Optimal stock–enhancement of a spatially distributed renewable resource

Thorsten Upmann [a,b,e,*], Hannes Uecker [c], Liv Hammann [c], Bernd Blasius [a,d]

[a] Helmholtz-Institute for Functional Marine Biodiversity, University of Oldenburg (HIFMB), Ammerländer Heerstraße 231, Oldenburg 23129, Germany
[b] Faculty of Business Administration and Economics, Bielefeld University, Germany
[c] Institut für Mathematik, University of Oldenburg, Campus Wechloy, Carl-von-Ossietzky-Straße 9–11, Oldenburg 26111, Germany
[d] Institute for Chemistry and Biology of the Marine Environment, University of Oldenburg, Campus Wechloy, Carl-von-Ossietzky-Straße 9–11, Oldenburg 26111, Germany
[e] CESifo, München, Germany

A B S T R A C T

We study the economic management of a renewable resource, the stock of which is spatially distributed and moves over a discrete or continuous spatial domain. In contrast to standard harvesting models where the agent can control the take-out from the stock, we consider the case of optimal stock enhancement. In other words, we model an agent who is, either because of ecological concerns or because of economic incentives, interested in the conservation and enhancement of the abundance of the resource, and who may foster its growth by some costly stock–enhancement activity (*e.g.*, cultivation, breeding, fertilizing, or nourishment). By investigating the optimal control problem with infinite time horizon in both spatially discrete and spatially continuous (1D and 2D) domains, we show that the optimal stock–enhancement policy may feature spatially heterogeneous (or patterned) steady states. We numerically compute the global bifurcation structure and optimal time-dependent paths to govern the system from some initial state to a patterned optimal steady state. Our findings extend the previous results on patterned optimal control to a class of ecological systems with important ecological applications, such as the optimal design of restoration areas.

1. Introduction

The ongoing decline of natural resources means that the problem of the optimal management of ecological systems is of ever increasing importance. The identification of optimal control measures is a challenging task because any intervention in the system will usually have immediate and intertemporal non-linear effects on the states of the ecosystem, and thus on the set of policies available in the future. This issue becomes even more intricate for ecological systems that extend in

---

space, with state variables distributed either over a set of discrete habitat patches or over a continuous spatial domain. The economic management of such spatially extended systems poses a particular challenge because the time path of control measures has to be chosen at every point in space. This raises the question of how management activities should be optimally distributed in space in the course of time. Specifically, is it beneficial to spread control measures as much as possible, while allocating an equal share of management effort to every location? Or, could there be situations where it is best to focus control measures to particular locations, at the cost of neglecting other areas? The answer to these questions is not only of theoretical interest but also has immediate practical consequences (*e.g.*, for the optimal design of nature reserves and restoration areas).

To find optimal spatio-temporal management strategies, optimal control theory has been applied to different ecological fields, such as the management of invasive species (Blackwood et al., 2010; Albers et al., 2018), epidemic spread (Tildesley et al., 2006; Rowthorn et al., 2009), spatial pollution in shallow lakes (Brock and Xepapadeas, 2008; Grass and Uecker, 2017), semiarid vegetation systems (Brock and Xepapadeas, 2010; Uecker, 2016), and travelling-and-harvesting problems (Behringer and Upmann, 2014; Belyakov et al., 2015; Upmann and Behringer, 2020). However, to date, most ecological applications of optimal control theory focus on the case of optimal fisheries (*e.g.* Herrera et al., 2016; Kelly et al., 2019; Grass et al., 2019). In particular, Neubert (2003), Neubert and Herrera (2008) and Ding and Lenhart (2009a) find spatially heterogeneous distributions of the fishing effort to be optimal, and show that no-take regions (*i.e.*, the concentration of fishing activities outside those localized regions) are typically part of an optimal fishery management strategy.[1]

Yet, these authors assume that the stocks equal zero at the boundary (Dirichlet boundary conditions), postulating that the fish are instantly killed as soon as they touch the shore (*i.e.*, totally hostile boundary), which is a questionable assumption from an ecological perspective. From a theoretical point of view, spatially heterogeneous solutions do not come as a surprise because they do not correspond to a true pattern formation but are rather enforced by the choice of Dirichlet boundary conditions; for, this type of boundary conditions negates the homogeneity of the spatial domain and implies an inhomogeneous steady state distribution of the fish stocks (unless they become extinct in the whole domain). So, with Dirichlet boundary conditions, the spatial heterogeneity of the optimal solution is directly imposed by the model. Moreover, in these models, the advantage of heterogeneous control measures is only gained in the vicinity of the boundary, so that with increasing size of the habitat the relative merit of heterogeneous control vanishes (see Moeller and Neubert, 2013). To avoid those ecologically and theoretically unsatisfactory effects, other authors dispense with the assumption of a hostile boundary and instead propose no-flux at the boundary of the domain (Neumann boundary conditions). This models a situation where stocks can live at, but cannot transgress, the boundary of the habitat, and thus captures situations where natural (*e.g.*, shores) or man-made barriers (*e.g.*, highways) constitute boundaries of the habitat. In this case, spatially homogeneous solutions are compatible with the model—and the question of whether or not spatially patterned solutions are optimal becomes meaningful and significant.

The possibility of spatially patterned solutions evokes much interest because their emergence contrasts with the intuition suggesting that diffusion contributes to spatially homogeneity, and thus to the stability of the homogeneous solution. Notably, Brock and Xepapadeas (2008, 2010) systematically investigate the emergence of heterogeneous solutions in spatial optimal control problems with an infinite time horizon. These authors show that heterogeneity may arise in a Turing like bifurcation where a spatially homogeneous steady state becomes unstable in the presence of diffusion. This version of diffusion-induced instability differs from the classical Turing mechanism (Turing, 1952) because here the instability results from optimizing behaviour. (For an excellent presentation of the Turing mechanism, see Murray, 2003, Ch. 2.) Brock and Xepapadeas (2008) term this type of local instability of an optimally controlled system with diffusion, that is, the instability of the optimal steady states, *optimal diffusion-induced instability*. Consequently, this type of instability in spatial optimal control models is a true pattern formation: while in the absence of diffusion the optimal steady state is necessarily homogeneous, or flat, a spatially homogeneous steady state may no longer be optimal in the presence of diffusion. If diffusion becomes sufficiently strong, the homogeneous solution becomes unstable and converges to a patterned steady state; such a steady state could be optimal, and is then called a *patterned optimal steady state* (POSS). In a POSS, the agent finds it optimal to concentrate the (fishing, harvesting) activity on some spots, and to curtail it on the rest of the region. Intuitively, the observation that spatially heterogeneous, and thus spatially focused fishing efforts may be optimal, can be explained by the fact that locally reduced fishing rates generate source–sink dynamics, which helps fish stocks to recover when they are overexploited: a high diffusion rate makes the high abundances in no-take zones quickly available in the catching zones.

While most studies of optimal diffusion–induced instabilities are concerned with describing the possible existence of a POSS, there has been far less work to address the problem of how to dynamically control the system towards an optimal steady state given some initial spatial configuration. As shown by Uecker (2016) and Grass and Uecker (2017), this is further complicated by the fact that typically multiple steady states exist, many of which are not stable; while even in systems that exhibit locally stable patterned control states, these are not necessarily optimal. Another open question concerns the generality of this effect (*i.e.*, of the emergence of POSSs). As shown in Brock and Xepapadeas (2008, Theorem 1), optimal diffusion–induced instabilities occur only under specific features of the model. Namely, the literature finds a POSS only when the control actions generate an externality (*e.g.*, Moeller and Neubert, 2013, explore the case when fishing damages the

---

[1] These findings have subsequently been confirmed and further extended in both spatially continuous (Ding and Lenhart, 2009b) and spatially discrete (two-patch or multi-patch) models (Moberg et al., 2015; Sanchirico et al., 2006; Sanchirico and Wilen, 1999; 2001).

habitat), or in multidimensional models where the control reduces the stock (*e.g.*, the harvesting model of Brock et al., 2014). In contrast, little is known about whether or not optimal diffusion–induced instability, and subsequently the emergence of a POSS, extends to other classes of models in resource and environmental economics—let alone in other areas of economics.[2]

In this paper, we investigate the management of a spatially distributed renewable natural resource where an agent or a policy maker is directly interested in the abundance of a resource, rather than in its yield or depletion, and is able to contribute to its growth and prosperity by means of some spatially targeted stock–enhancement policy (*e.g.*, cultivation, breeding, fertilizing, or nourishment). This management problem is complimentary to the familiar optimal harvesting problem, which has largely been analysed in the economics literature, where the resource is reaped to maximize the discounted profit or the utility stream of an agent.[3] Here, we model a situation where an agent has an immediate interest in ecological conservation and resource preservation, but has little or no interest in consuming the stock. Think of a population of penguins, which are hardly edible and have, once killed, little economic value, yet people protect and support their stocks.— We have in mind a policy maker (or an environmental agency) who is concerned about low abundances of species, such as species that are threatened and endangered with extinction,[4] and thus aims at protecting and rebuilding those stocks.

To derive and characterize the spatial and intertemporal characteristics of the optimal stock–enhancement policy, we apply Pontryagin's maximum principle. In addition, we explore the sensitivity of this policy with respect to system parameters. We show that although the interest of the agent is different from the familiar harvesting model, the presence of spatial couplings may induce an optimal diffusion–induced instability, which leads to spatially heterogeneous optimal stock–enhancement policies (*i.e.*, leading to a POSS. While a POSS in a harvesting model means the simultaneous presence of regions with high fishing activities and of protected areas with low take-out levels to safeguard sustainable natural resources, a POSS in a stock–enhancement model means the simultaneous presence of regions with high efforts of breeding and ecological restoration and of those with low efforts where the ecosystem basically remains in its natural conditions.

After introducing the model in a non–spatial setting (Section 2), we demonstrate the emergence of a POSS in two alternative spatial geometries, namely a system of two coupled discrete habitat patches (Section 3) and a spatially continuous, diffusive system in one- or two-dimensional bounded domains (Sections 4 and 5). The reasons for the different settings are as follows: in the non–spatial setting, we introduce the basic dynamics and objective; moreover, we can apply the ODE theory from Tauchnitz (2015) to rigorously derive the canonical system as the necessary first order optimality condition. This theory also applies to the two–patch ODE model, which is the simplest spatial case. For the PDE cases, the derivation of the canonical systems proceeds somewhat formally (see Appendix); the 1D case is numerically less expensive than 2D, and the results are easier to visualize, yet, as we show, they naturally extend to the 2D case. For all three spatial geometries, we use the numerical continuation and bifurcation tool `pde2path` (Uecker et al., 2014) to compute an overview of the steady states of the canonical systems, their stability, their spatial structure and the corresponding objective values. Importantly, in the interesting parameter range, our spatial models have multiple, homogeneous and heterogeneous, steady states. Then, we use the algorithms from Uecker and de Witt (2019) to compute paths from some initial spatial distributions of the resource to a specified target steady state. These paths are economically fundamental because they describe how to govern the system optimally to the specified steady state. They thus model the policies to be followed during a transition period, a period that may last for quite a long time and may thus, from a political perspective, be quite important. While our main result is that in a large parameter domain the optimal policy is to govern the system to a POSS, we also show how this steady state can be reached in an optimal way when we start at some historically given non-optimal situation.

2. The basic model of stock–enhancement

To begin with, and to introduce the terminology, we consider the non-spatial version of our model. We model the dynamics of a renewable resource of stock size x which is growing according to the logistic growth function $g(x) = ux\left(1 - \frac{x}{K}\right)$ with carrying capacity $K > 0$ and time-dependent growth rate $u(t) > 0$ that can be controlled externally by the agent. We assume that the stock is additionally reduced by some constant exogenous loss process with rate $\delta \geq 0$, describing mortality effects that are not already captured in the logistic growth term (*e.g.*, environmental stressors or constant harvesting efforts). Without loss of generality, we can set $K = 1$, which means that we measure the stock in units of its carrying capacity, yielding the ordinary differential equation (ODE)

$$\dot{x}(t) = f(x(t), u(t)) \equiv u(t)x(t)(1 - x(t)) - \delta x(t), \tag{1a}$$

with $x(0) = x_0 > 0$ representing the stock at the beginning of the planning period $\mathcal{T} \equiv [0, \infty)$. For any constant $u > \delta$, the stock equilibrates to the steady–state level $x_\infty = (u - \delta)/u < 1$, while for $u < \delta$ the stock goes extinct, $x_\infty = 0$. Thus, we model a resource under strong loss processes or environmental stress that needs external support $u > \delta$ to survive.

The agent derives utility from the presence of higher stocks, and is thus interested in the stocks *in situ*. More precisely, utility is logarithmic in x, which represents the idea that utility is increasing and concave. The agent is able to enhance the

---

living conditions of the resource by some stock–enhancement activity, such as cultivation, breeding, fertilizing, or nourishment. The cost of this activity is increasing and convex, namely quadratic, in u. Then, the instantaneous utility of the agent is

$$J_c(x(t), u(t)) \equiv \log(x(t)) - \frac{\gamma}{2} u^2(t), \tag{1b}$$

with γ the marginal cost of the stock–enhancing activity, and the optimization problem of the agent reads as

$$\max_{u \in \mathcal{U}} J(x(\cdot), u(\cdot)), \quad \text{where} \quad J(x(\cdot), u(\cdot)) \equiv \int_0^\infty e^{-\rho t} J_c(x(t), u(t)) \, dt, \tag{1c}$$

subject to (1a), where $\rho > 0$ denotes the discount rate of the agent. (The technical details about the class $\mathcal{U}$ of admissible controls are given in Appendix A. Moreover, in Appendix B we comment on the minor changes if $\log(x)$ in (1b) is replaced by $e_\eta(x)$ with e_η from the the isoelastic family $e_\eta(x) = (x^{1-\eta} - 1)/(1 - \eta)$ with η near 1.)

To derive the optimal policy we define the Hamiltonian

$$H(x, u, \lambda) \equiv J_c(x, u) + \lambda f(x, u) = \log(x) - \frac{\gamma}{2} u^2 + \lambda(ux(1 - x) - \delta x),$$

where λ denotes the costate (or shadow price) of x. The standard intertemporal transversality condition is (again see Appendix A for technical details)

$$\lim_{t \to \infty} e^{-\rho t} x(t) \lambda(t) = 0. \tag{2}$$

The first order necessary conditions for an optimal control u^* and associated solution x^* are rebuilding those stocks. To derive and characterize

$$H(x^*(t), u^*(t), \lambda(t)) = \max_u H(x^*(t), u, \lambda(t)), \tag{3a}$$

$$\dot{\lambda}(t) = \rho \lambda(t) - \partial_x H(x^*(t), u^*(t), \lambda(t)), \tag{3b}$$

jointly with (1a). Eq. (3a) says that u^* locally maximizes H, Eq. (3b) represents the costate (or adjoint) equation. Here, because H is concave in u, the local maximization (3a) yields (dropping the superscript $*$ of x)

$$u^*(t) = \frac{\lambda(t)}{\gamma} x(t)(1 - x(t)). \tag{4}$$

So, the optimal stock–enhancing activity is proportional to the instantaneous growth of the stock, $x(t)(1 - x(t))$, with a factor of proportionality given by the ratio of the value of the resource $\lambda(t)$ and the marginal cost of enhancement γ. By substituting u^* into the state Eq. 1a and the costate Eq. 3b, we obtain the *canonical system* for the state x and the costate λ:[5]

$$\dot{x} = \frac{\lambda}{\gamma} x^2 (1 - x)^2 - \delta x, \tag{5a}$$

$$\dot{\lambda} = \lambda(\delta + \rho) - \frac{\lambda^2 x(1 - 2x)(1 - x)}{\gamma} - \frac{1}{x}, \tag{5b}$$

with the initial condition $x(0) = x_0$ for the state, and the transversality condition (2) for the costate. We have eliminated the control from (5), and thus henceforth use the notation $X \equiv (x, \lambda)$. Finally, by substituting $u^* = u(X)$ into the objective function $J = J(x, u^*) = J(x, u(X))$, we may write, with minor sloppiness, $J(X)$.

We call a steady state of the canonical system (5), $X_\infty \equiv (x_\infty, \lambda_\infty)$, a *canonical steady state* (CSS); and a solution $t \mapsto X(t) \equiv (x(t), \lambda(t))$ of (5) such that $x(0) = x_0$ and $\lim_{t \to \infty} X(t) = X_\infty$ a *canonical path* (CP) from the initial state x_0 to the specified CSS X_∞. For our numerical computations, we subsequently assume a moderate value of the discount rate, $\rho = 0.025$; to make the enhancement activity costly, we set $\gamma = 10$, which amounts to a marginal cost of u equal to 5; and finally, since the growth rate can be controlled by the agent, it seems natural to set, for the moment, the death rate δ equal to unity (but we will use δ as a bifurcation parameter later anyway). So, our parameter specification is: $\delta = 1$, $\rho = 0.025$, and $\gamma = 10$. Using this, the canonical system (5) has a unique CSS X_∞, given by the stock $x_\infty = 0.0596$, the shadow prize $\lambda_\infty = 189.61$ and the associated control $u_\infty^* = 1.063$. As observed earlier, in the steady state the control must surmount the death rate: $u > \delta$. This CSS yields an instantaneous utility $J_c(X_\infty) = -8.47$ and a total discounted profit $J(X_\infty) = \frac{1}{\rho} J_c(X_\infty) = -338.95$.

The two-component ODE system (5) can also be conveniently discussed in the x–λ phase plane, see Fig. 1(a). The organizing centre is the CSS X_∞, which is a saddle point with stable manifold $W_s(X_\infty)$ (green line) and unstable manifold $W_u(X_\infty)$ (orange line). We want to stress that the canonical system (5) is not an initial value problem because only the initial state

---

[5] The (maximised) Hamiltonian (3a) is not concave in the state variable x, therefore we cannot apply standard sufficiency theorems.

of the stock value x_0 is given, while the initial costate $\lambda_0 \equiv \lambda(0)$ is free. Given x_0, the transversality condition (2) implies that optimal solutions of (5) must converge towards X_∞, because all other solutions diverge super-exponentially. For general initial values (x_0, λ_0) that are not in the stable manifold $W_s(X_\infty)$ the associated solution of (5) will diverge to $+\infty$, and thus, given x_0, the agent needs to find an intersection (which here is unique) of the line $x = x_0$ with $W_s(X_\infty)$.[6] For these initial values, the solution follows the CP along the stable manifold towards the CSS.

More generally, given a $2n$ dimensional canonical system with possibly several CSS, by the *domain of attraction* of a CSS $X_\infty \in \mathbb{R}^{2n}$ we mean the set

$$\mathcal{A}(X_\infty) := \{x \in \mathbb{R}^n : \text{there exists, by suitable choice of } \lambda_0 \in \mathbb{R}^n, \text{ a CP from } x \text{ to } X_\infty\}.$$

A CSS X_∞ is called *globally stable* if $\mathcal{A}(X_\infty) = \mathbb{R}^n$ (or if $\mathcal{A}(X_\infty) = \mathcal{B}$ if the dynamics is restricted to some subset $\mathcal{B} \subset \mathbb{R}^n$, such as $\mathcal{B} = \mathbb{R}_+^n$). These definitions do not involve the optimality of a CP and if several CSS exist (as in the spatial examples below), then their domains of attraction may have non–empty intersections. Then, given an initial state $x \in \mathcal{A}(X_\infty^{(a)}) \cap \mathcal{A}(X_\infty^{(b)})$, the first step is to compare the value $J(x; X_\infty^{(a)})$ of the CPs to $X_\infty^{(a)}$ with the value $J(x; X_\infty^{(b)})$ of the CP to $X_\infty^{(b)}$. If $J(x; X_\infty^{(a)}) = J(x; X_\infty^{(b)})$, and if no better solution starting at x (*i.e.*, a control u yielding a higher value) exist, then x is called a Skiba point (Skiba, 1978), and the set of all Skiba points between $X_\infty^{(a)}$ and $X_\infty^{(b)}$ yields a manifold, called Skiba manifold.

Problem (1a)–(1c) has a unique CSS X_∞, with $\mathcal{A}(X_\infty) = \mathbb{R}_+$, cf. Fig. 1(a). Moreover, Fig. 1(b) shows the time–dependent CP $X(t) \equiv (x(t), \lambda(t))$ starting at $x_0 = 0.03$, and the path of the associated optimal control $u^*(t)$. Here, the initial stock is lower than its steady state; that is, $x_0 = 0.03 < 0.0596 = x_\infty$. This implies that at the beginning of the period, the agent has to spend a rather high cultivation effort to increase the stock, while this effort may be reduced as the stock increases over time. Likewise, the shadow price of the stock is initially quite high but monotonously decreases as the stock becomes more abundant.

3. Stock–enhancement in a two-patch model

Let us now consider a system of two diffusively coupled patches with stocks x_1 and x_2 that disperse between the patches with constant per-capita dispersal rate D. (Think of two penguin colonies at two distinct locations, where individual penguins can migrate from one patch to another, and net migration is towards the location of low biomass concentration.) The population dynamics of the two patches are identical (*i.e.*, they have the same carrying capacity K and exogenous loss rate δ) and can be controlled in the same way by patch-specific stock-enhancement activities u_1 and u_2. Assuming that the stocks are again measured in units of the carrying capacity, the state equation (1a) generalises to (suppressing the time argument)

$$\dot{x}_i = f_i(\mathbf{x}, u_i) \equiv u_i x_i (1 - x_i) - \delta x_i + D(x_{3-i} - x_i), \qquad i = 1, 2 \tag{6a}$$

with $\mathbf{x} \equiv (x_1, x_2)$. Letting $\mathbf{u} \equiv (u_1, u_2)$, a straightforward extension of (1b) yields the instantaneous utility

$$J_c(\mathbf{x}, \mathbf{u}) \equiv \sum_{i=1}^{2} J_{c,i}(x_i, u_i), \quad \text{with} \quad J_{c,i}(x_i, u_i) \equiv \log(x_i) - \frac{\gamma}{2} u_i^2.$$

The problem of the agent is

$$\max_{u \in \mathcal{U}^2} J(\mathbf{x}(\cdot), \mathbf{u}(\cdot)), \quad \text{where} \quad J(\mathbf{x}(\cdot), \mathbf{u}(\cdot)) \equiv \int_0^\infty e^{-\rho t} J_c(\mathbf{x}(t), \mathbf{u}(t)) \, \mathrm{d}t, \tag{6b}$$

subject to the stock dynamics (6a). Then, the Hamiltonian reads as

$$H(\mathbf{x}, \mathbf{u}, \lambda) \equiv J_c(\mathbf{x}, \mathbf{u}) + \lambda \cdot \mathbf{f}(\mathbf{x}, \mathbf{u}),$$

with the costate vector $\lambda \equiv (\lambda_1, \lambda_2)$, stock dynamics $\mathbf{f} \equiv (f_1, f_2)$, and the limiting intertemporal transversality condition $\lim_{t \to \infty} e^{-\rho t} \mathbf{x}(t) \cdot \lambda(t) = 0$, where again we can apply Tauchnitz (2015), cf. Appendix B. Maximising H with respect to $\mathbf{u}$ yields the optimal controls $u_i^* = \frac{\lambda_i}{\gamma} x_i (1 - x_i)$. As expected, we obtain the same rule for the optimal stock–enhancement policy as in the one-patch case. However, this does not rule out (as we shall see) that we may find solutions that differ from the optimal control in the one-patch case.

Now, by substituting the optimal controls into the state equations (6a) and the costate equations $\dot{\lambda}_i = \rho \lambda_i - \partial H / \partial x_i$, we obtain the canonical system

$$\dot{x}_i(t) = \frac{\lambda_i}{\gamma} x_i^2 (1 - x_i)^2 - \delta x_i + D(x_{3-i} - x_i), \tag{7a}$$

---

[6] In a 2-component ODE (scalar state ODE), this can be done by (essentially) integrating backward in time from the stable eigenspace but in higher dimensions—for example, in case of two patches (see Section 3), or in case of high dimensional ODEs obtained from spatial discretizations of PDEs (see Sections 4 and 5)—such computations of stable manifolds become infeasible. See Grass et al. (2008) and Uecker (2016) for comments on the numerical method for computing canonical paths by truncation to large T and suitable asymptotic boundary conditions, and how this relates to the saddle point property of CSSs.

Table 1

Properties of the CSSs of system (7) with parameter values (8). The system exhibits a flat CSS $\mathbf{X}_\infty^f$ and a pair of POSSs, $\mathbf{X}_\infty^p$ and $\widetilde{\mathbf{X}}_\infty^p$, which are identical up to the exchange of patch indices.

CSS	$x_{1,2}$	$\lambda_{1,2}$	$u_{1,2}^*$	J	SPP	Remarks
	0.0596	189.61	1.063			not optimal, dominated by the path to $\mathbf{X}_\infty^p$ with value $J_{CP} = -602.6$
Flat: $\mathbf{X}_\infty^f$				-677.9	no	
	0.0596	189.61	1.063			
	0.086	166.37	1.305			optimal, hence a POSS and "almost" globally stable
Patterned: $\mathbf{X}_\infty^p$, $\widetilde{\mathbf{X}}_\infty^p$				-601.1	yes	
	0.0196	82.5	0.158			
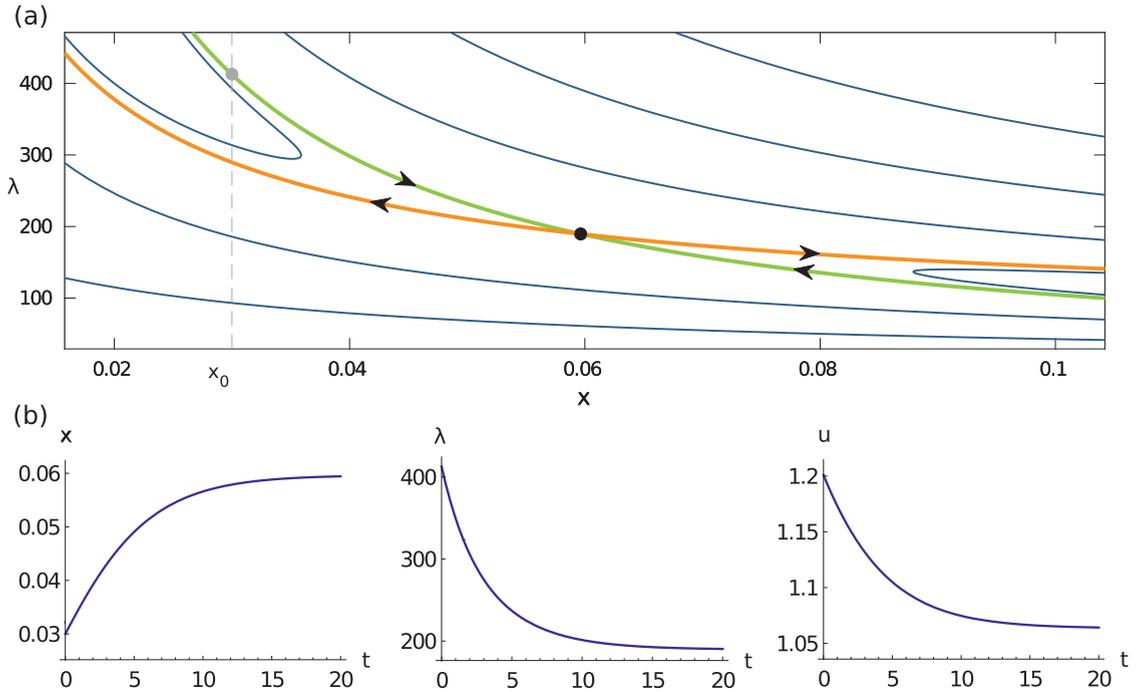
Fig. 1. Optimal solution of the non-spatial stock-enhancement system. (a) Phase portrait of the canonical system (5), showing the saddle point X_∞ (black point), and its stable (green) and unstable manifold (orange). Selected orbits are shown in blue. The vertical-dashed line marks the initial stock value $x_0 = 0.03$. In the optimal orbit, the initial costate λ_0 is chosen to be located exactly on the intersection of that line with the stable manifold of X_∞, yielding $\lambda_0 = 412.64$ (grey point). (b) Time dependence of the optimal orbit (x, λ) and of the optimal control u^* for the initial value $x_0 = 0.03$. The canonical path corresponds to the part of the stable manifold between the grey and the black point. (For the visibility of the colouring the reader is referred to the web version of this article.)

$$\dot{\lambda}_i(t) = \lambda_i(\delta + \rho) - \frac{\lambda_i^2 x_i(1 - 2x_i)(1 - x_i)}{\gamma} - \frac{1}{x_i} - D(\lambda_{3-i} - \lambda_i), \tag{7b}$$

with initial values $x_i(0) = x_{i,0}$. Remarkably, while the spatial coupling in the state equation (7a) is described by the 'diffusion' term $D(x_{3-i} - x_i)$, the coupling in the costate equation (7b) is described by the 'anti-diffusion' term $-D(\lambda_{3-i} - \lambda_i)$. Thus, while the diffusion-driven resource flow is from high towards low concentrations (*i.e.*, from the crowded to the less crowded area), an optimal stock–enhancing policy makes the shadow price of the resource 'move' from low to high prices.

Irrespective of the value of the dispersal rate $D \geq 0$, the canonical system (7) has the homogeneous (or flat) CSS: $\mathbf{X}_\infty^f = (x_\infty, x_\infty, \lambda_\infty, \lambda_\infty)$ (*i.e.*, twice the CSS X_∞ of the single patch model of Section 2). By setting $D = 0.25$ and using the same parameter values as in the previous section:

$$D = 0.25, \quad \rho = 0.025, \quad \gamma = 10, \quad \delta = 1, \tag{8}$$

we obtain $x_{1\infty} = x_{2\infty} = 0.0596$, $\lambda_{1\infty} = \lambda_{2\infty} = 189.61$, and $u_{1\infty}^* = u_{2\infty}^* = 1.0634$, which exactly replicates the values from the non-spatial model. Correspondingly, the instantaneous benefit and the total profit for this solution are simply doubled compared to the non-spatial model: $J_c(\mathbf{X}_\infty^f) = 2J_c(X_\infty) = -16.95$ and $J(\mathbf{X}_\infty^f) = 2J(X_\infty) = -677.9$. In addition, we find a heterogeneous (or patterned) CSS $\mathbf{X}_\infty^p = (x_1, x_2, \lambda_1, \lambda_2)$ for $i = 1, 2$, where the values of the stocks, $x_1 = 0.086$, $x_2 = 0.0196$, the costates, $\lambda_1 = 166.37$, $\lambda_2 = 82.505$, and the optimal controls, $u_1^* = 1.305$, $u_2^* = 0.158$, differ in the two patches. The resulting instantaneous benefit is $J_c(\mathbf{X}_\infty^p) = -15.03$, the discounted profit is $J(\mathbf{X}_\infty^p) = -601.1$, and hence the patterned CSS yields
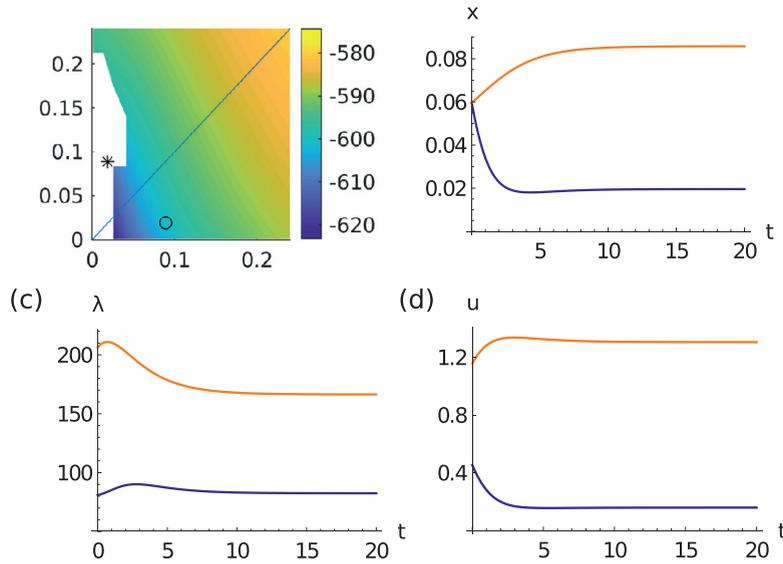
Fig. 2. (a) Domain of attraction $\mathcal{A}(\mathbf{X}_\infty^p)$, and values of the CP from $\mathbf{X}(0) \in \mathbb{R}_+^2$ to $\mathbf{X}_\infty^p$, marked by $\circ$ for the two patch model with parameter values (8). Initial states for which no CP to $\mathbf{X}_\infty^p$ could be found are left white. The mirrored $\widetilde{\mathbf{X}}_\infty^p$ is marked by $*$, and $\mathcal{A}(\widetilde{\mathbf{X}}_\infty^p)$ and the values of CPs are naturally obtained by mirroring the plot at the diagonal. (b)–(d) CP from the stock of $\mathbf{X}_\infty^f$ to $\mathbf{X}_\infty^p$; time dependence of the state variables, the costates and the optimal controls for both patches.

a higher utility compared to the flat CSS. Naturally, by exchanging the two patches, we find a second equivalent patterned CSS $\widetilde{\mathbf{X}}_\infty^p = (x_2, x_1, \lambda_2, \lambda_1)$.[7]

Having found these three CSSs, the next question refers to their stability and domains of attraction. That is, given some initial state $\mathbf{x}_0$, which of these three can be reached by a CP from $\mathbf{x}(0) = \mathbf{x}_0$, and if we can reach more than one, which of the CPs yields the highest profit. A simple visualization such as the phase plane plot in Fig. 1 is no longer possible because the canonical system of the two-patch model is 4–dimensional. Locally, a CSS X_∞ of a $2n$–dimensional canonical system has the *saddle point property* (SPP) if the dimension of its stable manifold $\dim W_s(X_\infty)$ equals n; correspondingly, the number $d \equiv n - \dim W_s(X_\infty)$ is called the defect of the CSS (see Grass et al., 2008), and these numbers can be computed by linearization of (6a) around X_∞. The dimension of the stable manifold, $\dim W_s(X_\infty)$, equals the number of eigenvalues of the linearization with negative real parts.

For our chosen parameter values, only the patterned CSS $\mathbf{X}_\infty^p$ (and naturally also $\widetilde{\mathbf{X}}_\infty^p$) has the SPP. In contrast, the flat CSS $\mathbf{X}_\infty^f$ does not have the SPP: the dimension of its stable manifold equals $\dim W_s(\mathbf{X}_\infty^f) = n_s = 1$, and thus its defect $d = 2 - 1 > 0$. Consequently, canonical paths to the flat CSS $\mathbf{X}_\infty^f$ only exist for particular values of $\mathbf{x}_0$ on a one-dimensional curve; namely, those for which there exists a $\boldsymbol{\lambda}(0)$ such that $(\mathbf{x}(0), \boldsymbol{\lambda}(0))$ is located in the stable manifold of $\mathbf{X}_\infty^f$. Such a defective CSS (*i.e.*, when $d > 0$) may still be optimal if a canonical path to another CSS does not exist.

Meanwhile, the patterned CSS $\mathbf{X}_\infty^p$ is "almost globally stable modulo symmetry". By this we mean that for all $\mathbf{x}_0 \in \mathbb{R}_+^2$, possibly except for a small set near $(0,0)$ (see below), initial costates $\boldsymbol{\lambda}(0)$ can be found such that a canonical path to $\mathbf{X}_\infty^p$ or $\widetilde{\mathbf{X}}_\infty^p$ (or both) exist. In Fig. 2(a), we show the values of the CPs from some arbitrary initial point $\mathbf{x}_0$ to $\mathbf{X}_\infty^p$, with the x–values of $\mathbf{X}_\infty^p$ marked by a $\circ$, and the white part indicating the set of initial states for which a CP to $\mathbf{X}_\infty^p$ could not be found. This includes, for instance, $\widetilde{\mathbf{X}}_\infty^p$, indicated by the $*$. However, $\mathcal{A}(\mathbf{X}_\infty^p)$ contains at least the coloured set, and naturally $\mathcal{A}(\widetilde{\mathbf{X}}_\infty^p)$ is obtained from mirroring along the diagonal, such that by symmetry the diagonal is a Skiba manifold between $\mathbf{X}_\infty^p$ and $\widetilde{\mathbf{X}}_\infty^p$. For $x_0 \in \mathcal{A}(\mathbf{X}_\infty^p) \cap \mathcal{A}(\widetilde{\mathbf{X}}_\infty^p)$, this Skiba manifold yields the (expected) result that below the diagonal it is optimal to go to $\mathbf{X}_\infty^p$, and above the diagonal to $\widetilde{\mathbf{X}}_\infty^p$. Together, $\mathcal{A}(\mathbf{X}_\infty^p) \cup \mathcal{A}(\widetilde{\mathbf{X}}_\infty^p)$ cover $\mathbb{R}_+^2$, except for a small set near $x = 0$, but we believe that this is due to our numerical method. We discretize the positive quadrant by a grid $(x_1, x_2)_{ij}$ and for each initial state of that grid aim to compute the CP to $\mathbf{X}_\infty^p$ with a fixed method, such as with a fixed truncation time T. This leaves open a region near $(x_1, x_2) = (0, 0)$ (where as in Fig. 1 the initial co–states go to ∞), but by tweaking the numerics (*e.g.*, using a finer grid of initial states, and adapting T), we can shrink the non-existence region near $(x_1, x_2) = (0, 0)$.

Suppose that in the past the two patches have been managed by two different agents (or firms), but that an integrated management of both patches will be considered from now on. Then, the question is how can we best govern the integrated system of the two patches to the optimal (which here means to the patterned, steady state)? However, this is a question about the best policy during the transition time from the flat CSS to the patterned CSS—or about the canonical path. Fig. 2(b)

---

[7] The CSSs and the CPs discussed were subsequently found using the toolbox `pde2path` by first treating (7) as a bifurcation problem for steady states and then computing the CPs by a truncated continuation problem in the initial states, but we postpone these bifurcation and computation issues to the next Section.

displays the CP from the initial stock $\mathbf{x}_0 = \mathbf{x}_\infty^f = (x_\infty, x_\infty)$ of the flat CSS (on the diagonal) to the patterned CSS $\mathbf{X}_\infty^p$. The discounted profit of this CP is $J_{\mathrm{CP}} = -602.6$, which is higher than the discounted profit $J(\mathbf{X}_\infty^f) = -677.9$ of the starting CSS. Thus, we may in fact characterize the patterned CSS $\mathbf{X}_\infty^p$ as a patterned *optimal* steady state (POSS), unique modulo the symmetry of interchanging patch 1 and 2.

These computations are based on the assumption of a perfectly homogeneous system. This symmetry is, of course, broken if any of the model parameters depends on the patch i, such as by assuming different exogenous loss rates $\delta_1 \neq \delta_2$. If the symmetry breaking is weak (*i.e.*, $\delta_1 \approx \delta_2$) then the flat CSS survives as an 'almost flat' CSS, and the two equivalent patterned CSS $\mathbf{X}_\infty^p$ and $\widetilde{\mathbf{X}}_\infty^p$ become two different patterned CSS.[8] It then becomes an interesting task to characterize their domains of attractions, which are also no longer related by simple symmetry. However, our point here is to deliberately focus on the symmetric (*i.e.*, spatially homogeneous) case. Our results show that even though the model is spatially homogeneous, positive dispersal rates may give rise to a spatially non-homogeneous (or patterned) optimal solution. This result is quite surprising, because dispersal, with movement from the crowded to the less crowded patch, *per se* tends to flatten out the stock over the spatial domain. Thus, a rather simple stock–enhancement model with two patches may bring about counter-intuitive optimal policies. The natural, though tentative, policy proposal calling for an equal distribution of the efforts (and thus of the cost) may be misled.

4. Stock–enhancement in a one-dimensional continuous space

We now extend our model to a continuous one–dimensional (1D) space, and assume that the resource is continuously distributed over a bounded interval $\Omega \subset \mathbb{R}$. The agent may then choose the stock–enhancing activity at any spatial location $z \in \Omega$; that is, the agent can select (at each instant of time) a function $u(\cdot, t)$ on Ω, rather than a scalar $u(t)$, as in the non-spatial model, or a 2D vector $\mathbf{u}(t)$, as in the two-patch model. Accordingly, we generalise the stock dynamics (6a) to the reaction–diffusion equation

$$\partial_t x(z, t) = u(z, t) x(z, t)(1 - x(z, t)) - \delta x(z, t) + D\Delta x(z, t) \tag{9a}$$

with diffusion constant D and initial condition (initial distribution) $x(z, 0) = x_0(z), \ \forall z \in \Omega$.[9] In (9a), we follow Fick's first law of diffusion, postulating that the flux of the biomass goes from regions of high to those of low concentration, where the magnitude of the flux is proportional to the concentration gradient (spatial derivative). Moreover, we assume homogeneous Neumann boundary conditions (BCs), modelling the idea that the stock can live at, but cannot traverse the boundary of the habitat, and thus implying the absence of any exterior in- or outflow (see Appendix A for generalization to Robin BCs):

$$\partial_{\mathbf{n}} x(z, t) = 0, \quad \forall z \in \partial\Omega, \tag{9b}$$

where $\partial_{\mathbf{n}}$ is the outward normal derivative (the derivative taken in the direction orthogonal to the boundary of Ω). In 1D, where the spatial domain is an interval, $\Omega = (a, b)$, we have, suppressing the time argument, $\partial_n x(a) = -x'(a)$ and $\partial_n x(b) = x'(b)$, and $\Delta x(z) = x''(z)$ (see also fn. 9).

Assuming again that the instantaneous pointwise utility is given by (1b)—that is, $J_c(x, u) \equiv \log(x) - \frac{\gamma}{2} u^2$—, we write the spatial average of J_c over Ω as

$$J_{ca}(t) \equiv \frac{1}{|\Omega|} \int_\Omega J_c(x(z, t), u(z, t)) \, \mathrm{d}z,$$

and the agent seeks to maximize the discounted average utility

$$\max_{u \in \mathcal{U}} J(x(\cdot), u(\cdot)) \quad \text{where} \quad J(x, u) \equiv \int_0^\infty e^{-\rho t} J_{ca}(x, u) \, \mathrm{d}t \tag{9c}$$

subject to (9a) and (9b), and with the optimal control as given in (4). The formal Hamiltonian for problem (9) is then given by

$$H \equiv J_{ca}(t) + \frac{1}{|\Omega|} \int_\Omega \lambda(z, t) \, f(x(z, t), u(z, t)) \, \mathrm{d}z, \tag{10}$$

where $f(x, u) \equiv h(x, u) + D\Delta x$ denotes the right-hand side of (9a).

To derive the canonical system for problem (9), we need to take the variational derivative of H with respect to x. To do so, we use integration by parts of the term $\lambda\Delta x$ occurring in H, and apply Pontryagin's maximum principle (see the Appendix for further comments) with the limiting intertemporal transversality condition

$$\lim_{t \to \infty} e^{-\rho t} \int_\Omega \lambda(z, t) x(z, t) \, \mathrm{d}z = 0. \tag{11}$$

---

[8] In the Appendix, we illustrate this for the spatially continuous model by breaking symmetry via boundary conditions.

[9] The Laplace operator $\Delta \equiv \nabla \cdot \nabla \equiv \nabla^2$ denotes the divergence of the gradient of a function f on Euclidean space. In Cartesian coordinates, it is the sum of second partial derivatives of f with respect to the independent variables, which here consist of the spatial coordinates: $\Delta f(z) = \sum_{i=1}^n \partial_{z_i}^2 f(z)$.

Table 2

Properties of the CSSs of the system (9) from the 1D spatially continuous model for $\delta = 1$. The system exhibits a flat CSS and two spatially heterogeneous patterned CSS (compare Fig. 3).

CSS	J	SPP	Remarks
flat f	−338.95	no	not optimal, dominated by the canonical path to p1 with value $J_{\text{CPf1}} = -277.77$
patterned p2	−327.16	no	not optimal, dominated by the canonical path to p1 with value $J_{\text{CP21}} = -279.82$
patterned p1	−277.80	yes	optimal, hence a POSS, with a large (probably global) domain of attraction

This procedure yields the canonical system

$$\partial_t x = ux(1-x) - \delta x + D\Delta x, \quad x(0) = x_0, \tag{12a}$$

$$\partial_t \lambda = \lambda(\delta + \rho) - \lambda u(1 - 2x) - \frac{1}{x} - D\Delta\lambda, \tag{12b}$$

with the optimal control $u = \frac{\lambda}{\gamma}x(1-x)$, and Neumann boundary conditions for x and λ. Thus, the canonical system (12) is now a system of partial differential equations (PDEs), where similar to the ODE case (7), the equation for the costate λ contains an anti-diffusion term. Again, the canonical system is not an initial value problem because only the initial values x_0 for the stocks are given. Also, the anti-diffusion in the costate equation would make an initial value problem for (12b) ill posed.

To solve the canonical system numerically, we use a spatial discretization of the PDE (12). If we take n spatial discretization points (in our numerics, we typically use n between 10^2 and 10^3 in 1D and $n \sim 10^4$ in 2D), then the canonical system is transferred to an ODE with $2n$ degrees of freedom. For small n (coarse meshes), the defect $d = n - \dim W_s(X_\infty)$ may depend on the discretization (on n), but for large n (fine meshes) d becomes independent of n (see Grass and Uecker, 2017, Appendix A). Using this asymptotic mesh–independence, we may also define the SPP by the criterion $d = 0$ for the PDE solutions, which would otherwise be delicate to define as $\dim W_s(X_{\text{PDE}}) = \infty$ and $\dim W_u(X_{\text{PDE}}) = \infty$ for systems of type (12).

We now can follow the same strategy as for the spatially discrete system, namely: first, we numerically compute the CSSs (*i.e.*, steady state solutions to (12)) and their instantaneous and discounted utilities; second, we check which of these CSSs have the SPP; and finally, we determine the canonical path from some initial state to a CSS with the SPP, and in case of multiple options of canonical paths compare their values.[10] Solutions to (12) are in general not unique, so that multiple CSSs may result, similar to the two-patch model, but now typically with a larger variety of CSSs. Thus, the important questions are again: which of the CSSs are optimal? And, in what sense?

To address these questions, we use the continuation and bifurcation software `pde2path` (Uecker et al., 2014) to compute the CSSs of the canonical system (12) in dependence of the mortality rate δ as our main bifurcation parameter (see Fig. 3). Subsequently, we use the algorithms of Uecker and de Witt (2019) to compute canonical paths. We characterize the CSSs by its spatially averaged profit J_{ca} and the spatially integrated values (L^1–norm) of the stocks $||x||_1$, the costates $||\lambda||_1$, and the controls $||u||_1$. For all $\delta > 0$, the system has a flat CSS; which is shown as the black branch f in Fig. 3. On this line we find two branching points, for $\delta_1 = 0.53$ and $\delta_2 = 0.76$, at which the new branches p1 and p2 of patterned CSSs emerge. At $\delta_3 = 1.1$, a secondary branch p2−1 of patterned CSSs branches off from p2. In Fig. 3(e)–(g), we show spatial profiles of sample solutions at $\delta = 1$, starting with the trivial flat case in 3(e). The CSSs in 3(f) and 3(g) from p1 and p2 are heterogeneous. Stock levels from CSSs on p1 are spatially decaying from the left to the right of the spatial domain, while stocks on p2 exhibit a spatially symmetric profile with elevated levels at the edges of the domain. We have described these CSSs modulo their spatial symmetries. Given that we model a homogeneous system, every solution to (12) has the $\mathbb{Z}_2$–symmetry $z \mapsto -z$. Moreover, given a patterned solution with wave number $k = m\pi/|\Omega|$, $m \in \mathbb{N}$ (such as p2 in Fig. 3 with $m = 2$), a spatial shift by π/k will generate another patterned solution, but henceforth we identify such states related by symmetry.

Our bifurcation analysis has shown that the canonical system (12) exhibits a flat CSS for all values of δ, and one, two or three patterned CSSs for $\delta > \delta_1$. Among these steady states, the patterned CSSs on branch p1 are optimal when mortality exceeds the threshold value $\delta > \delta_1$. These CSSs on p1 yield higher objective values than either the flat CSS on branch f or the patterned CSSs on branches p2 (which exists for $\delta > \delta_2$) and p2−1 (which exists for $\delta > \delta_3$). The more economical usage of the stock–enhancement activity on p1 is associated with a smaller (integrated) control $||u||_1$, a larger (integrated) stock abundance $||x||_1$, and a lower (integrated) shadow price $||\lambda||_1$, see Fig. 3.

Our numerical analysis further shows that for $\delta > \delta_1$, only the patterned branch p1 has the SPP, while the CSSs on the branches f and p2 are unstable (see Table 2). Now, let us assume (fixing the mortality rate at $\delta = 1$) a situation where for, say, historic reasons, the initial stocks are given by the spatially homogeneous distribution on branch f; that is, by $x(\cdot, 0) = 0.0596$, see Fig. 3(e). (Think of a naïve policy maker who, up to now, followed (the best) "flat policy".) Because the branch p1 has not only the highest discounted profit among the computed CSSs but also the SPP, it is natural to select the

---

[10] However, a numerical approximation of domains of attraction of CSS, or of the boundaries of such domains of attraction is no longer feasible, even for small n.
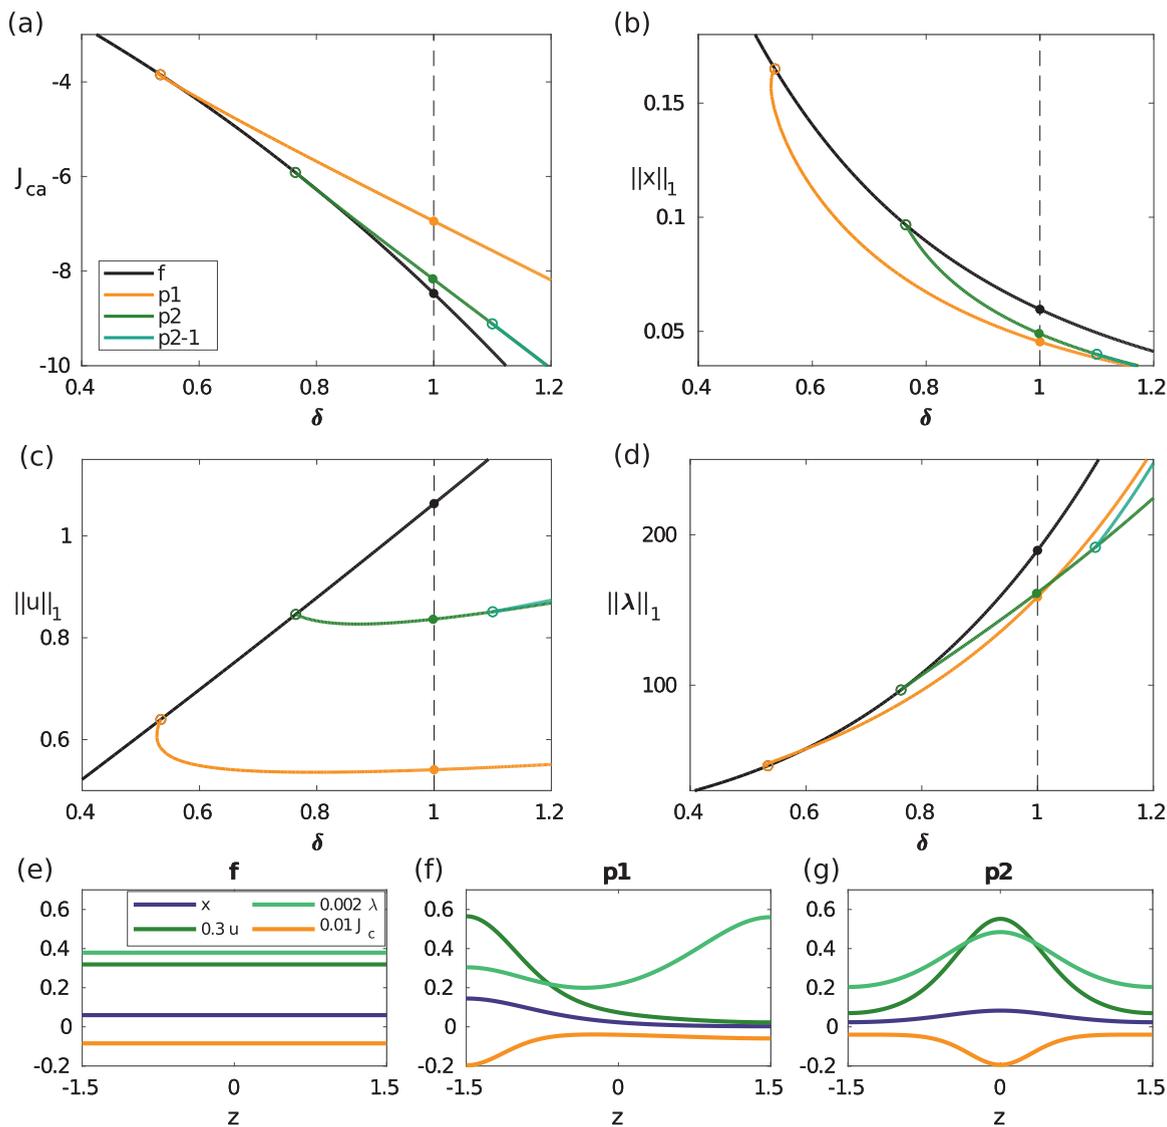
Fig. 3. Bifurcation analysis of the canonical system (12) in 1D. (a)–(d) Bifurcation diagrams of CSSs, showing (a) the spatially averaged utility J_{ca}, (b) the integrated amount of the stocks $||x||_1$, (c) the control $||u||_1$, and (d) the costates $||\lambda||_1$ in dependence of the mortality rate δ. Over the whole considered range $\delta \in [0.4, 1.2]$, the system exhibits a flat CSS (black branch f). Open circles depict branching points at which branches of patterned CSS are created: p1 at $\delta_1 = 0.53$, p2 at $\delta_2 = 0.76$, and p2-1 at $\delta_3 = 1.1$ as a secondary bifurcation from p2. (e)–(g) Spatial profiles of the CSSs on the branches f, p1, and p2 for $\delta = 1$, showing the state variable (blue), the scaled control (dark green), the scaled co-state (light green) and the scaled utility (orange). Parameter values as in (8) and spatial domain $\Omega = (-1.5, 1.5)$. (For the visibility of the colouring the reader is referred to the web version of this article.)

CSS on branch p1 in Fig. 3(f) as the target state of a desired canonical path. Our numerical analysis shows that this canonical path exists. This is illustrated in Fig. 4: Starting from the flat initial state, to reach the higher stock levels at the left end of the spatial domain, the agent invests more effort there than at the right-hand side of the domain, and the effort u is a decreasing function of the spatial coordinate z at $t = 0$, see Fig. 4(c). The discounted profit of this canonical path yields $J_{CPf1} = -277.77$, which is greater that that of the flat initial CSS $J_0 = -338.95$ and of the patterned target CSS $J_1 = -277.80$. Thus, starting in the flat CSS and governing the system towards the patterned CSS yields an even slightly higher objective value than starting directly in the patterned CSS. This can be explained by the fact that the average stock of the flat CSS is slighter higher than that of the patterned target CSS. Meanwhile, the average control of the canonical path at $t = 0$ is somewhat smaller than the average control of the target CSS. Given that the stock influences the discounted profit via the term $\log(x)$ (which increases strongly when $x < 1$) and the control via the term $-\frac{\gamma}{2}u^2$ (which does not decrease as strongly), the difference in the average stock weighs more and leads to a higher yield of the canonical path.

To check the robustness of these findings, we also conducted the same analysis for the cost parameter $\gamma = 5$ (instead of $\gamma = 10$ as before), keeping all other parameters unchanged (*i.e.* as given in (8)). In this situation, the bifurcations from the flat branch f happen at somewhat higher values ($\delta_1 = 0.811$ and $\delta_2 = 0.926$) but the qualitative structure of the bifurcation
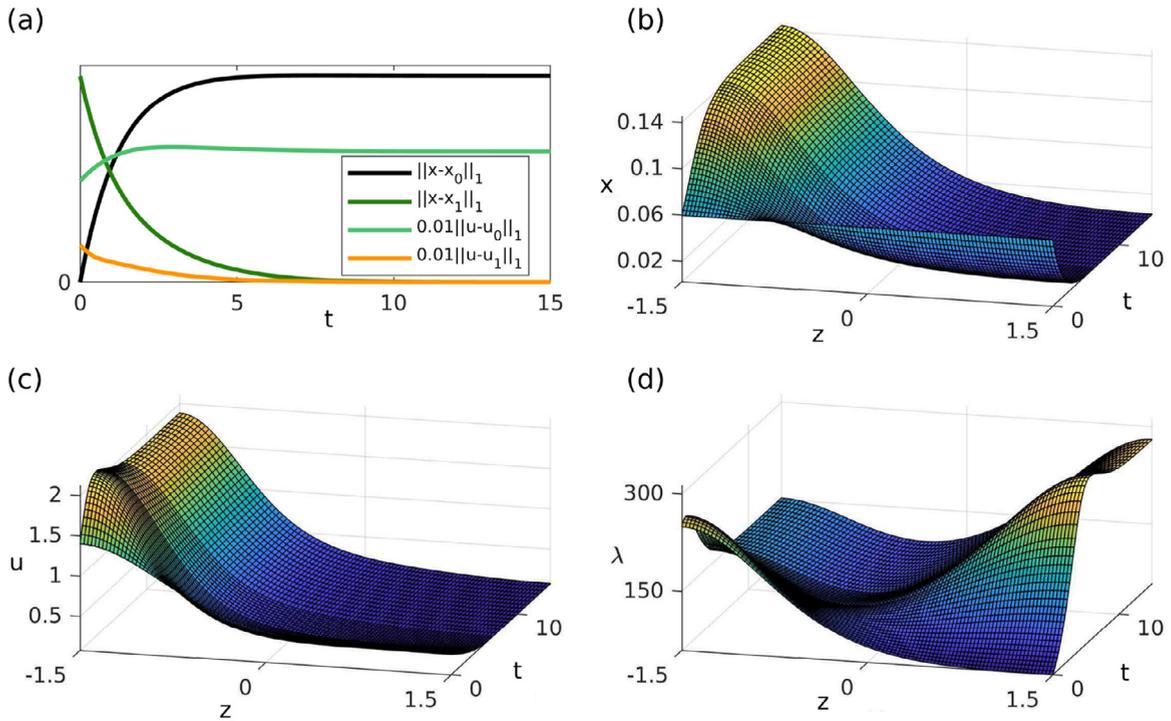
Fig. 4. Canonical path from the flat CSS on branch f to the patterned solution on branch p1 (see Fig. 3) for $\delta = 1$. (a) Temporal dependence of the distance between stocks $x(t)$, and control $u(t)$, to their starting values, x_0 and u_0, respectively, targeting values, x_1 and u_1, over time. (b)–(d) Temporal dependence of the corresponding spatial profiles of the stocks, $x(z,t)$, the control $u(z,t)$, and the costate $\lambda(z,t)$, showing the evolution towards a stationary heterogeneous solution.

diagrams, of the CSS, and of the canonical paths does not change. Moreover, we also systematically varied the discount rate ρ and the cost parameter γ, increasing the former up to $\rho = 0.1$ and decreasing the latter down to $\gamma = 1$—always finding qualitatively the same results. Similarly, in the Appendix we briefly illustrate the robustness of our results with respect to further changes of the objective function, and of the BCs in the spatial setting.

5. Stock–enhancement in a two-dimensional continuous space

Extending the 1D model of the previous section to 2D increases both the realism of the model—most natural habitats are better described as a two- than a one-dimensional space—and the numerical effort to solve it but the type of analysis remains completely analogous to the 1D case, and thus need not be repeated here. We keep the parameters (8) and choose the spatial domain $\Omega \subset \mathbb{R}^2$ to be the rectangle $\Omega = (-1.5, 1.5) \times (-1, 1)$, such that all CSS from Section 4 extend to CSS in the 2D domain by homogeneous extension in the second spatial direction. Additionally, we may expect the emergence of CSS that are heterogeneous in both directions; that is, genuine two-dimensional patterns.

As in the 1D–case, for sufficiently strong mortality rates, $\delta > \delta_1$, patterned CSSs bring about a higher profit than their homogeneous counterparts, see Fig. 5(a). In particular at $\delta = 1$, a true 2D pattern yields the highest profit of all CSSs in the 2D domain found in our numerical bifurcation analysis. For the sake of simplicity, we do not show all bifurcating branches but restrict our analysis to four different patterns, demonstrating the variety of emerging solutions. The solutions on branch p1 homogeneously extend those from p1 in 1D from Fig. 3e in the second direction, and thus also in 2D p1 branches off at $\delta_1 = 0.53$. The branch p1−1, which yields the highest profit at $\delta = 1$, arises in a secondary bifurcation from p1 at $\delta = 0.68$, where the solutions on p1 loose the SPP. As shown in Fig. 5a–b, the systems exhibit a series of other patterned branches but in the range $\delta \in [0.4, 1.2]$ they yield a lower profit than p1−1.

Of all of the branches that exist at $\delta = 1$, only the solutions on p1−1 and p3, which are both heterogeneous in both dimensions, possess the SPP and are hence possible target CSS of canonical paths. Assuming again that the initial stocks are given by the spatially homogeneous distribution on branch f, we can reach both the solution on p1−1 and the solution on p3 by a suitable choice of the canonical paths (Fig. 6). While for both canonical states, the state variable x starts with the same given values, the initial controls $u(0)$ are chosen differently to govern the dynamics to the two different solutions. This again demonstrates that we do not deal with an initial value problem and that the choice of u controls the system. The path from f to p1−1 yields a higher objective value than does the path from f to p3 ($J_{CP1} = -261.92$ vs. $J_{CP3} = -283.74$). Consequently, an agent with the aim of optimizing the objective function would decide to govern the system to the state on branch p1−1.
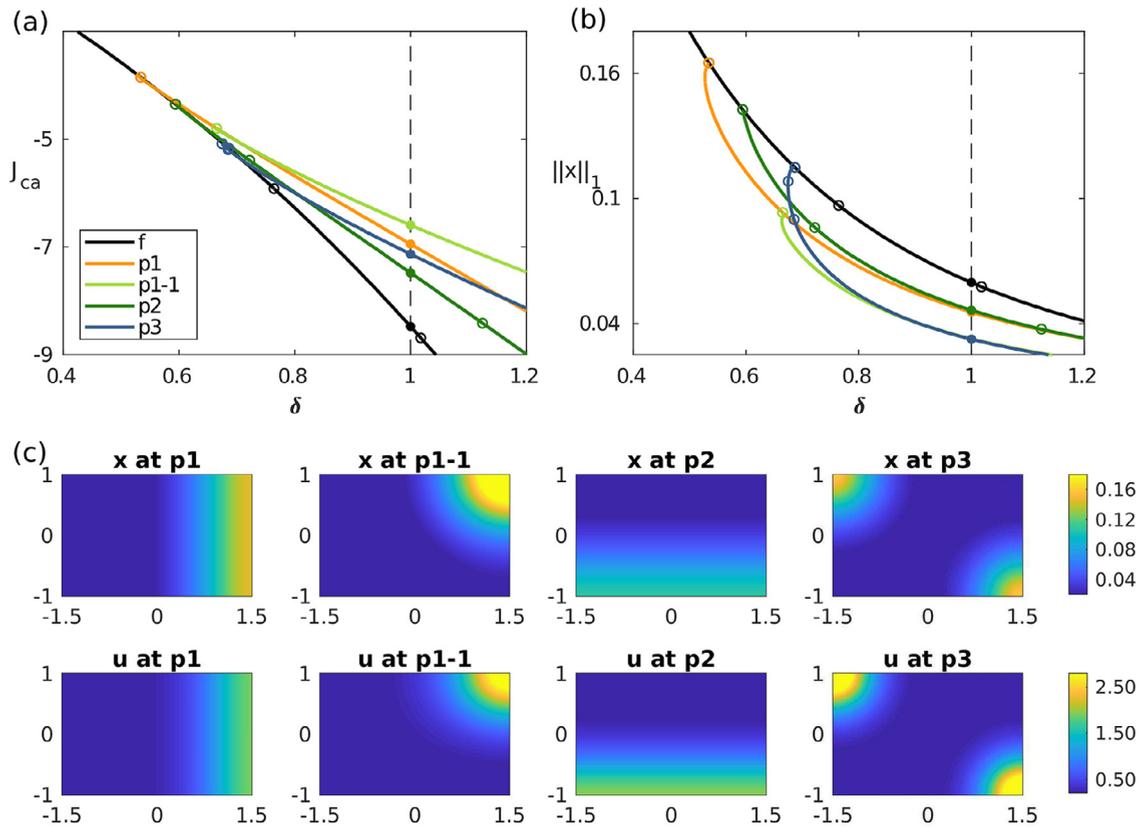
Fig. 5. Bifurcation analysis of the model in a 2D continuous space. (a)–(b) Bifurcation diagrams of CSSs, showing (a) the spatially averaged utility J_{ca} and (b) the integrated amount of the stocks $||x||_1$ in dependence of the mortality rate δ. Over the whole considered range $\delta \in [0.4, 1.2]$, the system exhibits a flat CSS (black branch f). Open circles depict branching points at which branches of patterned CSS are created (for simplicity only four patterned branches are shown): p1, p1-1 (secondary bifurcation from p1), p2, and p2-1 (secondary bifurcation from p2). (c) Spatial profiles of the solutions on the branches f, p1, and p2 for $\delta = 1$, showing the stocks, the scaled control, and the costate in colour coding. Parameter values as in (8) and spatial domain $\Omega = (-1.5, 1.5) \times (-1, 1)$.

6. Results and discussion

Today, the need to understand the effect of space for the optimal management of ecological systems is more pressing than ever. An augmented understanding could inform policy makers and stakeholder on where (and when) to best spend limited resources in their aim to conserve and to restore the biological diversity and ecosystem services. Hence, a thorough understanding of optimal spatio-temporal management of ecological systems has immediate practical consequences for the design of nature reserves, protected areas or restoration areas, for example. Modelling approaches can help in this quest, providing an important tool to illuminate the ecological and economic trade-offs involved.

Brock and Xepapadeas (2008, 2010) have shown that even in spatially homogeneous domains, it may be optimal to have spatially inhomogeneous policy measures. The interaction of economic objectives with the dynamics of spatially extended natural resources that are subject to reaction–diffusion processes can lead to the existence of a *patterned optimal steady state* (POSS), where spatially heterogeneous management strategies are economically beneficial. Subsequently, other authors extended that work and showed that POSSs may emerge in various models with two species (Brock et al., 2014; Uecker, 2016), or in models with strong externalities (Moeller and Neubert, 2013).

While many studies on the optimal control of spatially extended systems deal with optimal fisheries, harvesting and environmental externalities, in this paper we have studied a model of a one species stock population where the agent is able to increase the stock levels by means of stock–enhancement activities. Putting this model in a spatial framework, we compute optimal control policies and find that even though the model is spatially homogeneous, optimal stock–enhancement may be spatially inhomogeneous. These patterned optimal solutions emerge in both discretely coupled habitat patches and in spatially continuous one-dimensional and two-dimensional habitats when the stock mortality rate is above a critical threshold. Thus, spatially heterogeneous controls are optimal exactly in those situations in which the system is under strong stress. The computation of optimal paths from a given initial configuration towards a targeted steady state is as essential as the identification of the steady states in a first step because practical policies require transitional control measures to govern the ecological system to the targeted steady state.
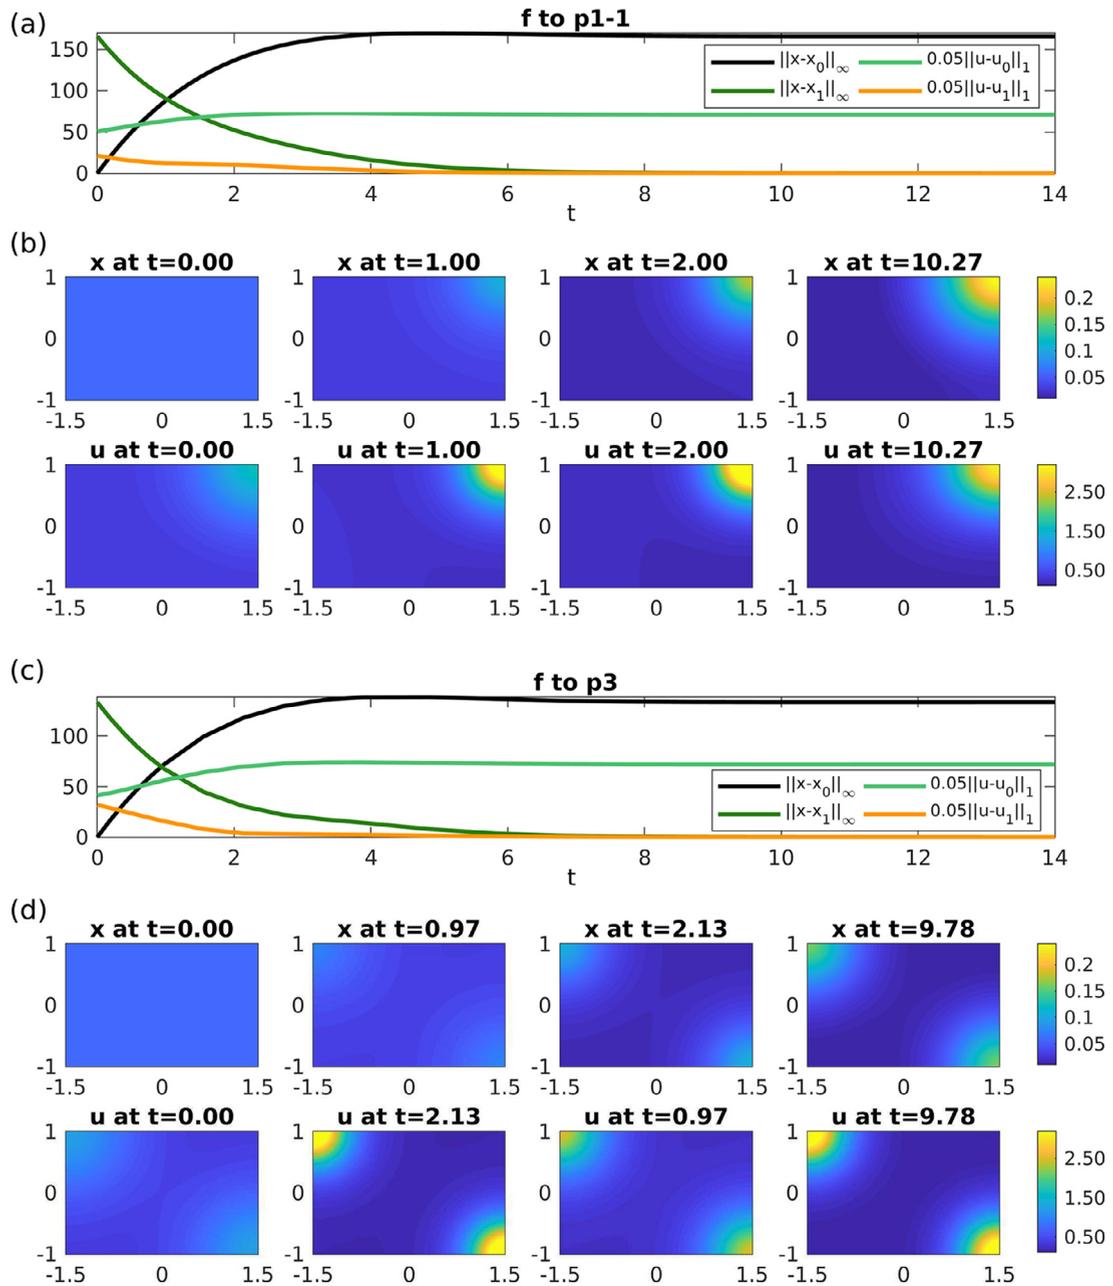
Fig. 6. (a)–(b) Canonical path from the CSS on branch f to branch p1-1 in a 2D spatial domain (see Fig. 5) for $\delta = 1$. (a) Temporal dependence of the distance between stocks $x(t)$ and controls $u(t)$ to their starting values, x_0 and u_0, respectively to their targeting values, x_1 and u_1, over time. (b) Evolution of spatial patterns along the optimal path, shown for stocks x (top row) and control u (bottom row), at selected time points. (c)–(d) Canonical path from the solution on branch f to the solution on branch p3 for $\delta = 1$.

Our major result is the presence of patterned steady states in a simple two patch and in a simple PDE stock–enhancement model, as well as the computation of the optimal paths to achieve the optimal steady states, *viz.*, the existence of POSSs. These patterned solutions are by no means a pathological theoretical finding but they may emerge for reasonable parameter settings. Specifically, increasing the realism of our model by extending the 1D model to a 2D spatial domain, brings about the emergence of steady states that are heterogeneous in both directions (*i.e.*, true two-dimensional patterns). In this case, more than one locally stable patterned solutions exist. We thus believe that this work is an important step to further understand the class of economic–ecological models exhibiting spatially patterned solutions. Even if the calculated optimal policy cannot be practically implemented because of a possible infeasibility of suitable spatio-temporally differentiated con-

trol measures, the calculated POSSs should give policy makers at least good indications about the approximate target state and the approximate transitional policy measures necessary to achieve that target.

Finally, we placed our model within the area of optimal management of renewable natural resources, but our model features great similarities with dynamic models of advertisement, market penetration and the spread of information. Namely, stock–enhancement policies appear in goodwill formation models, sales–advertising response models, models of adopters (and repeaters) of new products, models with word–of–mouth dynamics with information seeking, or in models of customer experience and information dissemination in social networks. In all these models, the agent (firm or policy maker) has an interest in enhancing the stock of informed "consumers", differentiated either by their geographic location, their position within a network or by their preferences, and where information or goodwill may spread in a diffusive way and may be controlled by advertisement, information or education. In view of these similarities, we believe that our analysis might also be valuable for research in the fields of advertisement, information economics, and political campaigns.

Acknowledgements

Appendix A. Modifications: Other utility functions, and other BCs

We briefly illustrate that our methods and results are also robust with respect to changes of the utility function J_c, or the boundary conditions for the spatially continuous case.

A1. Iso-elastic utility function

We start by generalizing J_c from $J_c(x,u) = \log(x) - \frac{\gamma}{2}u^2$ to $J_c(x,u) = e_\eta(x) - \frac{\gamma}{2}u^2$, with e_η from the isoelastic family $e_\eta(x) = \frac{1}{1-\eta}(x^{1-\eta} - 1)$, and illustrate some results for the 2-patch case.[11] As $e'_\eta(x) = x^{-\eta}$, the only change in the canonical system occurs in the costate (adjoint) equation(s); that is, we now have

$$\dot{x}_i(t) = \frac{\lambda_i}{\gamma}x_i^2(1-x_i)^2 - \delta x_i + D(x_{3-i} - x_i), \tag{A.1a}$$

$$\dot{\lambda}_i(t) = \lambda_i(\delta + \rho) - \frac{\lambda_i^2 x_i(1-2x_i)(1-x_i)}{\gamma} - \frac{1}{x_i^\eta} - D(\lambda_{3-i} - \lambda_i). \tag{A.1b}$$

Our strategy to study (Appendix A.1) remains as before: first we find (via numerical continuation and bifurcation) the CSSs, and then we compute CPs from some initial states to those CSSs with the SPP. It turns out that the behaviour of (Appendix A.1) converges to that of (7) as $\eta \searrow 1$ or $\eta \nearrow 1$, as expected. Otherwise by applying the same parameters as in Fig. 2, and using $\eta < 1$ and $\eta > 1$ with $|\eta - 1|$ being small, we have a flat CSS, $\mathbf{X}_\infty^f \approx \mathbf{X}_\infty^f(\eta)$, and two patterned CSSs, $\mathbf{X}_\infty^p(\eta)$ and $\widetilde{\mathbf{X}}_\infty^p(\eta)$, with the SPP. Fig. 7 shows the analogues of Fig. 2(a). Compared to the case $\log x$, the main effect of setting $\eta < 1$ is that the patterned CSSs move closer to the origin. For $\eta > 1$, we have the opposite effect. The qualitative features of our model remain the same, unless η substantially differs from 1 (with other parameters fixed) because we then lose the patterned CSSs.

A2. Robin boundary conditions

For the PDE case, we may also discuss the robustness with respect to the BCs. As an example, we modify the Neumann BCs (9b) and now use Robin BCs:

$$D\partial_\mathbf{n}x(z,t) + qx = g, \tag{A.2}$$

with $q, g : \partial\Omega \to \mathbb{R}$. In principle, `pde2path` also supports non-linear BCs, *e.g.*, when q, g also depend on x, u, but for simplicity we stick to (Appendix A.2). To derive the canonical system, we consider the Lagrangian

$$\mathcal{L}(x,\lambda,u) := \int_0^\infty e^{-\rho t}\left\{\frac{1}{|\Omega|}\int_\Omega J_c(x,u) - \lambda(\partial_t x - D\Delta x - f(x,u))dz\right\}dt. \tag{A.3}$$

---

[11] Similar changes occur for the 1-patch model, and for the spatially continuous case; moreover, the dependence of J_c on u can be modified similarly (including changes in the parameter γ), as already indicated at the end of Section 4.
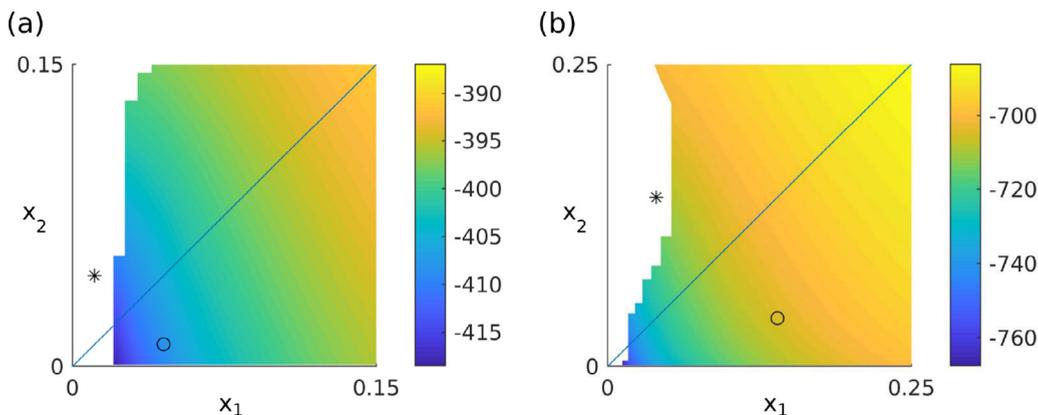
Fig. 7. Values of CPs to $\mathbf{X}_\infty^p$ (indicated by $\circ$) for (Appendix A.1) with (a) $\eta = 0.75$ (at $\delta = 0.8$) and (b) $\eta = 1.25$ (at $\delta = 1$). The respective mirrored $\widetilde{\mathbf{X}}_\infty^p$ is marked by $*$.

Integration by parts in t yields

$$-\int_0^\infty e^{-\rho t}\int_\Omega \lambda \partial_t v \ dz \ dt = \int_\Omega \lambda(z,0)x(z,0)dx + \int_0^\infty e^{-\rho t}\int_\Omega (\partial_t\lambda - \rho\lambda)x \ dz \ dt, \tag{A.4}$$

where we used the transversality condition (11), namely $\lim_{t\to\infty} e^{-\rho t}\int_\Omega \lambda(z,t)x(z,t)dz = 0$ by which the boundary terms at $t = \infty$ vanish. Next, using the boundary condition (Appendix A.2) we obtain, by integration by parts in z,

$$\int_\Omega \lambda D\Delta x \ dz = \int_\Omega (D\Delta\lambda)x \ dz + \int_{\partial\Omega} \lambda(g - qx) - x(D\partial_\mathbf{n}\lambda)dS, \tag{A.5}$$

where dS is the surface element (in 1D, $\int_{\partial\Omega} h(z)dS = h(l) - h(-l)$). Hence

$$\mathcal{L}(x,\lambda,u) = \frac{1}{|\Omega|}\int_\Omega x(z,0)\lambda(z,0)dz$$
$$+ \frac{1}{|\Omega|}\int_0^\infty e^{-\rho t}\left\{\int J_c(x,u) + (\partial_t\lambda - \rho\lambda + D\Delta\lambda)x + \lambda f(x,u)dz\right.$$
$$\left. - \int_{\partial\Omega} \lambda(g - qx) - xD\partial_\mathbf{n}\lambda dS\right\}dt. \tag{A.6}$$

The first variation of $\mathcal{L}$ with respect to x ($\mathcal{L}$ in the form (Appendix A.6)), applied to a test function $\phi \in C^\infty(Q)$ with $\phi(\cdot, 0) = 0$, yields

$$\partial_x\mathcal{L}\phi = \frac{1}{|\Omega|}\int_0^\infty e^{-\rho t}\left\{\int_\Omega\left(\partial_x J_c + \partial_t\lambda - \rho\lambda + D\Delta\lambda + \lambda\partial_x f\right)\phi dz - \int_{\partial\Omega}(\lambda q + D\partial_\mathbf{n}\lambda)\phi dS\right\}dt.$$

Therefore, by density of $C^\infty(Q)$ in $L^2(Q)$, $Q = \Omega \times [0,\infty)$, and by density of $\partial_n C^\infty(\Omega)$ in $L^2(\partial\Omega)$, the necessary first order optimality condition $\partial_x\mathcal{L}\phi = 0$ yields the costate equation $\partial_t\lambda - \rho\lambda + D\Delta\lambda + \partial_x f(x,u)\lambda + \partial_x J_c = 0$, and the boundary conditions $D\partial_\mathbf{n}\lambda + q\lambda = 0$. Similarly, the condition $\partial_\lambda\mathcal{L} = 0$ ($\mathcal{L}$ in the form (Appendix A.3)) yields the state equation. Thus, altogether we have (formally, because we assumed the existence and sufficient smoothness of optimal controls, and used (11), see also Appendix B) generalized the canonical system (12) to Robin BCs for x and λ, namely

$$D\partial_\mathbf{n}x + qx = g, \quad D\partial_\mathbf{n}\lambda + q\lambda = 0. \tag{A.7}$$

The Hamiltonian density $H = e^{-\rho t}(J_c + \lambda(f(x,u) + D\Delta x))$ is related to the Lagrangian density $L = e^{-\rho t}(J_c - \lambda(f(x,u) + D\Delta x - \partial_t x))$ as $H = L + e^{-\rho t}\lambda\partial_t x$, or after integration by parts in t, $H = L - e^{-\rho t}(-\rho\lambda + \partial_t\lambda)x$, which yields the Hamiltonian form $\partial_t\lambda = \rho\lambda - \partial_x H$ of the adjoint equation.

Naturally, for $q \neq 0$ or $g \neq 0$ we no longer have a homogeneous (or flat) CSS, so that finding a first "trivial" branch (like the homogeneous branch for $q = g = 0$), and then detecting bifurcations of non-trivial branches, becomes much more difficult, as does the classification of branches as trivial or not. We limit ourselves to the 1D case with 'mild' Robin terms, namely $q(-l) = g(-l) = 0$ (Neumann BC at the left boundary), $q(l) = 0.01$ and $g(l) = 0$. More generally, $g, q \neq 0$ can be considered as perturbations of the Neumann BCs, with $g > 0$ representing an inflow, and $q > 0$ representing a dampening of this inflow, at the right boundary. Fig. 8(a) shows a basic bifurcation diagram for this case. Using a homogeneous initial guess, we first find a solution at small δ, which has the SPP.[12] Then we again continue in δ, call the branch (which is no

---

[12] Of course, we could as well do continuation in, e.g., first q then g from the flat solution with $q = g = 0$ to $q = 0.01$, $g = 0$; this is the standard strategy, if initial solutions cannot be found otherwise. Simplify the problem by setting some terms to zero, then continuously go back to the original problem.
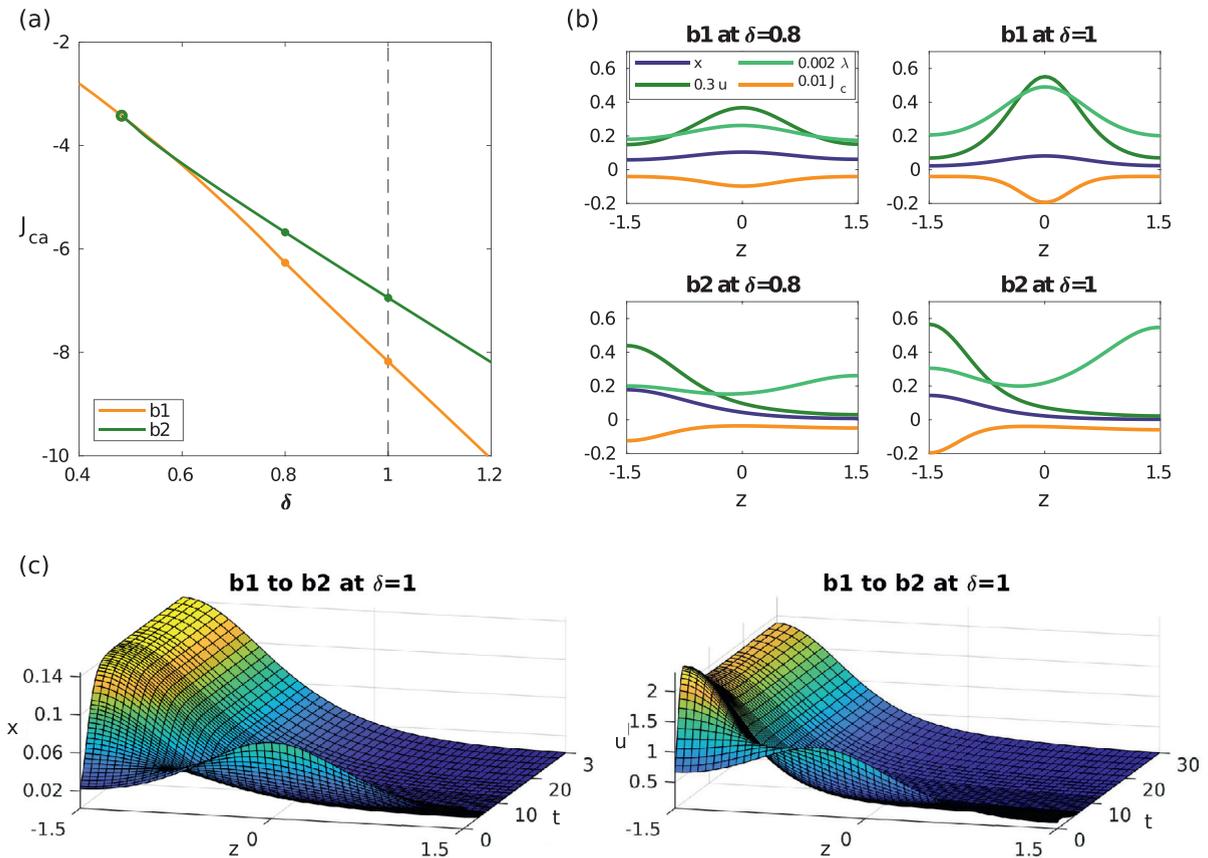
Fig. 8. (a) Bifurcation diagram and (b) sample solutions for (12) with the Robin BCs (Appendix A.7), $q = 0.01$, $g = 0$, all other parameters as in Fig. 3. (c) CP at $\delta = 1$ from b1 ($J = -327.02$) to b2 ($J = -227.75$), with $J = -279.87$ for the CP.

longer "trivial") b1 (black), and find a bifurcation point at $\delta \approx 0.483$, where the solutions on b1 lose the SPP. Moreover, the bifurcating branch b2 at higher values of δ has the SPP and larger J–values than the branch b1. The sample plots in Fig. 8(b) *inter alia* show that solutions on b1 are still "flatter" than those on b2. In Fig. 8(c) we display the CP from b1 to b2 at $\delta = 1$. and find that in this way we can increase the profit from -327.02 to -279.87. In principle, a corresponding analysis also works in 2D for small values of $|q|$, $|g|$ (and everything else kept fixed). Yet, the distinctive feature of the POSSs in a spatially homogeneous system is lost, or at least perturbed, because now all CSSs must have some spatial dependency anyway.

Appendix B. Functional analytic framework for the ODE case

In problem (1c) we choose the admissible controls from the set $\mathcal{U} = L^\infty(\mathbb{R}_+, (0, \overline{u}))$, where $\overline{u} > 0$ is some sufficiently large constant. With this choice we can apply the infinite time horizon version of Pontryagin's Maximum Principle from Tauchnitz (2015). From the modelling point of view it can be justified as follows. Very large u cannot be optimal. A larger value of u yields larger stock x, but the price is the quadratic decrease of the profit in (1b), which cannot be compensated by the logarithmic increase in x. Similarly, a very small value of u cannot be optimal because $\lim_{x \to 0} \log(x) = -\infty$. Thus, while the choice $\mathcal{U} = L^\infty(\mathbb{R}_+, (0, \overline{u}))$ formally gives control constraints, these never become active. Moreover, from the optimization (3a) we find that u is actually smooth, not just L^∞. The technical details for using the "natural" limiting intertemporal transversality condition in (2) are as follows. Let $L^2(\mathbb{R}_+, \mathbb{R}, e^{-\rho t})$ be the (exponentially) weighted L^2 space $L^2(\mathbb{R}_+, \mathbb{R}, e^{-\rho t}) = \{x \in L^2_{loc}(\mathbb{R}_+, \mathbb{R}) : \int_0^\infty x^2(t) e^{-\rho t} dt < \infty\}$. The set $\mathcal{A}_{adm}$ consists of admissible pairs $(x(\cdot), u(\cdot)) \in H^1(\mathbb{R}_+, \mathbb{R}, e^{-\rho t}) \times L^\infty(\mathbb{R}_+, U)$ such that (1a) holds and the integral in (1c) is finite, where $H^1(\mathbb{R}_+, \mathbb{R}, e^{-\rho t}) \equiv \{x, \dot{x} \in L^2(\mathbb{R}_+, \mathbb{R}, e^{-\rho t})\}$. Clearly, for any given $u \in L^\infty(\mathbb{R}_+, U)$ the solution x of (1a) is bounded and hence in $H^1(\mathbb{R}_+, \mathbb{R}, e^{-\rho t})$, and since the right hand side f of (1a) is locally Lipschitz the crucial assumptions (A1), (A2) from Tauchnitz (2015) are fulfilled.

In the two-patch model (6), choosing $\mathcal{U} = [L^\infty(\mathbb{R}_+, (0, \overline{u}))]^2$, the theory from Tauchnitz (2015) applies with very minor modifications. In the PDE model (9), we have to assume that optimal solutions exist, are sufficiently smooth, and satisfy the limiting transversality conditions (11). This is reasonable, but a genuine *assumption*, as we do not have a general existence

theory for PDE constrained infinite time horizon optimal control problems. Alternatively, we may take the point of view that for computations we will use a spatial discretization of x, u in (9), and thus may consider (9) as a symbolic notation for the high dimensional ODE constrained optimal control problem that we obtain from this spatial discretization. We may then again apply the theory from Tauchnitz (2015) to this ODE–OC problem, and derive a (high-dimensional) ODE canonical system. This ODE canonical system is the spatial discretization of the PDE canonical, which we have formally derived. Although this does not allow us to lift the ODE theory to the PDE case because the necessary estimates are not uniform in the fineness of the mesh, it at least shows the consistency of our approach. See Grass et al. (2019) for detailed discussion of related problems, including hints to the literature for the rigorous treatment of PDE–OC problems over finite time horizons, and special infinite-horizon cases where rigorous justifications of PDE canonical systems and transversality conditions can be obtained.

References

Albers, H.J., Hall, K.M., Lee, K.D., Taleghan, M.A., Dietterich, T.G., 2018. The role of restoration and key ecological invasion mechanisms in optimal spatial–dynamic management of invasive species. Ecol. Econ. 151, 44–54.

Behringer, S., Upmann, T., 2014. Optimal harvesting of a spatial renewable resource. J. Econ. Dyn. Control 42, 105–120.

Belyakov, A.O., Davydov, A.A., Veliov, V.M., 2015. Optimal cyclic exploitation of a renewable resources. J. Dyn. Control Syst. 21 (3), 475–494.

Blackwood, J., Hastings, A., Costello, C., 2010. Cost-effective management of invasive species using linear-quadratic control. Ecol. Econ. 69 (3), 519–527.

Boucekkine, R., Camacho, C., Fabbri, G., 2013. Spatial dynamics and convergence: the spatial AK model. J. Econ. Theory 148 (6), 2719–2736.

Boucekkine, R., Fabbri, G., Federico, S., Gozzi, F., 2019. Geographic environmental Kuznets curves: the optimal growth linear-quadratic case. Math. Model. Natural Phenomena 14 (1), 18.

Boucekkine, R., Fabbri, G., Federico, S., Gozzi, F., 2019. Growth and agglomeration in the heterogeneous space: a generalized AK approach. J. Econ. Geogr. 19 (6), 1287–1318.

Brock, W.A., Xepapadeas, A., 2008. Diffusion-induced instability and pattern formation in infinite horizon recursive optimal control. J. Econ. Dyn. Control 32 (9), 2745–2787.

Brock, W.A., Xepapadeas, A., 2010. Pattern formation, spatial externalities and regulation in coupled economic-ecological systems. J. Environ. Econ. Manag. 59 (2), 149–164.

Brock, W.A., Xepapadeas, A., Yannacopoulos, A.N., 2014. Optimal control in space and time and the management of environmental resources. Ann. Rev. Resour. Econ. 6, 33–68.

Desmet, K., Rossi-Hansberg, E., 2015. On the spatial economic impact of global warming. J. Urban Econ. 88, 16–37.

Ding, W., Lenhart, S., 2009. Optimal harvesting of a spatially explicit fishery model. Nat. Resour. Model. 22 (2), 173–211.

Ding, W., Lenhart, S., 2009. Optimal harvesting of a spatially explicit fishery model. Nat. Resour. Model. 22 (2), 173–211.

Grass, D., Caulkins, J.P., Feichtinger, G., Tragler, G., Behrens, D.A., 2008. Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption, and Terror. Springer, Berlin.

Grass, D., Uecker, H., 2017. Optimal management and spatial patterns in a distributed shallow lake model. Electron. J. Differ. Equ. 1, 1–21.

Grass, D., Uecker, H., Upmann, T., 2019. Optimal fishery with coastal catch. Nat. Resour. Model. 32 (4), 1–32.

Herrera, G.E., Moeller, H.V., Neubert, M.G., 2016. High-seas fish wars generate marine reserves. Proc. Natl. Acad. Sci. U.S.A. 113 (14), 3767–3772.

Kelly Jr., M.R., Neubert, M.G., Lenhart, S., 2019. Marine reserves and optimal dynamic harvesting when fishing damages habitat. Theor. Ecol. 12, 131–144.

Moberg, E.A., Shyu, E., Herrera, G.E., Lenhart, S., Lou, Y., Neubert, M.G., 2015. On the bioeconomics of marine reserves when dispersal evolves. Nat. Resour. Model. 28 (4), 456–474.

Moeller, H.V., Neubert, M.G., 2013. Habitat damage, marine reserves, and the value of spatial management. Ecol. Appl. 23 (5), 959–971.

Murray, J.D., 2003. Mathematical Biology — II: Spatial Models and Biomedical Applications. Interdisciplinary Applied Mathematics, Mathematical Biology, Vol. 18, 3rd Springer, New York.

Neubert, M.G., 2003. Marine reserves and optimal harvesting. Ecol. Lett. 6 (9), 843–849.

Neubert, M.G., Herrera, G.E., 2008. Triple benefits from spatial resource management. Theor. Ecol. 1 (1), 5–12.

Rowthorn, R.E., Laxminarayan, R., Gilligan, C.A., 2009. Optimal control of epidemics in metapopulations. J. R. Soc. Interface 6 (41), 1135–1144.

Sanchirico, J.N., Malvadkar, U., Hastings, A., Wilen, J.E., 2006. When are no-take zones an economically optimal fishery management strategy? Ecol. Appl. 16 (5), 1643–1659.

Sanchirico, J.N., Wilen, J.E., 1999. Bioeconomics of spatial exploitation in a patchy environment. J. Environ. Econ. Manag. 37 (2), 129–150.

Sanchirico, J.N., Wilen, J.E., 2001. A bioeconomic model of marine reserve creation. J. Environ. Econ. Manag. 42 (3), 257–276.

Sanchirico, J.N., Wilen, J.E., Coleman, C., 2010. Optimal rebuilding of a metapopulation. Am. J. Agric. Econ. 92 (4), 1087–1102. doi:10.1093/ajae/aaq045.

Segura, J., Hilker, F.M., Franco, D., 2019. Enhancing population stability with combined adaptive limiter control and finding the optimal harvesting–restocking balance. Theor. Popul. Biol. 130, 1–12.

Skiba, A.K., 1978. Optimal growth with a convex-concave production function. Econometrica 46 (3), 527–539.

Tauchnitz, N., 2015. The Pontryagin maximum principle for nonlinear optimal control problems with infinite horizon. J. Optim. Theory Appl. 167 (1), 27–48.

Tildesley, M.J., Savill, N.J., Shaw, D.J., Deardon, R., Brooks, S.P., Woolhouse, M.E.J., Grenfell, B.T., Keeling, M.J., 2006. Optimal reactive vaccination strategies for a foot-and-mouth outbreak in the UK. Nature 440 (7080), 83–86.

Turing, A.M., 1952. The chemical basis of morphogenesis. Philos. Trans. R. Soc. Lond. Series B Biol. Sci. 237 (641), 37–72.

Uecker, H., 2016. Optimal control and spatial patterns in a semiarid vegetation system. Nat. Resour. Model. 29 (2), 229–258.

Uecker, H., Wetzel, D., Rademacher, J.D.M., 2014. `pde2path`: A Matlab Package for Continuation and Bifurcation in 2D Elliptic Systems. In: Numerical Mathematics: Theory, Methods and Applications, 7, pp. 58–106.

Uecker, H., de Witt, H., 2019. Infinite time horizon spatially distributed optimal control problems with `pde2path`: algorithms and tutorial examples. abs/1912.11135v1.

Upmann, T., Behringer, S., 2020. Harvesting of a Renewable Resource at Discrete Patches. Theor. Ecol. 13, 459–480.