(t)$, effort will move from patch j to patch i (and vice versa). Unlike the stock, we do not constrain effort to move only between neighboring patches. Additionally, we assume that $\lambda_{\text{rel}} > \lambda_{\text{ent}},\ \lambda_{\text{exit}}$, such that movement between patches occurs more quickly than exiting/entering the fishery does.

Finally, our model must account for what happens to effort in patches that become closed. When a patch containing effort is closed by the manager, we assume that the effort is instantaneously redistributed into patches that are open, representing how real spatial closures often result in effort shifting into other areas (Agardy et al. 2011). As a simple way of representing this displacement, we redistribute effort from newly closed patches to open patches in proportion to the amount of effort currently in each open patch.

2.3. Management goals and rewards

The reward given to the manager for year d is a weighted sum of the total harvest and total habitat quality over the course of the year,

$$\text{(15)}\quad \text{Reward}\,(d) = \int_{d}^{d+1} w_{\text{harv}} \left[\sum_{i=1}^{n} \text{Harvest}_i\,(t)\right] + w_{\text{hab}} \left[\sum_{i=1}^{n} h_i\,(t)\right]\, dt$$

The manager wishes to maximize the total discounted reward (with discount factor δ) on an infinite time horizon. Notably, the total discounted reward can also be written as a weighted sum of the total discounted harvest and the total discounted habitat. Both forms are shown in eq. 16.

$$\text{(16)}\quad \text{Total discounted reward} = \sum_{d=0}^{\infty} \delta^{d}\ \text{Reward}\,(d)$$

$$= w_{\text{harv}} \underbrace{\sum_{d=0}^{\infty} \delta^{d} \left(\int_{d}^{d+1} \sum_{i=1}^{n} \text{Harvest}_i\,(t)\ dt\right)}_{\text{total discounted harvest}} + w_{\text{hab}} \underbrace{\sum_{d=0}^{\infty} \delta^{d} \left(\int_{d}^{d+1} \sum_{i=1}^{n} h_i\,(t)\ dt\right)}_{\text{total discounted habitat}}$$

Since we cannot actually simulate an infinite time horizon, we estimate the total discounted reward on a large, finite time horizon t_{f}. For our selected discount factor of $\delta = 0.95$, we found $t_{\text{f}} = 150$ years to be sufficient.

In investigating spatial closure policies, we will consider both conventional single-objective harvest maximization ($w_{\text{harv}} = 1$, $w_{\text{hab}} = 0$) and multiobjective optimization with nonzero weights on both harvest and habitat. While our optimization focus is on maximizing yield in the single-objective setting, we still pay attention to the habitat conditions occurring under each policy. Notably, some policies may result in more or less habitat degradation than others, even though habitat protection was not explicitly included in the optimization goal. We call this a "secondary habitat benefit". This secondary benefit could, for instance, be seen as a tie-breaker for policies that perform similarly in yield.

In the multiobjective setting, our general approach will be to vary w_{hab} over a range of values in [0,1] (with $w_{\text{harv}} = 1 - w_{\text{hab}}$), and look at the resulting policies produced in each scenario. Plotting the total discounted harvest vs. total discounted habitat for each policy allows us to visualize the range of possible outcomes as different management goals are prioritized. We call a policy “Pareto optimal” if no other policy of the same class (e.g., comparing DSC policies with DSC policies) produces both a larger discounted harvest and a larger discounted habitat (or ties in one and is larger in the other). Connecting all the Pareto optimal policies with a line allows us to visualize the “Pareto frontier”, a representation of the outer boundary of what can be obtained using only that class of policy. Rather than deciding on exact mathematical weights (w_{hab} and w_{harv}) to define an optimization goal, a manager could visually inspect a Pareto plot, and after considering different trade-offs, pick the "best" policy for their fishery.

Finally, we also briefly consider a stochastic setting, in which there is variability in yearly recruitment. In this setting, the manager’s goal becomes to maximize the *expected* total discounted reward. The procedure for this setting is further discussed in Supplementary Material Section S3.

2.4. Summary of optimization methods

For DSCs, we define a policy as a function that outputs the manager’s action (which patches are opened/closed) at each state. Ideally, we would like to find a DSC policy that maximizes the total discounted reward. In simpler models (e.g., where there are only a finite number of possible states), this could be done using an algorithm such as value iteration (Sutton and Barto 2018). However, due to the complexity of our model, we must rely on approximate optimization methods instead, with the goal of finding near-optimal DSC policies. Specifically, we use an approximate dynamic programming (ADP) algorithm called “sampling-based fitted value iteration” (as described in Munos and Szepesvári 2008). While classic value iteration proceeds in iterations that require calculations at every state in the state space, sampling-based fitted value iteration calculates each iteration at a sample of states, and then uses a functional approximation to generalize to other states. This ADP algorithm can thus be applied to models with large and continuous state spaces, such as ours. Interested readers may refer to Appendix Section A.1 and Fig. A1 for a more thorough description of the algorithm and our implementation details, but these specifics are not necessary to understand our results. Because the ADP algorithm involves random sampling, each run will produce slightly different DSC policies; however, we found that the difference in policy performance between runs was generally very small (see Section A.1).

All simulations and optimization were performed in Microsoft R Open 4.0.2 (R Core Team 2020). ODEs were solved with the ode function from the deSolve package, using the Adams method (Soetaert et al. 2010). For the functional approximation in the ADP algorithm, we used a GAM (generalized additive model), implemented with the mgcv package (Wood 2017).

2.5. Parameters

We will examine a fishery consisting of $n = 5$ patches in a line. The key parameters that we will vary in our analysis are the strength of the habitat–stock survival effect (ϕ_{s}) and the strength of the habitat–stock movement effect (ϕ_{m}), along with the stock movement rate α (to represent either a sedentary stock or a mobile stock). Although we vary ϕ_{s}, ϕ_{m} over several values, throughout our analysis we pay particular attention to two scenarios: "no habitat–stock effects" (meaning $\phi_{\text{s}} = \phi_{\text{m}} = 0$) and "very strong habitat–stock effects" (meaning $\phi_{\text{s}} = \phi_{\text{m}} = 1$).

Other parameters will remain fixed in the "base case", which will be used for the majority of our analysis. Table 1 gives a summary of our model parameters and the chosen value(s) for each in the base case. We selected the stock parameter values to represent a typical harvested marine taxa where the early life stages contribute to a common juvenile pool. Similarly, habitat parameter values were selected to be moderate to represent a typical habitat impacted

Table 1. A summary of model parameters and their selected values in the base case.

Parameter	Meaning	Base case value(s)	Units
n	Number of patches	5 (arranged in a line)	patches
ϕ_s	Strength of habitat–stock survival effect	Varied in {0, 0.5, 0.8, 1}	unitless
ϕ_m	Strength of habitat–stock movement effect	Varied in {0, 0.5, 0.8, 1}	unitless
r	Growth parameter for stock	0.8	year^{-1}
κ	Recruitment term scaling parameter	*	unitless
μ	Stock natural mortality rate	0.2	year^{-1}
α	Stock movement rate	0 (sedentary), 1 (mobile)	year^{-1}
$\Upsilon_{i,j}$	Weights for fish movement	see Section 2.2	unitless
q	Catchability parameter	*	year^{-1}
β	Habitat recovery rate	0.2	year^{-1}
$h_{\max}$	Maximum (per patch) habitat quality	*	unitless
γ	Habitat impact parameter	*	year^{-1}
p	Price per unit harvest	*	unitless
c	Cost per unit effort	1	year^{-1}
λ_{ent}	Rate of effort entry into fishery	10	unitless
λ_{exit}	Rate of effort exit from fishery	10	unitless
λ_{rel}	Rate of effort movement between patches	50	unitless
x_i^*	Stock open-access equilibrium	0.1	unitless
h_i^*	Habitat open-access equilibrium	0.25	unitless
f_i^*	Effort open-access equilibrium	1	unitless
$\tilde{x}_i^*$	Stock no-fishing equilibrium	1	unitless
$\tilde{h}_i^*$	Habitat no-fishing equilibrium	1	unitless
$\tilde{f}_i^*$	Effort no-fishing equilibrium	0	unitless
t_f	Time horizon used to estimate infinite horizon reward	150	years
δ	Discount factor	0.95	unitless
w_{hab}	Reward weight for habitat protection	0 (single objective), varied in [0, 1] (multiobjective)	unitless
w_{harv}	Reward weight for harvest	$1 - w_{\text{hab}}$	unitless

Note: A "*" under "base case values(s)" indicates that the parameter is found analytically using the equations in Section A.3 such that the chosen values for the equilibria ($x_i^*, h_i^*, f_i^*, \tilde{x}_i^*, \tilde{h}_i^*, \tilde{f}_i^*$) result.

Table 2. Key parameters changed for each of the four alternative cases explored.

Case description	r	β	h_i^* (γ)
Base case	0.8	0.2	0.25 (0.6)
(C1) Slower habitat recovery	0.8	0.1	0.1429 (0.6)
(C2) Less damaging gear	0.8	0.2	0.4 (0.3)
(C3) Slower habitat recovery and less damaging gear	0.8	0.1	0.25 (0.3)
(C4) Decreased stock productivity	0.45	0.2	0.25 (0.6)

Note: The remaining parameters take their base case values, given in Table 1. r is the stock growth parameter, β is the habitat recovery rate, and h_i^* is the open-access habitat equilibrium (from which the habitat impact parameter γ is calculated using eq. A6).

by bottom-towed fishing gear. We will also consider several alternate cases in which some key parameters will differ from the base case (see Table 2). In the alternative cases, we vary the habitat recovery rate, the damage caused by fishing gear (equivalently, the sensitivity of the habitat), and the productivity of the stock. Additional details about parameter selection for the base case and alternative cases are given in Section A.2.

Two sets of equilibria are of central importance to our model: the open-access equilibria (x_i^*, h_i^*, f_i^*), occurring when all patches are open, and the no-fishing equilibria ($\tilde{x}_i^*$, $\tilde{h}_i^*$, $\tilde{f}_i^*$), occurring when all patches are closed. The formulas for each equilibria are given and discussed in Appendix Section A.3. We fix the values for these equilibria, and then solve for some parameters (those marked with a "$*$" in Table 1) to result in the chosen values. For example, the habitat impact parameter γ is determined based on the value for the open-access habitat equilibrium h_i^* using eq. A6. The base case values for each equilibria are listed in Table 1. Note that we set x_i^* to a small value to represent that the fishery is overfished under open-access conditions. Similarly, h_i^* is set to a small value to represent that continually fishing at open-access effort levels is very damaging to the seafloor. The open-access equilibrium will also serve as our initial state in simulations (in other words, we assume that the fishery was previously unmanaged): $x_i(0) = x_i^*$, $h_i(0) = h_i^*$, $f_i(0) = f_i^*$.

2.6. Comparison with other policy classes

Section 2.4 discussed our procedure for finding near-optimal policies of DSCs. Our analysis will focus on how these DSC policies compare in performance with other classes of policies, which are described here.

2.6.1. Marine protected areas

For MPAs (i.e. fixed spatial closures), we assume that the manager decides which patches are closed at $t = 0$, and does not change the closures throughout the entire time horizon. In an n-patch model, there are $2^n - 1$ possible configurations of MPAs (since the manager is not allowed to close all patches), so we simply simulate each one and select the configuration that maximizes the total discounted reward. Comparing our DSC policies with the best MPA will let us identify when significant gains could be achieved by allowing spatial closures to change over time.

As a quick note, theoretically, our DSC policies should always perform at least as well as the best MPA, because making decisions identical to the best MPA is always an option. However, because ADP is an *approximate* optimization algorithm, there were a few scenarios where the best MPA slightly outperformed the best DSC policy.

2.6.2. Heuristic spatial policies

While our DSC policies represent an upper bound for the performance of (yearly changing) spatial closures, they may be quite challenging to implement in real life. These policies require a large amount of information to develop, which may not be available, especially for fisheries with limited management resources available. Additionally, they may produce quite complicated patterns of spatial closures, which may be undesirable from a management perspective. Thus, we also consider some simple heuristic policies of rotating spatial closures on a fixed schedule. These heuristic policies rotate through all n patches in order with one patch open at a time; each patch is open for a fixed length of y^{rot} years (and closed otherwise). Since the fishery begins in a depleted state, an initial period of y^{rec} years is also included to give the fishery time to recover before the rotations begin. For the first y^{rec} years, one patch is permanently open and the rest are closed (as not all n patches are allowed to be closed at once). We vary $y^{\text{rot}} \in \{1, 2, 3\}$ and $y^{\text{rec}} \in \{2, 3, \ldots, 8\}$ to get an assortment of reasonable rotational policies, and compare their harvest and habitat protection performance to the best DSC policy and the best MPA. Beyond these simple heuristic policies of rotating spatial closures, we also briefly examine how heuristic policies could utilize additional information about the optimal DSC policies to achieve improved performance.

2.6.3. Nonspatial maximum sustainable yield

Finally, as a performance benchmark, we also compare our DSC policies to nonspatial management at MSY harvest rates. Maximum sustainable yield is the largest catch that can be sustained in a fishery indefinitely (Maunder 2008). For our nonspatial MSY benchmark, we assume that the manager has total control over equilibrium effort, but manages the fishery in a nonspatial manner (the effort level in every patch is the same). Note that under this management regime, the fishery is *no longer open-access*. The calculation of MSY is shown in Section A.4.

To allow for a direct comparison to our nonspatial MSY benchmark, we will consider the *average* harvest and habitat produced by our DSC policies, taken over the time interval $[10, t_f]$ years. Although our DSC policies are trained to maximize the total discounted reward rather than the average reward, we find that they still perform well in the latter. As our DSC policies assume an initially depleted fishery, we ignore the first 10 years to allow the fishery time to recover.

3. Results

3.1. Maximizing yield: base case

3.1.1. Demonstration and characteristic policies

We begin by considering a single-objective management goal, where the manager's only aim is to maximize yield ($w_{\text{hab}} = 0$, $w_{\text{harv}} = 1$), under base case parameters. To demonstrate our procedure and explore the characteristics of the resulting policies, Fig. 2 presents the best DSC and MPA in four scenarios: a sedentary stock with no habitat–stock effects ($\phi_{\text{m}} = \phi_{\text{s}} = 0$), a sedentary stock with very strong habitat–stock effects ($\phi_{\text{m}} = \phi_{\text{s}} = 1$), a mobile stock with no habitat–stock effects, and a mobile stock with very strong habitat–stock effects.

We found that the DSC policies generally had two distinct sections: an initial "recovery" period lasting a few to several years, followed by a "rotational" period lasting the remainder of the time horizon. During the "recovery" period, a single patch was permanently open and the rest were closed, giving the fishery a chance to recover from its initially depleted state. (Recall that the manager was not allowed to close all n patches at once.) The recovery period was longer for the mobile stock, and in the very strong habitat–stock effects scenarios. During the "rotational" period, which patch(es) were open changed through time. The length of the openings in the rotational period varied from 1 year to several years, as did the time between openings. It was common for only one patch to be open at a time during the rotational period (although see Fig. 2D). This is likely due to our chosen parameter values: notably, we assumed that the stock is overfished under open-access conditions, and thus keeping more patches closed at once to maintain a larger stock size may often be optimal. Keeping more patches closed also results in more habitat protection (and subsequently, greater recruitment if $\phi_{\text{s}} > 0$). The strength of the habitat–stock effects and the mobility of the stock both qualitatively affected the behavior of the optimal DSC policies in the rotational period. For the sedentary stock (Figs. 2A and 2B), very strong habitat–stock effects generally increased the number of years each patch was closed between openings, and the number of years each patch was consecutively open. For the mobile stock (Figs. 2C and

Fig. 2. The best dynamic spatial closure (DSC) policy from single-objective optimization (harvest only; $w_{\mathrm{hab}} = 0$), and corresponding best marine protected area (MPA), in four stock mobility/habitat effects scenarios. The behavior of the best DSC and MPA was strongly affected by both the mobility of the stock and the strength of the habitat–stock effects. "No habitat–stock effects" implies $\phi_{\mathrm{m}} = \phi_{\mathrm{s}} = 0$, while "very strong habitat–stock effects" implies $\phi_{\mathrm{m}} = \phi_{\mathrm{s}} = 1$. Other parameters take their base case values (Table 1). Only the first 50 years of decisions are shown for clarity.

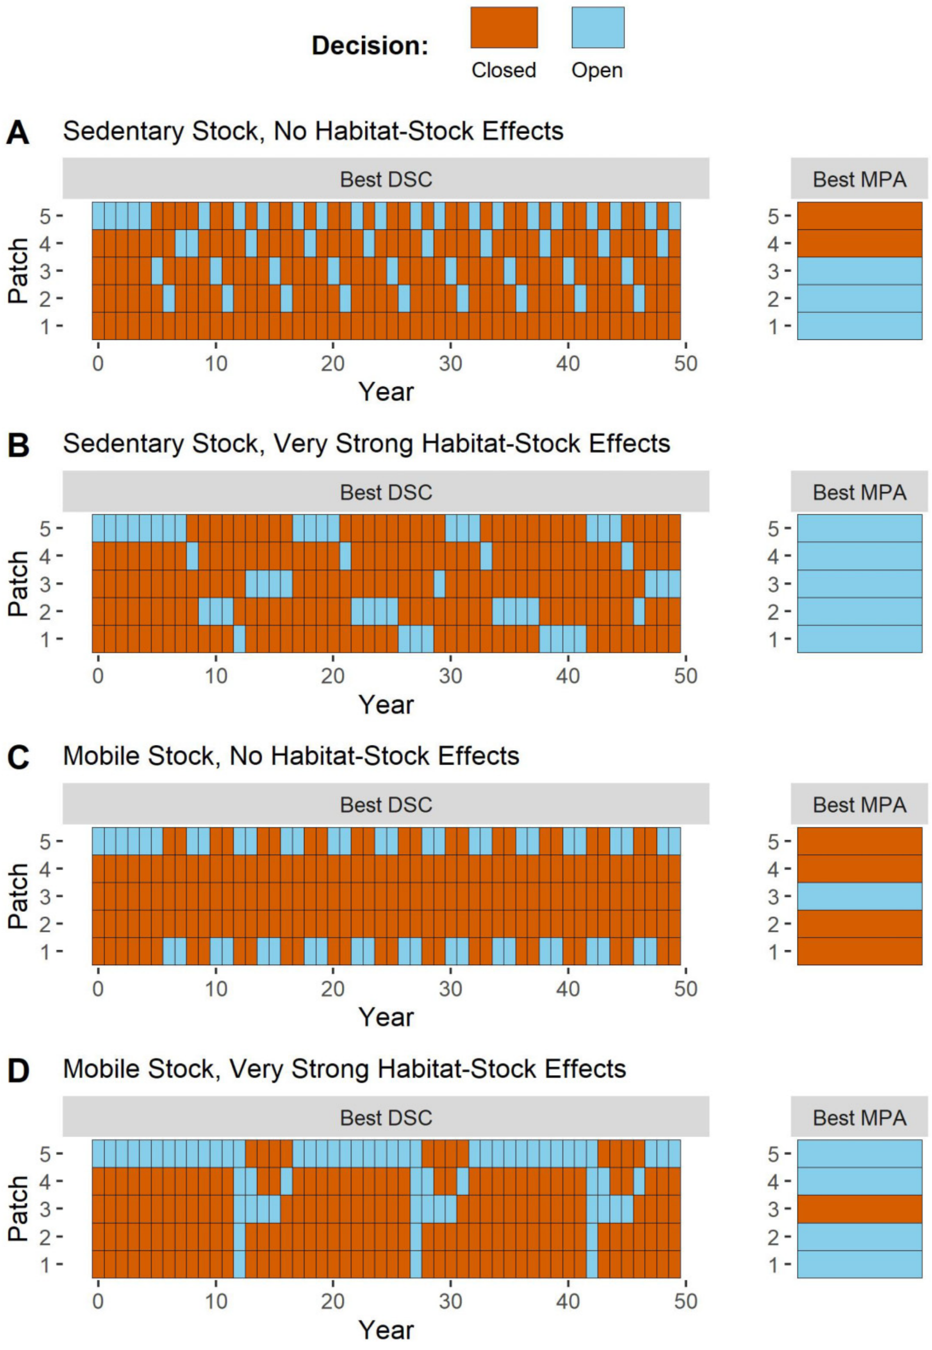

2D), very strong habitat–stock effects shifted the optimal DSC policy from one with several permanent closures to a pulse fishing strategy, with the entire fishery opened once every several years. Finally, the strength of the habitat–stock effects and the mobility of the stock also affected the best MPA configuration. More patches were closed in the optimal MPA for the mobile stock, and in the no habitat–stock effect scenarios (Fig. 2).

3.1.2. Relative performance across policies

For both yield and secondary habitat benefits, the optimality of DSCs versus MPAs generally increased as the strengths of the habitat–stock effects (ϕ_{s}, ϕ_{m}) increased (Figs. 3A and 3B). When there was no effect of habitat on stock movement ($\phi_{\mathrm{m}} = 0$), the best DSC policy achieved similar or only slightly larger harvests than the best MPA (Fig. 3A). However,

Fig. 3. Relative performance of dynamic spatial closures (DSCs) under base case parameters (Table 1) for a single-objective optimization goal (harvest only; $w_{\mathrm{hab}} = 0$). The optimality of DSCs versus marine protected areas (MPAs) and the nonspatial maximum sustainable yield (MSY) benchmark strongly depends on the strengths of both habitat–stock effects (survival and movement). (A) Ratio of discounted harvest of the best DSC policy to the best MPA. (B) Ratio of discounted habitat of the best DSC policy to the best MPA. (C) Ratio of postrecovery average harvest of the best DSC policy to nonspatial MSY. (D) Ratio of postrecovery average habitat of the best DSC policy to equilibrium habitat under nonspatial MSY management.

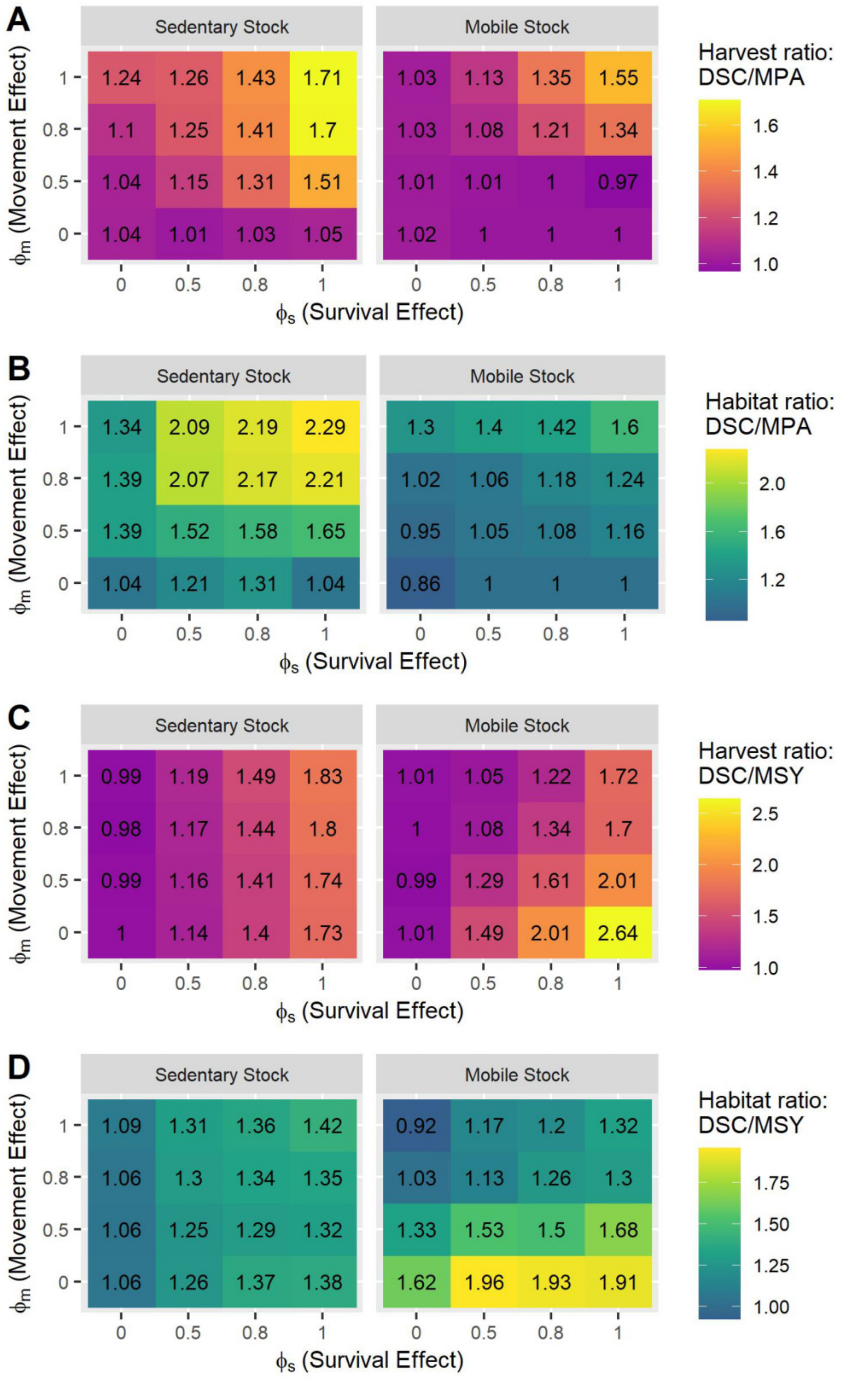

when the habitat–stock movement effect was present ($\phi_m > 0$), DSCs were often able to produce significantly larger discounted harvests: up to 71% (sedentary stock) and 55% (mobile stock) greater than that produced by the best MPA (Fig. 3A). Additionally, DSCs achieved larger secondary habitat benefits than MPAs in many scenarios, notably for the sedentary stock (Fig. 3B). DSCs resulted in discounted habitat up to 129% (sedentary stock) and 60% (mobile stock) greater than that resulting from the best MPA (Fig. 3B).

We observed that when fish have a strong preference for habitat, they tended to aggregate in areas that had been closed for a while (and therefore had a high habitat quality), potentially decreasing yield if they were unreachable by fishers (Fig. S1). This was a key driver in the performance difference between DSCs and MPAs. DSCs were able to allow access to areas with high densities of fish by changing the placement of closures through time, helping to alleviate this problem. On the other hand, using MPAs (i.e. fixed spatial closures) resulted in a trade-off between closing more patches to protect more of the stock and habitat, and closing fewer patches and making fish more accessible to harvest. The optimal number of patch closures in the best (harvest-maximizing) MPA decreased as the strength of the habitat–stock movement effect increased, with this pattern especially strong for the mobile stock (Fig. S2). In fact, there were some scenarios (e.g., Fig. 2B) for which the strong habitat–stock movement effect made it so that the best MPA configuration was to close no patches at all.

Now comparing DSCs with our nonspatial MSY benchmark, we find that for both yield and secondary habitat benefits, the optimality of DSCs generally increased as the strength of the habitat–stock survival effect (ϕ_s) increased (Figs. 3C and 3D). When there was no effect of habitat on survival ($\phi_s = 0$), DSCs could match but not increase yield versus nonspatial MSY (Fig. 3C). However, when habitat did affect survival ($\phi_s > 0$), DSCs could produce average postrecovery harvests greater than nonspatial MSY (up to 83% for the sedentary stock and 164% for the mobile stock; Fig. 3C). The ability of DSCs to protect more habitat than nonspatial management at MSY harvest rates explains why DSCs were able to produce average yields greater than nonspatial MSY when the habitat–stock survival effect was present. DSCs were able to achieve more habitat protection than our nonspatial MSY benchmark in most cases: up to 42% (sedentary stock) and 96% (mobile stock) greater (Fig. 3D). On the other hand, the habitat–stock movement effect had varied impacts on the optimality of DSCs versus nonspatial MSY (with respect to both yield and habitat protection). For the sedentary stock, the habitat–stock movement effect had little impact on the relative performance of DSCs; for the mobile stock, the relative performance of DSCs generally increased as the strength of the habitat–stock movement effect decreased (Figs. 3C and 3D).

3.1.3. Heuristic spatial policies

In scenarios with no habitat–stock effects ($\phi_m = \phi_s = 0$), our heuristic policies of rotating spatial closures achieved similar harvests as the best DSC and best MPA, although the amount of habitat protection varied more greatly (Fig. S3). On the other hand, under very strong habitat–stock effects ($\phi_m = \phi_s = 1$), our heuristic rotational policies often achieved less harvest than the best DSC (upwards of 25% less in the worst case), especially for the mobile stock (Figs. 4A and 4B). The heuristic rotational policies also often achieved less habitat protection under very strong habitat–stock effects (Figs. 4A and 4B). Despite this, all heuristic policies achieved notably better performance than the best MPA under very strong habitat effects (Figs. 4A and 4B). This suggests that even in information-limited systems where it is hard to identify and implement optimal dynamic patterns of spatial closures, simple policies of rotating spatial closures may be able to capture many of their benefits.

We also investigated how using additional information from the DSC policies could produce better heuristic policies. For example, the best DSC policy for the "mobile stock, very strong habitat effects" scenario (Fig. 2D) suggests a pulse fishing strategy, as opposed to a rotational strategy. We thus consider a range of heuristic pulse fishing policies for this scenario, wherein all patches are open to fishing every y^{pulse} years, and between that a single patch is open (to meet the management requirement that not all patches can be closed at once). We vary $y^{\text{pulse}} \in \{8, 10, \ldots, 20\}$ years, representing a wide range of pulse lengths around that of the best DSC policy (15 years, Fig. 2D). An initial recovery period of y^{rec} years is again included, but since the best DSC policy suggests a longer recovery period (12 years, Fig. 2D), we vary $y^{\text{rec}} \in \{8, 9, \ldots, 16\}$ years. Compared with the heuristic rotational policies, we observe improved performance across this range of heuristic pulse policies, with many achieving similar or even slightly better harvests than the best DSC (Fig. 4C). This "informed heuristics" approach demonstrates how optimality analysis might be used to guide simple heuristic policies that would be easier to implement and manage.

3.1.4. Stochastic setting

Finally, we briefly consider a stochastic environment with variable recruitment (described in Supplementary Material Section S3). In this setting, we compared the performance of DSC policies that take recruitment variability into account when making decisions with DSC policies that ignore recruitment variability. Ultimately, we found that the two policy types performed similarly in the stochastic setting (Fig. S4), indicating that it was not crucial to take this type of variability into account for decision-making. The results for the stochastic setting are further discussed in Section S3.

3.2. Maximizing yield: alternative cases

Here we briefly summarize the main results found by exploring the four alternative cases outlined in Table 2. Additional details are given in Supplementary Material Section S4 and Figs. S5–S8. First, we found that slower habitat recovery rates had little impact on the optimality of DSCs versus MPAs, but in most cases significantly increased the performance of DSCs versus the nonspatial MSY management benchmark. On the other hand, less damaging gear decreased the benefits of using DSCs versus both MPAs and nonspatial MSY management. Finally, we found that decreasing stock productivity had little effect on the relative optimality of the different policy classes.

3.3. Multiobjective management

Now we turn our attention to a multiobjective management goal, where both harvest and habitat protection can have non-zero weight (w_{harv}, w_{hab}) in the reward (eq. 16). Here we focus our attention on two scenarios under base case parameters (Table 1): a sedentary stock with no habitat–stock

Fig. 4. The performance of several heuristic policies and the best marine protected area (MPA), relative to the performance of the best dynamic spatial closure (DSC) policy, for a single-objective optimization goal (harvest only; $w_{\text{hab}} = 0$) under very strong habitat–stock effects ($\phi_s = \phi_m = 1$). Other parameters take their base case values (Table 1). (A, B) A range of simple rotational heuristic policies were generated by varying the initial recovery period and the length of patch openings (y^{rec} and y^{rot}; see Section 2.6.2). These heuristic policies sometimes achieve smaller discounted harvests than the best DSC policy, but still notably outperform the best MPA. (C) Since the best DSC policy for a mobile stock with very strong habitat–stock effects (Fig. 2D) suggests a pulse fishing strategy with a longer recovery period, a range of heuristic pulse policies were tested, varying the initial recovery period and how often pulses occur (y^{rec} and y^{pulse}; see Section 3.1.3). By using this additional information, the heuristic policies were able to achieve similar performance to the best DSC.

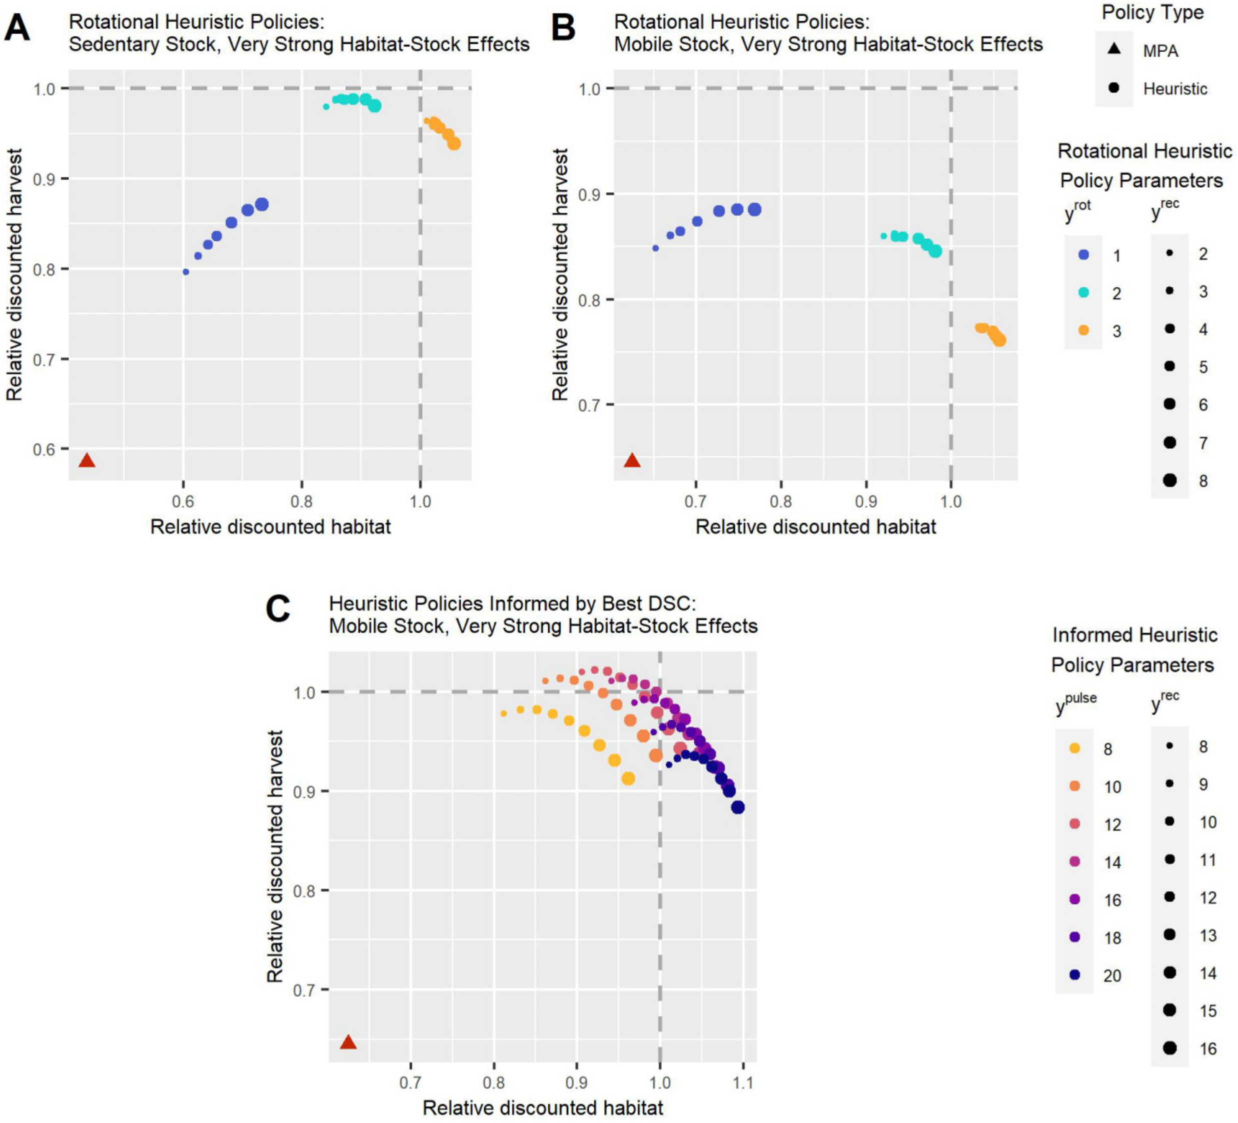

effects ($\phi_m = \phi_s = 0$) and a sedentary stock with very strong habitat–stock effects ($\phi_m = \phi_s = 1$).

For a sedentary stock with no habitat–stock effects, no matter the weights used for habitat protection (w_{hab}) versus harvest (w_{harv}), DSCs are only able to achieve marginal improvements over the best MPA (Fig. 5A). This matches our observations from the single-objective management case. On the other hand, for a sedentary stock with very strong habitat–stock effects, we find that the manager can obtain significant benefits by using DSCs instead of MPAs across a range of weights (Fig. 5B). Matching the single-objective case, when the management goal is mostly focused on harvest maximization (w_{hab} is small), DSC policies significantly outperform MPAs. As the manager's goal shifts more toward protecting habitat, however, DSCs begin to lose their edge (the optimality of DSCs versus MPAs decreases as the habitat weight w_{hab} increases). In extreme situations where the habitat weight is sufficiently large, the optimal DSC policy is in fact an MPA. Still, by varying the weights, we see that DSCs are able to offer a variety of options that provide different trade-offs between yield and habitat protection.

Fig. 5. Pareto plots summarizing the results of multiobjective optimization (where both harvest and habitat protection are included as management goals). Each circle represents a dynamic spatial closure (DSC) policy generated under a different set of weights (w_{hab}, w_{harv}). Triangles represent all possible marine protected areas (MPAs); note that MPAs with the same number of patches closed overlap in these examples. A policy is "Pareto optimal" if no other policy of the same class obtains both a larger total discounted harvest and total discounted habitat. To represent the Pareto frontiers (which show the trade-offs between the two management goals), Pareto optimal MPAs are connected with a dashed line while Pareto optimal DSCs are connected with a solid line. Both panels consider a sedentary stock with base case parameters (Table 1). (A) No habitat–stock effects ($\phi_{\mathrm{m}} = \phi_{\mathrm{s}} = 0$). (B) Very strong habitat–stock effects ($\phi_{\mathrm{m}} = \phi_{\mathrm{s}} = 1$).

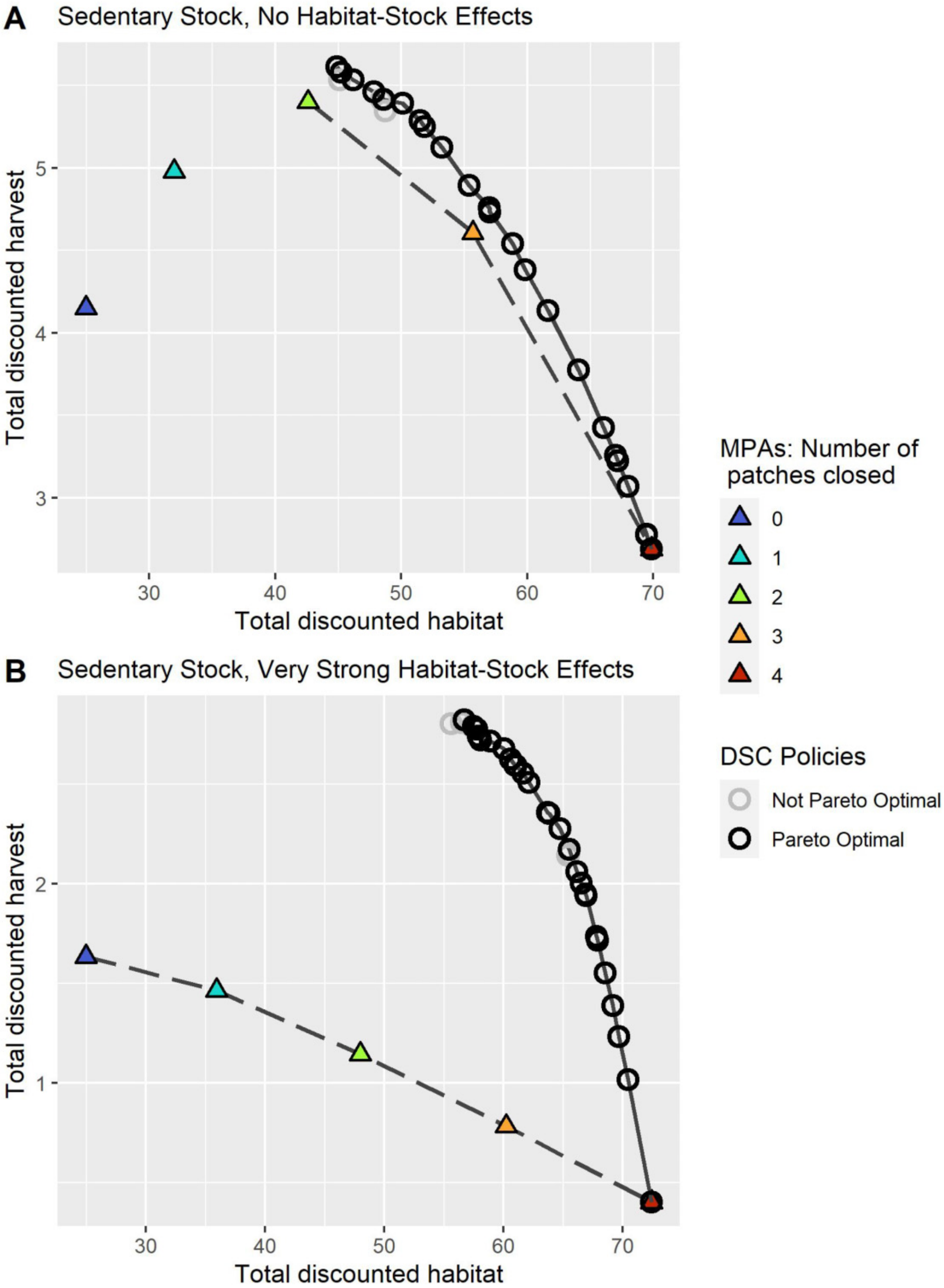

To visualize how DSC policies change as the weight given to habitat protection changes, for each of the two scenarios discussed previously, we selected four example policies to present for the no habitat–stock effects scenario (Fig. S9) and the very strong habitat–stock effects scenario (Fig. 6). As discussed in Section 3.1.1, the policies generally had two distinct sections: a "recovery" period followed by a "rotational" period. We found that the length of the recovery period increased as the habitat weight (w_{hab}) increased. For example, in the no habitat–stock effects (very strong habitat–stock effects) scenario, the recovery period lasted 5 (8) years when $w_{\mathrm{hab}} = 0$, but 11 (11) years when $w_{\mathrm{hab}} = 0.05$ (Figs. 6A and 6B and S9A and S9B). As the habitat weight increased, we also found that switches in the rotational period became less frequent, and

Fig. 6. The best dynamic spatial closure (DSC) policy from multiobjective optimization (for both harvest and habitat protection), and corresponding best marine protected area (MPA), as the weight for habitat protection in the reward increases (w_{hab}, with $w_{\mathrm{harv}} = 1 - w_{\mathrm{hab}}$). Increasing the weight for habitat protection strongly affected both the length of the recovery period and the behavior in the rotational period for the DSC policies, and the number of patches closed in the best MPA. All panels consider a sedentary stock under very strong habitat–stock effects ($\phi_{\mathrm{m}} = \phi_{\mathrm{s}} = 1$); other parameters take their base case values (see Table 1). Only the first 50 years of decisions are shown for clarity. Note that MPAs with the same number of patches closed perform identically for a sedentary stock.

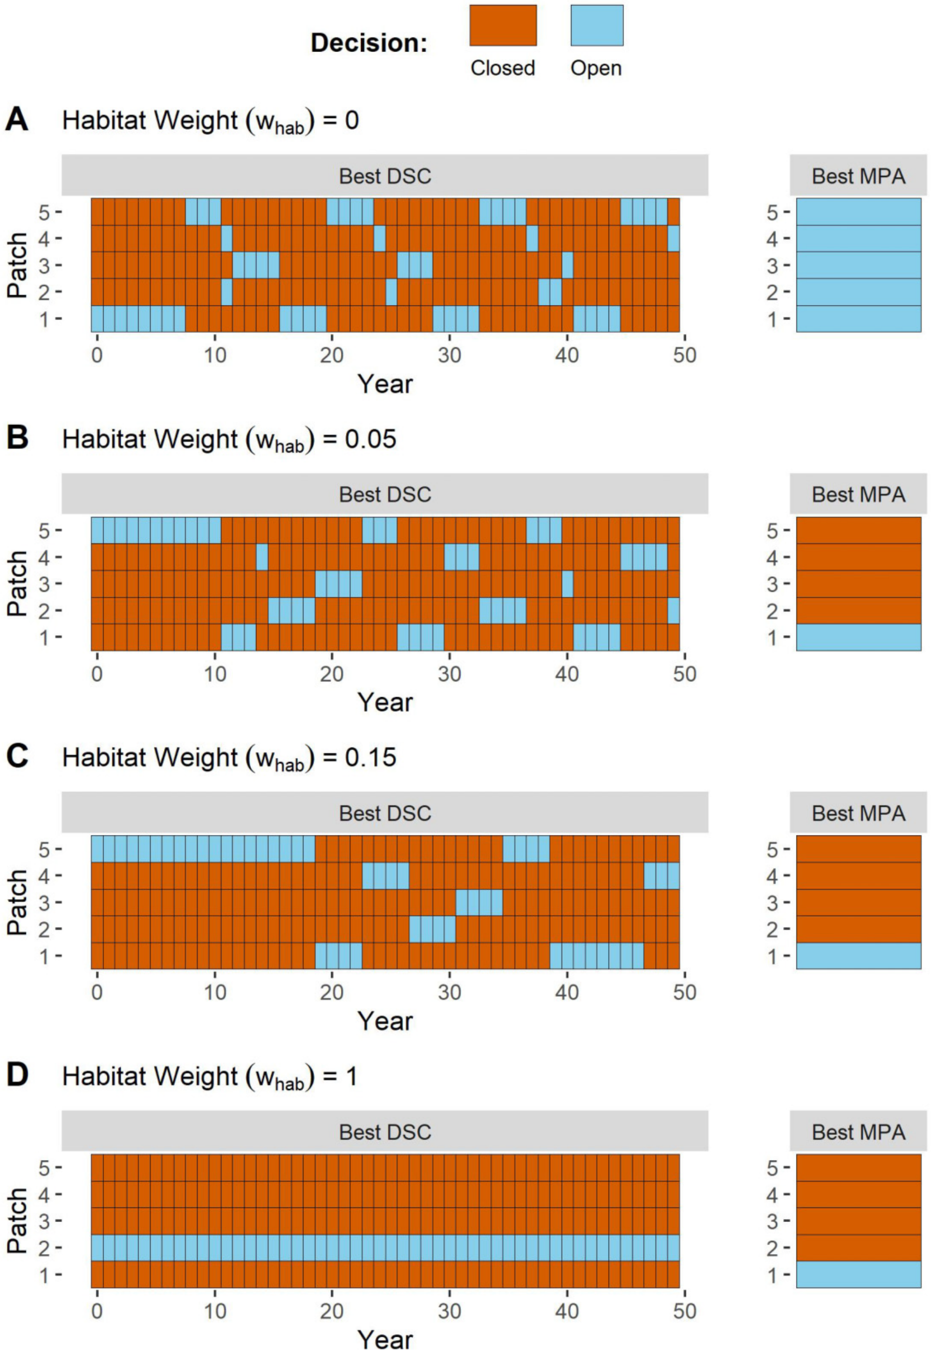

the lengths of patch closures generally became longer. When the habitat weight became sufficiently large, there were no switches at all: in both scenarios, the optimal DSC policy was an MPA with one patch open (Figs. 6D and S9D).

In both scenarios, the number of patches closed in the best MPA increased as the habitat weight increased (Figs. 6 and S9; note that MPAs with the same number of patches closed perform identically for a sedentary stock). When the habitat weight was sufficiently large, the best MPA always involved a single open patch (recall that the manager was not allowed to close all *n* patches). In the very strong habitat–stock effects scenario (Fig. 6), this shift occurred very quickly; note that because of the near-linear shape of the MPA Pareto frontier in Fig. 5B, there were very few weights for which an MPA with one, two, or three patch closures was optimal.

4. Discussion

Our approach allowed us to evaluate the effectiveness of DSCs for balancing harvest with habitat protection in a context that has previously focused predominantly on MPAs (i.e., fixed spatial closures). We found that DSCs offered improved yield and greater habitat protection versus MPAs across many scenarios for a fishery with fishing-induced habitat damage. To date, on-the-ground implementations of DSCs have predominantly focused on targeted benthic stocks such as scallops or urchins (Cohen and Foale 2013; Plagányi et al. 2015; O'Boyle et al. 2017). The best DSC policy still outperformed the best MPA in some scenarios for a mobile stock (notably under strong habitat–stock effects), demonstrating that they may warrant further attention for such taxa, although the benefits were generally greater for a sedentary stock. We also found that DSCs were always able to match or increase yield versus the nonspatial MSY benchmark, along with generally providing greater habitat protection. This is despite several restrictions faced by our DSC policies (the fishery was open-access with no direct way to control effort, not all *n* patches could be closed at once, etc.). Overall, we found DSCs to be effective at meeting management goals in an open-access setting, demonstrating their potential value when management resources are limited (see also Reithe et al. 2014).

How strongly the stock is affected by habitat quality greatly impacted the performance of DSCs. For instance, the relative benefits of DSCs (versus both MPAs and the nonspatial MSY benchmark) were typically greater when habitat quality had a larger effect on early life stage survival. Notably, DSCs only produced higher average yields than nonspatial MSY when this survival effect was present, due to the increased habitat protection DSCs provided. This result broadly agrees with past studies of spatial management (Moeller and Neubert 2013, 2015; Alqawasmeh and Lutscher 2019; Kelly et al. 2019). The affinity that fish had for less disturbed areas was also important. We found that MPAs performed poorly relative to DSCs in terms of both harvest and habitat protection when fish had a strong preference for high-quality habitat (because fish clustered in closed areas; see also Alqawasmeh and Lutscher 2019). DSCs were able to alleviate this problem by allowing access to previously closed areas. As there is empirical evidence supporting an interaction between habitat quality and fish movement in some systems (Claudet et al. 2010; Grüss et al. 2011), this connection may deserve further attention when considering the effectiveness of spatial management policies. While changes in other factors (habitat recovery rate, gear damage, etc.) affected the quantitative performance of DSCs, the qualitative trends caused by habitat–stock effects were consistent, which may have important management implications for any stock thought to be highly affected by habitat quality.

Identification of optimal DSC policies is useful for determining the theoretical maximum performance of spatial management strategies; however, this approach may present a high information burden to implement in practice, requiring information on stock status, habitat distributions, and fishery dynamics. To investigate the potential for implementing DSCs with simpler approaches, we also explored the harvest and habitat conservation performance of heuristic rotating closure policies. These heuristic policies performed robustly across a suite of rotational schedules tested (achieving comparable or greater harvests than the best MPA), demonstrating their potential in systems where limited information is available for spatial siting analyses. We further showed how improved heuristic spatial policies could be developed when additional information is available (e.g., using a pulse rather than rotational strategy when suggested by the optimal DSC policy). As heuristic spatial policies were able to capture many of the benefits achievable by optimal DSC policies, they may provide a practical yet effective management option for real-world fisheries.

We found that DSCs offered both greater harvest and habitat protection than the nonspatial MSY benchmark in many scenarios, especially those with slower habitat recovery and more damaging gear. Some types of habitat may take even longer to recover postfishing than our slowest simulated habitat recovery rates (e.g., coral, Rooper et al. 2011), for which spatial management is likely to be even more important. Other benefits spatial closures may offer over nonspatial controls include being easier to manage, being able to protect dedicated portions of fish populations, and providing less disturbed areas for scientific study (Hilborn et al. 2004). However, it is important to consider how spatial closures and shifts in effort may impact fishers/fishing communities when considering spatial management policies (Hilborn et al. 2004). DSCs would require both enforcement and cooperation from fishers to be effective (as with MPAs; Sethi and Hilborn 2008). Although we focus on how well spatial closures perform alone, they are likely to produce better overall results when combined with other management methods (when such methods are feasible). Future research into DSCs might examine their performance alongside harvest/effort controls, a question which has mostly been considered for MPAs in this setting (but see Kelly et al. 2019).

Given the potential benefits of DSCs identified by our analysis, it is important to explore the robustness of our findings and note where future research might be valuable. First, there are many potential forms for the habitat–stock effects that could be considered (e.g., besides early life stage survival, habitat quality might impact population growth via effects on somatic growth, natural mortality, or so forth). Further theoretical investigation of habitat–stock interactions represented by a variety of functional forms may reveal important drivers for the effectiveness of spatial management policies; additional empirical studies of the impacts of habitat quality on a variety of taxa would provide an important foundation for this study. Second, we chose to focus on a single stock and its links to habitat to reduce the size of the state space. Our model could be expanded in the future to incorporate multispecies dynamics, an important consideration for ecosystem-based fisheries management, at the cost of an increased computational burden. State models that incorporate structured populations may also be of interest in cases where the size or age of fish is important for fisheries harvest/conservation goals.

As our aim was to represent a general case of a fishery with fishing-induced habitat damage while limiting the computational burden, we chose to use a simple spatial configuration (five patches in a line). While it would be possible to model additional patches or more complex spatial arrangements, we do not expect it to change our overall conclusions (e.g., that DSCs outperform MPAs under strong habitat–stock effects). Still, there may be real-world complexities and questions relevant to managers requiring a model with more complex spatial processes. For instance, if the fishery forms a habitat mosaic with heterogeneous quality, or if different life stages utilize different areas, the optimal DSC/MPA might look very different. Our framework could be modified to examine such questions (e.g., by letting $h_{\max}$ differ between patches). While we focused on two crucial fisheries management objectives, maximizing harvest and protecting habitat, investigation of the ability for DSCs to meet other goals (economic, conservation, social, etc.) may reveal additional benefits or drawbacks of this management tool. Lastly, while we found our results were robust to stochastic recruitment (a key source of variability; Thorson et al. 2014), there are other types of stochasticity that may be relevant for a variety of systems. For instance, examining stochasticity in habitat quality or negative shocks resulting in large mortality events (Aalto et al. 2019; White et al. 2021) might reveal situations where taking variability into account is essential for spatial decision-making.

Our work provides important contributions to the study of fisheries with fishing-induced habitat damage. The transition to ecosystem-based fisheries management recognizes the importance of management objectives beyond maximizing yield, a key one of which is maintaining seafloor habitat integrity. We found DSCs often resulted in greater habitat protection compared with other policies, even when habitat protection was not included in the optimization goal. Furthermore, we demonstrated how DSCs offer a wide variety of options to meet different goals balancing harvest and habitat protection. While DSCs have not been widely considered as management options in comparison to MPAs, our work demonstrates that they may show promise across a range of systems and taxa for meeting fisheries management goals. The need for a flexible, dynamic management approach may be especially important under rapid climate-driven ocean changes (Lynch et al. 2021).

Acknowledgements

We thank Alex Vladimirsky for providing valuable assistance regarding optimization methodology and feedback during the model development stage. We are also grateful to the Cornell Ecotheory group for their helpful feedback on the manuscript.

Article information

History dates

Received: 25 August 2022
Accepted: 22 December 2022
Version of record online: 28 February 2023

Copyright

Data availability

Novel code and output (Poulton et al. 2022) are available on Figshare (doi:10.6084/m9.figshare.20288523.v1, https://doi.org/10.6084/m9.figshare.20288523.v1).

Author information

Author ORCIDs

Anna J. Poulton https://orcid.org/0000-0003-3066-365X
Suresh A. Sethi https://orcid.org/0000-0002-0053-1827
Stephen P. Ellner https://orcid.org/0000-0002-8351-9734
T. Scott Smeltz https://orcid.org/0000-0002-8578-8540

Author contributions

Conceptualization: AJP, SAS, SPE, TSS
Formal analysis: AJP
Methodology: AJP, SAS, SPE, TSS
Software: AJP
Supervision: SAS, SPE
Visualization: AJP
Writing – original draft: AJP
Writing – review & editing: AJP, SAS, SPE, TSS

Competing interests

The authors declare there are no competing interests.

Funding information

This work was supported by funding from the NSF Research Training Group program (Dynamics, Probability, and Partial Differential Equations in Pure and Applied Mathematics, DMS-1645643).

Supplementary material

Supplementary data are available with the article at https://doi.org/10.1139/cjfas-2022-0198.

References

Aalto, E.A., Micheli, F., Boch, C.A., Espinoza Montes, J.A., Woodson, C.B., and De Leo, G.A. 2019. Catastrophic mortality, allee effects, and marine protected areas. Am. Nat. **193**(3): 391–408. doi:10.1086/701781.

Agardy, T., di Sciara, G.N., and Christie, P. 2011. Mind the gap: addressing the shortcomings of marine protected areas through large scale marine spatial planning. Mar. Policy, **35**(2): 226–232. doi:10.1016/j.marpol.2010.10.006.

Akpalu, W., and Bitew, W.T. 2014. Optimum reserve size, fishing induced change in carrying capacity, and phenotypic diversity. J. Bioecon. **16**(3): 289–304. doi:10.1007/s10818-014-9178-8.

Alqawasmeh, Y., and Lutscher, F. 2019. Movement behaviour of fish, harvesting-induced habitat degradation and the optimal size of marine reserves. Theor. Ecol. **12**(4): 453–466. doi:10.1007/s12080-019-0411-x.

Auster, P.J. 1998. A conceptual model of the impacts of fishing gear on the integrity of fish habitats. Conserv. Biol. **12**(6): 1198–1203. doi:10.1046/j.1523-1739.1998.0120061198.x.

Barbier, E.B. 2017. Marine ecosystem services. Curr. Biol. **27**(11): R507–R510. doi:10.1016/j.cub.2017.03.020.

Bradshaw, C., Collins, P., and Brand, A.R. 2003. To what extent does upright sessile epifauna affect benthic biodiversity and community composition? Mar. Biol. **143**(4): 783–791. doi:10.1007/s00227-003-1115-7.

Cabral, R.B., Halpern, B.S., Lester, S.E., White, C., Gaines, S.D., and Costello, C. 2019. Designing MPAs for food security in open-access fisheries. Sci. Rep. **9**(1): 8033. doi:10.1038/s41598-019-44406-w.

Carvalho, P.G., Jupiter, S.D., Januchowski-Hartley, F.A., Goetze, J., Claudet, J., Weeks, R., et al. 2019. Optimized fishing through periodically harvested closures. J. Appl. Ecol. **56**(8): 1927–1936. doi:10.1111/1365-2664.13417.

Cashion, T., Al-Abdulrazzak, D., Belhabib, D., Derrick, B., Divovich, E., Moutopoulos, D.K., et al. 2018. Reconstructing global marine fishing gear use: catches and landed values by gear type and sector. Fish. Res. **206**: 57–64. doi:10.1016/j.fishres.2018.04.010.

Cinner, J., Marnane, M.J., McClanahan, T.R., and Almany, G.R. 2005. Periodic closures as adaptive coral reef management in the Indo-Pacific. Ecol. Soc. **11**(1): 31. doi:10.5751/ES-01618-110131.

Claudet, J., Osenberg, C.W., Domenici, P., Badalamenti, F., Milazzo, M., Falcón, J.M., et al. 2010. Marine reserves: fish life history and ecological traits matter. Ecol. Appl. **20**(3): 830–839. doi:10.1890/08-2131.1.

Cohen, P.J., and Foale, S.J. 2013. Sustaining small-scale fisheries with periodically harvested marine reserves. Mar. Policy, **37**: 278–287. doi:10.1016/j.marpol.2012.05.010.

Collie, J.S., Escanero, G.A., and Valentine, P.C. 1997. Effects of bottom fishing on the benthic megafauna of Georges Bank. Mar. Ecol. Prog. Ser. **155**: 159–172. doi:10.3354/meps155159.

Commonwealth of Australia. 2021. Assessment of the Queensland Sea Cucumber Fishery.

Costanza, R., d'Arge, R., de Groot, R., Farber, S., Grasso, M., Hannon, B., et al. 1997. The value of the world's ecosystem services and natural capital. Nature, **387**(6630): 253–260. doi:10.1038/387253a0.

Costello, M.J., and Ballantine, B. 2015. Biodiversity conservation should focus on no-take Marine Reserves. Trend. Ecol. Evol. **30**(9): 507–509. doi:10.1016/j.tree.2015.06.011.

European Commission. 2008. Directive 2008/56/EC of the European Parliament and of the Council of 17 June 2008 establishing a framework for community action in the field of marine environmental policy (Marine Strategy Framework Directive).

Falcó, C., and Moeller, H.V. 2022. Optimal spatial management in a multiuse marine habitat: balancing fisheries and tourism. Nat. Res. Model. **35**(1): e12309. doi:10.1111/nrm.12309.

Foden, J., Rogers, S., and Jones, A. 2011. Human pressures on UK seabed habitats: a cumulative impact assessment. Mar. Ecol. Prog. Ser. **428**: 33–47. doi:10.3354/meps09064.

Grabowski, J.H., Bachman, M., Demarest, C., Eayrs, S., Harris, B.P., Malkoski, V., et al. 2014. Assessing the vulnerability of marine benthos to fishing gear impacts. Rev. Fish. Sci. Aquacult. **22**(2): 142–155. doi:10.1080/10641262.2013.846292.

Grüss, A., Kaplan, D.M., and Hart, D.R. 2011. Relative impacts of adult movement, larval dispersal and harvester movement on the effectiveness of reserve networks. PLoS One, **6**(5): 1–14. doi:10.1371/journal.pone.0019960.

Hart, D.R. 2006. When do marine reserves increase fishery yield? Can. J. Fish. Aquat. Sci. **63**(7): 1445–1449. doi:10.1139/f06-071.

Hastings, A., and Botsford, L.W. 1999. Equivalence in yield from marine reserves and traditional fisheries management. Science, **284**(5419): 1537–1538. doi:10.1126/science.284.5419.1537.

Hiddink, J.G., Johnson, A.F., Kingham, R., and Hinz, H. 2011. Could our fisheries be more productive? Indirect negative effects of bottom trawl fisheries on fish condition. J. Appl. Ecol. **48**(6): 1441–1449. doi:10.1111/j.1365-2664.2011.02036.x.

Hilborn, R., Stokes, K., Maguire, J.J., Smith, T., Botsford, L.W., Mangel, M., et al. 2004. When can marine reserves improve fisheries management? Ocean Coast. Manage. **47**(3): 197–205. doi:10.1016/j.ocecoaman.2004.04.001.

Hutchings, J.A., and Reynolds, J.D. 2004. Marine fish population collapses: consequences for recovery and extinction risk. BioScience, **54**(4): 297–309. doi:10.1641/0006-3568(2004)054[0297:MFPCCF]2.0.CO;2.

Kelly, M.R., Neubert, M.G., and Lenhart, S. 2019. Marine reserves and optimal dynamic harvesting when fishing damages habitat. Theor. Ecol. **12**(2): 131–144. doi:10.1007/s12080-018-0399-7.

Le Quesne, W.J.F., and Codling, E.A. 2009. Managing mobile species with MPAs: the effects of mobility, larval dispersal, and fishing mortality on closure size. ICES J. Mar. Sci. **66**(1): 122–131. doi:10.1093/icesjms/fsn202.

Lester, S., Halpern, B., Grorud-Colvert, K., Lubchenco, J., Ruttenberg, B., Gaines, S., et al. 2009. Biological effects within no-take marine reserves: a global synthesis. Mar. Ecol. Prog. Ser. **384**: 33–46. doi:10.3354/meps08029.

Lynch, A.J., Thompson, L.M., Beever, E.A., Cole, D.N., Engman, A.C., Hawkins Hoffman, C., et al. 2021. Managing for RADical ecosystem change: applying the Resist–Accept–Direct (RAD) framework. Front. Ecol. Environ. **19**(8): 461–469. doi:10.1002/fee.2377.

Maunder, M. 2008. Maximum sustainable yield. In Encyclopedia of ecology. Edited by S.E. Jørgensen and B.D. Fath. Academic Press, Oxford. pp. 2292–2296.

Moeller, H.V., and Neubert, M.G. 2013. Habitat damage, marine reserves, and the value of spatial management. Ecol. Appl. **23**(5): 959–971. doi:10.1890/12-0447.1.

Moeller, H.V., and Neubert, M.G. 2015. Economically optimal marine reserves without spatial heterogeneity in a simple two-patch model. Nat. Res. Model. **28**(3): 244–255. doi:10.1111/nrm.12066.

Munos, R., and Szepesvári, C. 2008. Finite-time bounds for fitted value iteration. J. Mach. Learn. Res. **9**(27): 815–857. doi:10.1145/1390681.1390708.

Myers, R.A., Fuller, S.D., and Kehler, D.G. 2000. A fisheries management strategy robust to ignorance: rotational harvest in the presence of indirect fishing mortality. Can. J. Fish. Aquat. Sci. **57**(12): 2357–2362. doi:10.1139/f00-250.

National Research Council. 2001. Marine protected areas: tools for sustaining ocean ecosystems. The National Academies Press, Washington, DC.

National Research Council. 2002. Effects of trawling and dredging on seafloor habitat. The National Academies Press, Washington, DC.

New England Fishery Management Council. 2003. Final Amendment 10 to the Atlantic Sea Scallop Fishery management plan with a supplemental environmental impact statement, regulatory impact review, and regulatory flexibility analysis.

Nichols, R., Yamazaki, S., and Jennings, S. 2018. The role of precaution in stock recovery plans in a fishery with habitat effect. Ecol. Econ. **146**: 359–369. doi:10.1016/j.ecolecon.2017.11.033.

NOAA National Marine Fisheries Service. 2007. Magnuson–Stevens Fishery Conservation and Management Act. Public Law 94-265.

O'Boyle, R., Roddick, D., and Sinclair, M. 2017. A review of rotational management of fisheries and its application to canadian east coast arctic surfclam (mactromeris polynyma). Can. Sci. Adv. Secret. Res. Doc. **066**: 1–33.

Peterson, C.H., Summerson, H.C., and Fegley, S.R. 1987. Ecological consequences of mechanical harvesting of clams. Fish. Bull. **85**: 281–298. doi:10.17615/2rb3-rm41.

Pikitch, E.K., Santora, C., Babcock, E.A., Bakun, A., Bonfil, R., Conover, D.O., et al. 2004. Ecosystem-based fishery management. Science, **305**(5682): 346–347. doi:10.1126/science.1098222.

Plagányi, É. E., Skewes, T., Murphy, N., Pascual, R., and Fischer, M. 2015. Crop rotations in the sea: increasing returns and reducing risk of collapse in sea cucumber fisheries. Proc. Natl. Acad. Sci. U.S.A. **112**(21): 6760–6765. doi:10.1073/pnas.1406689112.

Poulton, A.J., Sethi, S.A., Ellner, S.P., and Smeltz, T.S. 2022. Optimal dynamic spatial closures can improve fishery yield and reduce fishing-induced habitat damage. Code and results. Figshare. doi:10.6084/m9.figshare.20288523.v1.

R Core Team. 2020. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.

Ramirez-Llodra, E., Tyler, P.A., Baker, M.C., Bergstad, O.A., Clark, M.R., Escobar, E., et al., 2011. Man and the last great wilderness: human impact on the deep sea. PLoS One, **6**(8): e22588. doi:10.1371/journal.pone.0022588.

Reithe, S., Armstrong, C.W., and Flaaten, O. 2014. Marine protected areas in a welfare-based perspective. Mar. Policy, **49**: 29–36. doi:10.1016/j.marpol.2014.04.002.

Rooper, C.N., Wilkins, M.E., Rose, C.S., and Coon, C. 2011. Modeling the impacts of bottom trawling and the subsequent recovery rates of sponges and corals in the Aleutian Islands, Alaska. Cont. Shelf Res. **31**(17): 1827–1834. doi:10.1016/j.csr.2011.08.003.

Russ, G.R. 2002. Yet another review of marine reserves as reef fishery management tools. In Coral reef fishes. Edited by P.F. Sale. Academic Press, San Diego. pp. 421–443.

Sala, E., Mayorga, J., Bradley, D., Cabral, R.B., Atwood, T.B., Auber, A., et al. 2021. Protecting the global ocean for biodiversity, food and climate. Nature, **592**(7854): 397–402. doi:10.1038/s41586-021-03371-z.

Sanchirico, J.N., and Wilen, J.E. 2001. A bioeconomic model of marine reserve creation. J. Environ. Econ. Manage. **42**(3): 257–276. doi:10.1006/jeem.2000.1162.

Scharf, F.S., Manderson, J.P., and Fabrizio, M.C. 2006. The effects of seafloor habitat complexity on survival of juvenile fishes: species-specific interactions with structural refuge. J. Exp. Mar. Biol. Ecol. **335**(2): 167–176. doi:10.1016/j.jembe.2006.03.018.

Sethi, S.A., and Hilborn, R. 2008. Interactions between poaching and management policy affect marine reserves as conservation tools. Biol. Conserv. **141**(2): 506–516. doi:10.1016/j.biocon.2007.11.005.

Smeltz, T.S., Harris, B.P., Olson, J.V., and Sethi, S.A. 2019. A seascape-scale habitat model to support management of fishing impacts on benthic ecosystems. Can. J. Fish. Aquat. Sci. **76**(10): 1836–1844. doi:10.1139/cjfas-2018-0243.

Smith, V.L. 1969. On models of commercial fishing. J. Polit. Econ. **77**(2): 181–198. doi:10.1086/259507.

Soetaert, K., Petzoldt, T., and Setzer, R.W. 2010. Solving differential equations in R: package deSolve. J. Stat. Software, **33**(9): 1–25. doi:10.18637/jss.v033.i09.

Sutton, R.S., and Barto, A.G. 2018. Reinforcement learning: an introduction. 2nd ed. The MIT Press, Cambridge, MA.

Thorson, J.T., Jensen, O.P., and Zipkin, E.F. 2014. How variable is recruitment for exploited marine fishes? A hierarchical model for testing life history theory. Can. J. Fish. Aquat. Sci. **71**(7): 973–983. doi:10.1139/cjfas-2013-0645.

Watling, L., and Norse, E.A. 1998. Disturbance of the seabed by mobile fishing gear: a comparison to forest clearcutting. Conserv. Biol. **12**(6): 1180–1197. doi:10.1046/j.1523-1739.1998.0120061180.x.

Watling, L., Findlay, R.H., Mayer, L.M., and Schick, D.F. 2001. Impact of a scallop drag on the sediment chemistry, microbiota, and faunal assemblages of a shallow subtidal marine benthic community. J. Sea Res. **46**(3): 309–324. doi:10.1016/S1385-1101(01)00083-1.

White, E.R., Baskett, M.L., and Hastings, A. 2021. Catastrophes, connectivity and Allee effects in the design of marine reserve networks. Oikos, **130**(3): 366–376. doi:10.1111/oik.07770.

Wood, S.N. 2017. Generalized additive models: an introduction with R. 2nd ed. Chapman and Hall/CRC.

Appendix A

A.1 Approximate dynamic programming

Here we describe the optimization procedure used to find policies of DSCs. As we consider both deterministic and stochastic environments, the procedure here is written in general terms that work in both cases (e.g., we consider the expected total discounted reward, which would simplify to just the total discounted reward in the deterministic case).

Let π be a policy, a function that outputs the manager's action at each state: given a state $\mathbf{s}(d)$ at the beginning of year d, the action taken by the manager is $\mathbf{a}(d) = \pi(\mathbf{s}(d))$. The performance of a policy is quantified by its value function V^{π}, which gives the expected total discounted reward given that the system begins at state $\mathbf{s}$ and decisions are made according to the policy π (Sutton and Barto 2018).

$$\text{(A1)}\quad V^{\pi}(\mathbf{s}) = \mathbb{E}\left[\sum_{d=0}^{\infty} \delta^{d}\ \text{Reward}(d)\ \mid \mathbf{s}(0) = \mathbf{s},\ \text{decisions made by } \pi\right]$$

Notably, the value function can be written recursively in the form of Bellman's equation (Sutton and Barto 2018), where Reward(0) and new state $\mathbf{s}'$ result from running the model for 1 year starting in state $\mathbf{s}$ and using action $\pi(\mathbf{s})$.

$$\text{(A2)}\quad V^{\pi}(\mathbf{s}) = \mathbb{E}\left[\text{Reward}(0) + \delta V^{\pi}(\mathbf{s}')\ \mid \mathbf{s}(0) = \mathbf{s}, \mathbf{a}(0) = \pi(\mathbf{s})\right]$$

Ideally, we would like to identify an optimal policy: a policy $\pi*$ that maximizes the total discounted reward ($\forall \mathbf{s} \in \mathcal{S},\ V^{\pi^*}(\mathbf{s}) = \max_{\tilde{\pi}} V^{\tilde{\pi}}(\mathbf{s})$). However, due to the complexity of our model (notably the large, continuous state space), we must turn to approximate optimization methods. Specifically, we use an ADP algorithm called "sampling-based fitted value iteration" (as described in Munos and Szepesvári 2008), the general procedure of which is given in Fig. A1.

As an initial value function, we select $\widehat{V}_0 = 0$ uniformly. With this as the initial value function, $\widehat{V}_1(\mathbf{s})$ estimates the maximum expected reward obtained in 1 year, starting from state $\mathbf{s}$. Expanding this, it can be seen that generally, $\widehat{V}_k$ estimates the maximum expected discounted reward from a state when looking ahead k years. Intuitively, K (the total number of training iterations) represents how many years our resulting policy will look ahead when trying to make optimal decisions. We selected $K = 30$, and did not find increasing K to result in better policies; with future years discounted at rate δ, the first 30 years are ≈80% of the (discounted) infinite future. In other words, with a discount factor of $\delta = 0.95$, we found that it was not necessary for the policy to consider the fishery more than 30 years in the future in order to make good decisions.

We chose to sample $J = 200$ states $\mathbf{s}_j$ during each training iteration. To randomly generate these states, where $\mathbf{s}_j = (x_{1,j}, \ldots, x_{n,j}, h_{1,j}, \ldots, h_{n,j}, f_{1,j}, \ldots, f_{n,j})$, we used the following procedure. First, $x_{i,j}$ are drawn *iid*, with probability 0.7 from $\text{unif}(0, \tilde{x}_i^*)$ and with probability 0.3 from $\text{unif}(x_i^* - 0.025, x_i^* + 0.025)$; recall that x_i^* is the open-access equilibrium and $\tilde{x}_i^*$ is the no-fishing equilibrium. We chose to sample often from near the open-access equilibrium, since open-access dynamics often dominate (as they occur whenever a patch is open). Then, we use the same procedure to draw values for $h_{i,j}$, substituting in the appropriate equilibria. Finally, for $f_{i,j}$, we simply let $f_{i,j} = f_i^*$ if $h_{i,j} \leq h_i^* + 0.025$, and $f_{i,j} = 0$ otherwise (our assumption being that if $h_{i,j}$ is greater than this, the patch was likely closed or had little effort during the previous year). Additionally, we always include the open-access equilibrium of the fishery ($x_{i,j} = x_i^*$, $h_{i,j} = h_i^*$, $f_{i,j} = f_i^*$) in every training set, as we found estimating the value of this state to be very important when training policies. This may be because the model ends up near this state frequently (whenever all n patches were open in the previous year); this state is also our initial condition, as we assume that we start in open-access equilibrium. Finally,

Fig. A1. Summary of the ADP algorithm we use to find near-optimal DSC policies: sampling-based fitted value iteration, from Munos and Szepesvári 2008.

ADP: Sampling-Based Fitted Value Iteration

1. Select an initial value function $\hat{V}_0$
2. Generate random states $\mathbf{s}_1, \ldots, \mathbf{s}_J$, and add them to the training set
3. For each state $\mathbf{s}_j$, calculate $\tilde{V}_j$ and add it to the training set:

   $$\tilde{V}_j = \max_{\mathbf{a}\in\mathscr{A}} \mathbb{E}[\text{Reward}(0) + \delta\hat{V}_k(\mathbf{s}') \mid \mathbf{s}(0) = \mathbf{s}_j, \mathbf{a}(0) = \mathbf{a}] \tag{A3}$$

   where Reward(0) and new state $\mathbf{s}'$ result from running the model for one year starting in state $\mathbf{s}(0)$ and using action $\mathbf{a}(0)$. In the stochastic case, the expected value is estimated as the mean over M_1 trials.
4. Train a functional approximation $\hat{V}_{k+1}$ using the training set data $(\mathbf{s}_j, \tilde{V}_j)_{j=1}^J$
5. Repeat steps 2-4 for K iterations
6. After obtaining $\hat{V}_K$, let π be the policy given by

   $$\pi(\mathbf{s}) = \text{argmax}_{\mathbf{a}\in\mathscr{A}} \mathbb{E}[\text{Reward}(0) + \delta\hat{V}_K(\mathbf{s}') \mid \mathbf{s}(0) = \mathbf{s}, \mathbf{a}(0) = \mathbf{a}] \tag{A4}$$

   where again, the expected value is estimated as the mean of M_1 trials in the stochastic case. Test the performance of π by using it to make decisions over a time horizon of t_f years (in the stochastic case, averaging over M_2 trials of the policy).

we make use of the symmetry of our problem to increase the size of the training set. As $\mathbf{s}_j = (x_{1,j}, \ldots, x_{n,j}, h_{1,j}, \ldots, h_{n,j}, f_{1,j}, \ldots, f_{n,j})$ and $\tilde{\mathbf{s}}_j = (x_{n,j}, \cdots, x_{1,j}, h_{n,j}, \cdots, h_{1,j}, f_{n,j}, \cdots, f_{1,j})$ will always result in the same reward, we were able to double the size of our training set without performing any additional simulations.

Using the training set data $(\mathbf{s}_j, \tilde{V}_j)_{j=1}^J$, the following GAM formula (implemented using the mgcv package in R) was used to train $\widehat{V}_{k+1}$:

$$\tilde{V}_j = \beta_0 + \sum_{i=1}^{n} \Big[S_{3i-2}\left(x_{i,j}\right) + S_{3i-1}\left(h_{i,j}\right) + S_{3i}\left(x_{i,j}\right) x_j + T_i\left(x_{i,j}, h_{i,j}\right) \Big] + \varepsilon_j \tag{A5}$$

where S are smooth functions using a thin plate regression spline basis, and T are tensor product interaction smooths using a cubic regression spline basis, $x_j = \sum_{i=1}^n x_{i,j}$, and $\varepsilon_j \sim N\left(0, \sigma_\varepsilon^2\right)$. The first two terms in the sum are just smooth functions of the stock size and the habitat quality in patch i. The last two terms include interactions between the stock size in patch i and the total stock size, and between the stock size and habitat quality of patch i. The basis dimension k was set to a low value in each term to limit the computational burden and the number of training points required; we did not find increasing k further to result in better performing policies. Other training settings used the defaults in the mgcv package.

Since the samples of states used by the ADP algorithm are randomly generated, each full run of the algorithm may produce DSC policies with slightly different properties and performance results. For instance, consider the DSC policies in Figs. 2A and S9A, which were produced under the exact same parameters. While these policies have some qualitative differences (e.g., a patch was permanently closed in Fig. 2 A but not in Fig. S9A), they performed similarly, with only a 1.13% difference in discounted reward. In the single objective (harvest-maximizing) deterministic setting, we performed three runs of the ADP algorithm for each parameter combination under consideration, and selected the DSC policy resulting in the highest total discounted harvest. Overall, we found the difference between runs to be small across parameter space (e.g., the mean percent difference between the performance of the best and worst DSC policies of the three runs in Section 3.1.2 was only 1.59%). Thus, in our further analyses (multiobjective and stochastic settings), we only performed one run of the ADP algorithm for each parameter combination.

A.2 Parameter selection

For the base case, we selected $\mu = 0.2$ and $r = 0.8$ to represent a medium productivity stock with a moderate (nonspatial) maximum sustainable yield exploitation rate of 20% (MSY/$x^{\mathrm{MSY}} = 20\%$ under no habitat–stock effects; see Section A.4 for details on calculating nonspatial MSY). The habitat recovery rate for the base case was also selected to be moderate ($\beta = 0.2$, resulting in a mean recovery time of five years), but we note that estimated real habitat recovery rates differ greatly based on the type of habitat (Grabowski et al. 2014; Rooper et al. 2011). As the ratio between cost per unit effort (c) and price per unit harvest (p) determines the open-access dynamics, we simply set $c = 1$. The rates of effort movement (λ_{ent}, λ_{exit}, λ_{rel}) were hand selected, with $\lambda_{\mathrm{rel}} > \lambda_{\mathrm{ent}}, \lambda_{\mathrm{exit}}$ representing that movement between patches occurs more quickly than entering/exiting the fishery does. The discount factor δ was also selected to be moderate, although this could vary greatly between fisheries based on how much current versus future goals are prioritized. The values for the habitat–stock effects parameters ($\phi_{\mathrm{m}}, \phi_{\mathrm{s}} = 0$, 0.5, 0.8, 1) were selected because in the base case, they respectively result in a fully recovered patch ($h_i\,(t) = \tilde{h}_i^*$) having a scaled habitat effect (eqs. 4 and 5) that is 1, 2, 3, or 4 times greater than that of a patch at open-access equilibrium ($h_i\,(t) = h_i^*$).

For the alternative cases (Table 2), we adjust three key factors: the habitat recovery rate, the damage caused by fishing gear (equivalently, the sensitivity of the habitat), and the productivity of the stock. To examine how more slowly recovering habitat may respond to spatial management, we consider cases where the habitat recovery rate is $\beta = 0.1$, half its base case value. To examine how improvements in fishing gear resulting in less habitat damage might impact our fishery, we consider cases where the value for the open-access habitat equilibrium (h_i^*) is selected to result in the habitat impact parameter (γ) being half its base case value, representing gear that is half as damaging. Finally, we consider one alternative case concerning a less productive stock where r is selected such that the resulting (nonspatial) maximum sustainable yield exploitation rate is only 10%, compared with 20% in the base case.

A.3 Model equilibria

Here we give formulas for the model equilibria. These were derived using the assumptions from Section 2.2 (notably that effort follows open-access dynamics, and that the fish movement weights $Y_{i,\,j}$ are as described).

The open-access equilibria (all patches are open) for the stock, habitat, and effort are as follows:

$$\text{(A6)}\quad x_i^* = \frac{c}{pq}, \qquad h_i^* = \frac{\beta h_{\max}}{\gamma f_i^* + \beta}, \qquad f_i^* = \frac{1}{q}\left(\frac{r}{1 + \frac{\kappa x_i^*}{1+2\phi_s\left(h_i^*-0.5\right)}} - \mu\right)$$

x_i^* is the stock level at which the instantaneous average profit is 0 (if $x_i\,(t) > x_i^*$, then $P_i(t) > 0$, resulting in effort entry and the stock being driven down; the reverse occurs if $x_i\,(t) < x_i^*$). h_i^* and f_i^* are found by setting $\frac{dh_i}{dt} = 0$ and $\frac{dx_i}{dt} = 0$, respectively, while assuming $a_i(t) = 1$ for all i.

The no-fishing equilibria (all patches are closed) for the stock, habitat, and effort are as follows:

$$\text{(A7)}\quad \tilde{x}_i^* = \frac{1 + 2\phi_s\,(h_{\max} - 0.5)}{\kappa}\left[\frac{r}{\mu} - 1\right], \qquad \tilde{h}_i^* = h_{\max}, \qquad \tilde{f}_i^* = 0$$

where $\tilde{f}_i^* = 0$ since no effort is allowed in closed patches. With no fishing activity, habitat can recover fully to its maximum quality, $h_{\max}$. Finally, $\tilde{x}_i^*$ is found by setting $\mathrm{d}x_i/\mathrm{d}t = 0$ with $a_i(t) = 0$ for all i.

The parameters marked with a "$*$" in Table 1 are solved by using eqs. A6 and A7 using the following procedure: (1) $h_{\max}$ is set to $\tilde{h}_i^*$, (2) κ is determined using the equation for $\tilde{x}_i^*$, (3) q is determined using the equation for f_i^*, (4) γ is determined using the equation for h_i^*, and finally (5) p is determined using the equation for x_i^*.

A.4 Calculating MSY

Assuming the manager controls $\widehat{f}_i^*$, the equilibrium per patch effort level, the resulting equilibrium for stock and habitat quality would be

$$\text{(A8)}\quad \widehat{x}_i^* = \frac{1 + 2\phi_s\left(\widehat{h}_i^* - 0.5\right)}{\kappa}\left[\frac{r}{\mu + q\widehat{f}_i^*} - 1\right], \qquad \widehat{h}_i^* = \frac{\beta h_{\max}}{\gamma \widehat{f}_i^* + \beta}$$

The manager will pick the $\widehat{f}_i^*$ that maximizes equilibrium yield (MSY). The optimal (per patch) effort level will be denoted by f_i^{MSY}, and the resulting (per patch) stock and habitat quality under MSY will similarly be denoted as x_i^{MSY} and h_i^{MSY}, respectively. The nonspatial maximum sustainable yield exploitation rate, used in parameterizing our model, is then defined as $\mathrm{MSY}/\sum_{i=1}^{n} x_i^{\mathrm{MSY}}$.