UvA-DARE (Digital Academic Repository)

Small-noise asymptotics of Hamilton-Jacobi-Bellman equations and bifurcations of stochastic optimal control problems

Grass, D.; Kiseleva, T.; Wagener, F.

DOI

[10.1016/j.cnsns.2014.09.029](https://doi.org/10.1016/j.cnsns.2014.09.029)

Publication date

2015

Document Version

Final published version

Published in

Communications in Nonlinear Science & Numerical Simulation

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Grass, D., Kiseleva, T., & Wagener, F. (2015). Small-noise asymptotics of Hamilton-Jacobi-Bellman equations and bifurcations of stochastic optimal control problems. *Communications in Nonlinear Science & Numerical Simulation*, 22(1-3), 38-54.
<https://doi.org/10.1016/j.cnsns.2014.09.029>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Small-noise asymptotics of Hamilton–Jacobi–Bellman equations and bifurcations of stochastic optimal control problems

Dieter Grass^a, Tatiana Kiseleva^{b,*}, Florian Wagener^c

^a Institute for Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

^b Department of Spatial Economics, Free University Amsterdam, Amsterdam, The Netherlands

^c Department of Economics and Econometrics, University of Amsterdam and Tinbergen Institute, Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Available online 22 October 2014

Keywords:

Small noise asymptotics
Stochastic control
Regime shifts
Bifurcations

ABSTRACT

We derive small-noise approximations of the value function of stochastic optimal control problems over an unbounded domain and use these to perform a bifurcation analysis of these problems. The corresponding zero-noise problems may feature indifference (shock, Skiba) points, that is, points of non-differentiability of the value function. Small-noise expansions are obtained in regions of regularity by a singular perturbation analysis of the stochastic Hamilton–Jacobi–Bellman equation; the expansions are matched at the boundaries of these regions to obtain an approximation over the whole state space. From this approximation, a functional geometric invariant is computed: in the presence of zero-noise indifference points, this invariant is multimodal. Regime switching thresholds of the optimally controlled dynamics are defined as those critical points where the invariant takes a local minimum. A change in the number of thresholds is a bifurcation of the dynamics. The concepts are applied to analyse the stochastic lake model.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Tipping points force economic decision makers to make up their minds. Consider the simplest situation, where the system dynamics is deterministic, the actions of the agent affect the system state and act as a system control, and there is a critical state at which the system can tip over into another regime if the agent does not change its action. In this context, typically three types of qualitatively different optimal action or ‘policy’ schedules exist. Calling the two regimes ‘red’ and ‘green’, the first type of policy always steers the system to the red regime; the second steers it to the green regime; and the third steers the system either to the green or the red, depending on the position of the initial state of the system. Which type actually obtains depends on the values of the system parameters. In the third situation, especially when time discounting is low, it is common that there are indifference states separating the basins of attraction of the two regimes: these are initial states where the optimisation problem has more than one solution. At such states the value function of the optimisation problem fails to be differentiable. Indifference states go under many names: shocks, Maxwell points, Skiba points [37], Dechert–Nishimura–Sethi–Skiba points [8,35], irregular points etc.

* Corresponding author at: De Boelelaan 1105, 1081 HV Amsterdam. Tel.: +31 205982436.

E-mail address: kiseleva.t@gmail.com (T. Kiseleva).

Tipping points and indifference states have not been studied widely by economists; as the underlying state dynamics are usually nonconcave, these systems cannot be treated by the familiar methods of convex analysis. Yet they have been found in a variety of contexts: in economic models of climate change see [30,29], growth theory see [40,35–37,8,18], illicit drug consumption [3], political economy [7,4] and others. Environmental economics is the biggest harbour for such models [2,5,6,9,19,20,25,31,38,44,39,27].

When the dynamics depends on one or more parameters, it is of interest to know the sets of parameter values for which indifference states occur in the system. A straightforward method to find these sets would be to put a grid over the parameter space, compute the optimal policy at the grid points, and check for the occurrence of an indifference state. Clearly, such an approach is computing intensive and slow. It is much faster to trace out the boundary of the parameter set for which there are indifference states; these boundaries indicate changes of the qualitative nature of the dynamics, that is, they are bifurcation curves. Along these lines, a bifurcation theory of optimal control problems has been developed recently [44,25,26].

In the present article we extend this bifurcative approach to small-noise stochastic optimal control problems. Our main contributions are that we give a natural definition of a ‘regime switching threshold’, which is the stochastic analogue of an indifference point, that we develop a natural bifurcation theory for optimally controlled stochastic systems, and that we illustrate these notions by applying them to the stochastic lake model.

Two major technical issues arise. First, in order to analyse parameter dependence of solutions, a method is needed for computing solutions quickly. For deterministic problems, initial value or boundary value methods based on the state-costate system address this issue. Based on these ‘zero-noise’ solutions, we develop ‘small-noise’ asymptotics for small stochastic perturbations, which indeed can be computed very efficiently using singular perturbation methods [32,41,42]. They obtained expansions for the value function in regions of ‘strong regularity’, that is for subregions of the state space where the value function of the corresponding unperturbed problem is smooth. These expansions are useful in situations without tipping points, where there is a single regime. But when there are several, each of them gives rise to an expansion, which then have to be matched at the boundaries [12,21,43].

The second issue is the absence of a natural bifurcation theory for stochastic dynamical systems. In the deterministic context the theory of (structural) stability and bifurcations of systems, which is based on the notion of topological equivalence, answers the question when two systems can be considered ‘qualitatively equal’. A system is stable if its equivalence class forms an open set in a suitable function topology; non-stable systems are called bifurcating.

For stochastic dynamics several equivalence notions have been proposed. One is based on the notion of ‘random dynamical systems’ of Arnold and his co-workers [1], taking as its starting point the interpretation of a stochastic process as a deterministic shift map acting on the space of realisations of the process. A random dynamical system is a deterministic skew system defined over this sequence space, for which the usual notion of equivalence can be defined. This approach has however conceptual problems [11].

A second approach, which we adopt, is based on ideas of Zeeman [46], who suggested to define equivalence of systems in terms of Morse equivalence of their ergodic invariant probability densities. While the direct application of this idea is problematic, densities not being invariant under coordinate transformations, there is a geometric invariant [45] that can be associated to a stochastic dynamical system. In special cases it reduces to the invariant probability density, and based on it a bifurcation theory can be developed [10,11].

The stochastic system spends most of its time close to the states that locally maximise the geometric invariant; as the noise strength decreases, these tends to the attracting steady states of the ‘zero-noise’ limit dynamics. In between these two maximising states, there is a local minimiser of the invariant, called by us the ‘regime switching threshold’; we introduce this as the natural analogue for stochastic systems of a deterministic indifference point.

As an illustration, we analyse the stochastic shallow lake model. Its deterministic counterpart has driven the initial development of the bifurcation theory of optimal control problems [44,25]. Dechert and O’Donnell [9] have considered a discrete time dynamic game involving stochastic shallow lake dynamics, but their noise process is rather restricted. We consider the model in continuous time and perturbed by a white noise term.

The article consists of two parts. In the first part, asymptotic expansions of the small-noise Hamilton–Jacobi–Bellman equations are derived for problems with several regimes. Section 2 is introductory and gives the framework in which we work. Asymptotic expansions within a regime are given in Section 3, while in Section 4 the matching of asymptotic expansions of neighbouring regimes is performed. In the second part of the article, these expansions are used to compute bifurcations of stochastic control problems with thresholds efficiently. In Section 5 stochastic bifurcation theory is discussed and regime switching thresholds are introduced. In Section 6 the concepts developed in this article are applied to the stochastic lake model. Section 7 concludes.

2. Formulation of the problem

Let $x(t)$ and $u(t)$ denote respectively the state of the system and the control applied to the system at time t . It is assumed that $x(t) \in X$ where $X \subset \mathbb{R}$ is an open interval and $u(t) \in D$ where $D \subset \mathbb{R}$ is a compact convex set with open interior. Let furthermore $w(t)$ denote a one-dimensional Wiener process. In the following, time arguments will generally be dropped, in order not to overburden the notation.

Let the state evolve according to the stochastic differential equation

$$dx = f(x, u)dt + \sqrt{2\varepsilon\sigma^2(x)}dw. \quad (1)$$

Here ε denotes the noise level, and the functions f and σ satisfy the following conditions:

- i f satisfies a uniform Lipschitz condition jointly in x and u ;
- ii f is linear in u : $f(x, u) = A(x) + B(x)u$, where the functions A , B are bounded on X together with their first order derivatives;
- iii σ satisfies a uniform Lipschitz condition in x ;
- iv $\sigma(x) \neq 0 \forall x \in X$.

The reader is referred to [13,28] for a detailed discussion of these conditions.

The state $x(t)$ is assumed to be completely observable at each time t . Thus the control policy may be represented as a function of the state $u = u(x)$. The control $u = u(x)$ is called *admissible* if u takes values in the compact convex set D and is locally Lipschitz.

The benefit functional associated with each u is given by

$$B[x_0, u] = \mathbb{E}_{x_0}^u \int_0^\infty g(x, u(x))e^{-\rho t} dt \quad (2)$$

where the discount rate $\rho > 0$, and where $\mathbb{E}_{x_0}^u$ is the conditional expectation operator given the initial state $x(0) = x_0$ and the control u . The integrand $g(x, u)$ is supposed to satisfy the following conditions:

- i $x \mapsto g(x, u(x))$ is bounded in any compact set for any admissible $u(x)$;
- ii $g(x, u)$ is locally Lipschitz jointly in x and u ;
- iii $\exists c < 0$ such that

$$\frac{\partial^2 g(x, u)}{\partial u^2} \leq c \quad (3)$$

for all admissible u .

We denote by $\mathcal{G}(\varepsilon)$ the problem to find an admissible control $u_\varepsilon^* = u^*(x, \varepsilon)$ that maximises the benefit functional (2) given the state dynamics Eq. (1) and the initial state $x(0) = x_0$.

Introduce the value function

$$V(x_0, \varepsilon) = \sup_u \mathbb{E}_{x_0}^u \int_0^\infty g(x, u)e^{-\rho t} dt \quad (4)$$

and the current value Hamiltonian

$$H(x, p, u) = g(x, u) + pf(x, u), \quad (5)$$

where p is a costate variable.

For each x and $p \in \mathbb{R}$ consider the *maximised current value Hamiltonian*

$$\mathcal{H}(x, p) = \max_u H(x, p, u) = \max_u [g(x, u) + pf(x, u)] = H(x, p, U(x, p)), \quad (6)$$

where

$$U(x, p) = \arg \max_u [g(x, u) + pf(x, u)]. \quad (7)$$

The function $U(x, p)$ is well defined as H is strictly concave in u because of (3).

The following theorem of Kushner [28] ensures the existence of an optimal maximising control.

Theorem 2.1. *Let the functions f , σ and g satisfy the conditions above. For every $\varepsilon > 0$ there exists an optimal control $u^*(x, \varepsilon)$ for $\mathcal{G}(\varepsilon)$. The corresponding value function $V(x, \varepsilon)$ has continuous second derivatives w.r.t. x in any compact set and solves the equation*

$$\varepsilon\sigma^2(x)V_{xx}(x, \varepsilon) + \mathcal{H}(x, V_x(x, \varepsilon)) - \rho V(x, \varepsilon) = 0. \quad (8)$$

Eq. (8) is the Hamilton–Jacobi–Bellman equation of the stochastic optimal control problem $\mathcal{G}(\varepsilon)$. If V solves this equation then the optimal control policy is given by

$$u^*(x, \varepsilon) = U(x, V_x(x, \varepsilon)). \quad (9)$$

Accordingly, to solve the stochastic problem it is sufficient to solve the second order ordinary differential Eq. (8).

We are mainly interested to solve (8) approximately for small positive values of ε , by considering the term $\varepsilon\sigma^2 V_{xx}$ as a perturbation of the equation $\mathcal{H}(x, V_x) - \rho V = 0$. As the perturbation term contains the highest order derivative of V occurring

in (8), the order of the unperturbed ($\varepsilon = 0$) and perturbed ($\varepsilon > 0$) equations are different, and we are facing a singular perturbation problem [21,43].

3. Single regimes

The present section develops a small-noise approximation of the Hamilton–Jacobi–Bellman Eq. (8) when the value function of the unperturbed equation is twice continuously differentiable. This corresponds to the ‘single regime’ situation mentioned in the introduction.

Fleming and Souganidis [14] constructed an asymptotic expansion of solutions of (8) in the situation that optimal trajectories of the unperturbed problem leave the state space in finite time. This is different from our context, where all such trajectories tend to a steady state. In contrast to their work, our approach in this section and the next will be purely formal, as we establish the algorithm how to compute the coefficients of the asymptotic expansion under the assumption that such an expansion exists, leaving the justification of the latter to future work. The differential equations for the coefficients derived in the present section parallel those found by Fleming and Souganidis; but whereas they determine the constants of integration from a Dirichlet boundary condition, in our approach these constants follow from a consideration of the optimal dynamics of the unperturbed problem at steady state.

Consider therefore the situation that the value function $v_x = V(x, 0)$ of the deterministic problem is C^2 and satisfies the dynamic programming equation

$$\mathcal{H}(x, v_x(x)) - \rho v(x) = 0 \quad (10)$$

for all $x \in X$. Note that this is an implicit first order differential equation for which we do not have any boundary conditions. Instead, solutions are selected using the following procedure. Let $v(x)$ solve (10); assuming it to be the value function of the optimal control problem, the state dynamics under optimal control satisfies

$$\dot{x} = \frac{\partial \mathcal{H}}{\partial p}(x, v_x(x)). \quad (11)$$

The resulting state trajectory $x = x(t)$ has to be defined for all t ; moreover, it has to satisfy the necessary transversality condition, which in the simplest ‘interior’ case reads as

$$\lim_{t \rightarrow \infty} e^{-\rho t} v_x(x(t)) = 0. \quad (12)$$

An extensive discussion of the necessity of the transversality condition has been given by Kamihigashi [24]. It turns out that this condition, together with the requirement that v be defined for all x , picks out the unique singular solution of (10); cf. Fig. 1.

The *shape optimal vector field* determines the state dynamics under optimal control; it is given as

$$\dot{x} = f^o(x) = \frac{\partial \mathcal{H}}{\partial p}(x, v_x(x)). \quad (13)$$

Fig. 1. The value function as the unique singular solution to the dynamic programming equation $\mathcal{H}(x, v_x) - \rho v = 0$. Dashed lines represent integral curves that do not define functions over the whole state space; thin drawn lines fail to give rise to state trajectories defined for all time or satisfying the transversality condition. The thickly drawn line represents the value function, which is a singular solution of the equation.

Assume that $\dot{x} = f^0(x)$ has a unique globally attracting equilibrium $x = x^*$; assume moreover that for all $k > 0$ there exists an asymptotic expansion

$$V(x, \varepsilon) = \sum_{j=0}^k \varepsilon^j v_j(x) + o(\varepsilon^k) = v_0(x) + \varepsilon v_1(x) + \dots + \varepsilon^k v_k(x) + o(\varepsilon^k), \quad (14)$$

uniformly on compact sets, with all $v_j, j = 0, \dots, k$, being C^2 functions. Substitution of this expression into (8), collecting terms with the same orders of ε , results in a recurrent series of equations for the $v_k(x)$. The first three of these are

$$\mathcal{H}(x, v'_0(x)) = \rho v_0(x), \quad (15a)$$

$$\mathcal{H}_p(x, v'_0(x)) v'_1(x) = \rho v_1(x) - \sigma^2(x) v''_0(x), \quad (15b)$$

$$\mathcal{H}_p(x, v'_0(x)) v'_2(x) = \rho v_2(x) - \sigma^2(x) v''_1(x) - \frac{1}{2} \mathcal{H}_{pp}(x, v'_0(x)) (v'_1(x))^2. \quad (15c)$$

Eq. (15a) is the dynamic programming equation of the unperturbed problem $\mathcal{G}(0)$; as already mentioned, it is a nonlinear implicit differential equation in $v_0(x) = V(x, 0)$. Eq. (15b) for v_1 has a singularity at $x = x^*$, since the coefficient multiplying the highest derivative is zero. As a consequence, it has only a single continuous solution that is defined on the whole state space. By requiring v_1 to be differentiable, the solution to the equation is determined. The Eq. (15b) and all subsequent equations are first order linear implicit differential equations of v_j with coefficients depending on $v_{j-1}, v_{j-2}, \dots, v_0$, determining v_j recursively.

It is convenient to replace (15a) by the linear implicit equation obtained by differentiating with respect to x and by introducing the variable $p_0(x) = v'_0(x)$. The result is the differential equation

$$\mathcal{H}_x(x, p_0(x)) + \mathcal{H}_p(x, p_0(x)) p'_0(x) = \rho p_0(x), \quad (16)$$

which is singular along the curve $\mathcal{H}_p(x, p) = 0$.

To study the solutions, we make the remark that around a non-singular point, the graph $p = p_0(x)$ of a solution of (16) is traced out by the curve $\gamma(t) = (\xi(t), \eta(t))$, whose component functions are solutions of the state-costate system

$$\dot{\xi} = \mathcal{H}_p(\xi, \eta), \quad \dot{\eta} = \rho \eta - \mathcal{H}_x(\xi, \eta). \quad (17)$$

This follows from differentiating the relation $\eta(t) = p_0(\xi(t))$ with respect to t and using the relations (17).

In terms of p_0 , the optimal vector field reads as $f^0(x) = \mathcal{H}_p(x, p_0(x))$. That the point $x = x^*$ is an attracting equilibrium of f^0 implies that close to x^* the curve γ has to lie on the stable manifold of the equilibrium $(x^*, p_0(x^*))$ of (17). This determines the curve locally and, after integrating (17), also globally.

Consider next Eq. (15b), which can be written as

$$f^0(x) v'_1(x) = \rho v_1(x) - \sigma^2(x) v''_0(x). \quad (18)$$

As v_1 is differentiable at x^* , it follows that necessarily $v_1(x^*) = \sigma^2(x^*) v''_0(x^*) / \rho$. The limit condition

$$\lim_{x \rightarrow x^*} v_1(x) = \sigma^2(x^*) v''_0(x^*) / \rho \quad (19)$$

then determines the constant of integration in Eq. (18). Proceeding in the same way the functions $v_2(x), \dots, v_k(x)$ are obtained iteratively.

Note that in [23] the value function $V(x, \varepsilon)$ is approximated by its Taylor expansions with respect to ε and x near a deterministic steady state. Summation of the coefficients corresponding to the same order of ε , say k , is then a Taylor approximation of the function $v_k(x)$ in a neighbourhood of the deterministic steady state x^* .

4. Multiple regimes

If there are several regimes, that is, if the optimal vector field f^0 of the deterministic problem has multiple equilibria which are divided by an indifference point, then Eq. (15a) has no classical solution, but it has a unique viscosity solution v_0 that is non-differentiable at the indifference point. Every equilibrium of f^0 gives rise, as in Section 3, to a separate asymptotic expansion of the form (14) of the solution to the Hamilton–Jacobi–Bellman Eq. (8). These expansions are usually defined only on neighbourhoods of the corresponding equilibria, which contain the basin of attraction of the defining equilibrium. In order to get a global expansions they have to be ‘glued’ together, using the theory of matched asymptotic expansions [12,21,43].

For definiteness, we assume that the optimal dynamics f^0 has two attracting equilibria that are separated by a single indifference point $\hat{x}$. In the terminology of the theory of matched asymptotic expansions, the interiors of the basins of attraction of the attracting equilibria are called *outer layers* (these correspond to the ‘regions of strong regularity’ in [14]). On these subdomains the function $v_0(x)$ is differentiable and $V(x, \varepsilon)$ and $V_x(x, \varepsilon)$ converge to $v_0(x)$ uniformly on compact subsets as $\varepsilon \rightarrow 0$.

The *inner layer* is an interval containing $\hat{x}$ whose length tends to 0 as $\varepsilon \rightarrow 0$. As v'_0 has a jump discontinuity at $\hat{x}$, and as $V(x, \varepsilon)$ converges uniformly to $v_0(x)$ in the outer layers, the smooth function $V_x(x, \varepsilon)$ needs to change fast in the inner layer. Consequently in the inner layer (8) is not well approximated by (15a) as $V_{xx}(x, \varepsilon)$ is large and the term $\varepsilon \sigma^2(x) V_{xx}$ in the stochastic Hamilton–Jacobi–Bellman Eq. (8) fails to be negligible. However, as the inner layer is small, the x variable is almost constant. A ‘local variable’ ξ is introduced to reflect this by setting $x = \hat{x} + \varepsilon \xi$ and expanding (8) with respect to ε .

This results in an explicit second order equation for the nondeterministic correction term. The constants of integration that determine the solution are obtained by matching the inner and outer expansions on the intersection of their domains of definition, the ‘transition layers’. Finally the inner and outer expansions are combined to yield a global asymptotic expansion of (8).

4.1. Step 1: outer layer

Denote the attracting equilibria of f^0 by x_L and x_R , and let $\hat{x}$ be the indifference point separating them. Then two asymptotic expansions of $V(x, \varepsilon)$, each associated to one of the equilibria, are obtained using the method described in Section 2.

Let the k ’th order outer expansion of $V(x, \varepsilon)$ that is defined by the equilibrium x_i be written as

$$V(x, \varepsilon) = V^{\text{out},k}(x, \varepsilon) + o(\varepsilon^k) = \begin{cases} V^{\text{out},L,k}(x, \varepsilon) + o(\varepsilon^k) & \text{if } x < \hat{x}; \\ V^{\text{out},R,k}(x, \varepsilon) + o(\varepsilon^k) & \text{if } x > \hat{x}, \end{cases} \quad (20)$$

where

$$V^{\text{out},i,k}(x, \varepsilon) = \sum_{j=0}^k \varepsilon^j v_j^i(x) = v_0^i(x) + \varepsilon v_1^i(x) + \dots + \varepsilon^k v_k^i(x) \quad (21)$$

with $i = L$ if $x < \hat{x}$ and $i = R$ if $x > \hat{x}$. The indifference point $x = \hat{x}$ satisfies

$$v_0^L(\hat{x}) = v_0^R(\hat{x}). \quad (22)$$

Set $\hat{v} = v_0^L(\hat{x}) = v_0^R(\hat{x})$. The function $V^{\text{out},k}(x, \varepsilon)$ can generally not be extended to a twice differentiable function; it is usually even discontinuous at $x = \hat{x}$.

Proceeding for each region as in the previous section, we obtain the following family of equations for the v_j^i :

$$\mathcal{H}(x, (v_0^i)'(x)) = \rho v_0^i(x), \quad (23a)$$

$$\mathcal{H}_p(x, (v_0^i)'(x))(v_1^i)'(x) = \rho v_1^i(x) - \sigma^2(x)(v_0^i)''(x), \quad (23b)$$

$$\mathcal{H}_p(x, (v_0^i)'(x))(v_2^i)'(x) = \rho v_2^i(x) - \sigma^2(x)(v_1^i)''(x) - \frac{1}{2} \mathcal{H}_{pp}(x, (v_0^i)'(x))((v_1^i)'(x))^2. \quad (23c)$$

4.2. Step 2: inner layer

To analyse (8) around $\hat{x}$, introduce the local variable

$$\xi = \frac{x - \hat{x}}{\varepsilon}. \quad (24)$$

Rewriting (8) in terms of ξ and denoting the solution in the inner layer as $V^{\text{in}}(\xi, \varepsilon)$ yields

$$\varepsilon^{-1} \sigma^2(\hat{x} + \varepsilon \xi) V_{\xi\xi}^{\text{in}} + \mathcal{H}(\hat{x} + \varepsilon \xi, \varepsilon^{-1} V_{\xi}^{\text{in}}) - \rho V^{\text{in}} = 0. \quad (25)$$

Write the k th order expansion of $V^{\text{in}}(\xi, \varepsilon) = V^{\text{in},k}(\xi, \varepsilon) + o(\varepsilon^k)$ using

$$V^{\text{in},k}(\xi) = \sum_{j=0}^k \varepsilon^j W_k(\xi) = \hat{v} + \varepsilon W_1(\xi) + \dots + \varepsilon^k W_k(\xi). \quad (26)$$

Then (25) reads as

$$\sigma^2(\hat{x}) W_1'' + \mathcal{H}(\hat{x}, W_1') - \rho \hat{v} = O(\varepsilon). \quad (27)$$

The leading terms $V^{\text{in},0}(\xi_0, \varepsilon) = W_0(\xi) = \hat{v}$ of the inner expansion and

$$V^{\text{out},0}(x, \varepsilon) = v_0^i(x) \quad (i = L \text{ if } x < \hat{x} \text{ and } i = R \text{ if } x > \hat{x}),$$

of the outer expansion coincide. Moreover, the inner expansion $V^{\text{in},1}(\xi, \varepsilon) = \hat{v} + W_1(\xi)$ approximates the solution $V(x, \varepsilon)$ in the inner layer as $\varepsilon \rightarrow 0$. Then the leading terms in the two approximations have to coincide if the limit $\varepsilon \rightarrow 0$ is taken. This gives us two conditions that are needed to determine the solution of (25) when $\varepsilon = 0$: these are derived from matching the inner and the outer expansions at the left and right transition layer respectively.

For the higher order coefficients $W_k(\xi)$, $k = 2, \dots$ of the expansion, we obtain a series of inhomogeneous linear second order differential equations whose coefficients depend on $W_1, W_2, \dots, W_{k-1}$ and their first and second order derivatives. For instance, the function W_2 satisfies

$$\sigma^2(\hat{x})W_2''(\xi) + \mathcal{H}_p(\hat{x}, W_1'(\xi))W_2'(\xi) = -(\sigma^2)'(\hat{x})W_1''(\xi)\xi - \mathcal{H}_x(\hat{x}, W_1'(\xi))\xi - \rho W_1(\xi). \quad (28)$$

4.3. Step 3: matching

Outer and inner solutions are matched by equating the inner and outer solutions in the transition layer. If the outer solutions are rewritten in terms of the inner layer coordinate ξ , they read as

$$V^{\text{out},k}(x, \varepsilon) = V^{\text{out},k}(\hat{x} + \varepsilon\xi, \varepsilon) = \sum_{m=0}^k \varepsilon^m v_m^i(\hat{x} + \varepsilon\xi) = \sum_{m=0}^k \varepsilon^m \sum_{\ell=0}^{k-m} \frac{\varepsilon^\ell}{\ell!} (v_m^i)^{(\ell)}(\hat{x}) \xi^\ell + o(\varepsilon^k) = \Gamma^{i,k}(\xi) + o(\varepsilon^k), \quad (29)$$

where $i = L$ if $\xi < 0$ and $i = R$ if $\xi > 0$, and where

$$\Gamma^{i,k}(\xi) = \sum_{j=0}^k \varepsilon^j \sum_{\ell=0}^j \frac{1}{\ell!} (v_{j-\ell}^i)^{(\ell)}(\hat{x}) \xi^\ell = \hat{v} + \varepsilon (v_1^i(\hat{x}) + (v_0^i)'(\hat{x})\xi) + \varepsilon^2 \left(v_2^i(\hat{x}) + (v_1^i)'(\hat{x})\xi + \frac{1}{2} (v_0^i)''(\hat{x})\xi^2 \right) + \dots$$

This is equated to the inner solution (26) in the transition layer, which is described by the coordinate ξ_α satisfying $x = \hat{x} + \varepsilon\xi_\alpha$. Note that $\xi = \varepsilon^{\alpha-1}\xi_\alpha$.

The values of the inner and the outer expansions in the transition layers have to be matched in the limit $\varepsilon \rightarrow 0$, therefore it is required that

$$\lim_{\varepsilon \rightarrow 0} \left(\sum_{\ell=0}^j \frac{1}{\ell!} (v_{j-\ell}^i)^{(\ell)}(\hat{x}) (\varepsilon^{\alpha-1}\xi_\alpha)^\ell - W_j(\varepsilon^{\alpha-1}\xi_\alpha) \right) = 0 \quad (30)$$

for fixed ξ_α and for $j = 1, 2, \dots, k$.

In particular, we find for W_1 the conditions

$$\lim_{\varepsilon \rightarrow 0} \left((v_0^i)'(\hat{x}) \varepsilon^{\alpha-1}\xi_\alpha + v_1^i(\hat{x}) - W_1(\varepsilon^{\alpha-1}\xi_\alpha) \right) = 0, \quad (31)$$

where $i = L$ for $\xi_\alpha < 0$ and $i = R$ for $\xi_\alpha > 0$.

For $\varepsilon = 0$ and $w_1 = W_1'$, Eq. (27) takes the form

$$\sigma^2(\hat{x})w_1' + \mathcal{H}(\hat{x}, w_1) - \rho\hat{v} = 0. \quad (32)$$

If $w_1(\tau, C_1)$ designates the general solution of (32), where $w_1(0, C_1) = C_1$, then (31) can be rewritten, according to the sign of ξ_α , as

$$W_1(0) - \int_{-\infty}^0 \left\{ w_1(\xi, C_1) - (v_0^L)'(\hat{x}) \right\} d\xi = v_1^L(\hat{x}), \quad (33)$$

$$W_1(0) + \int_0^{+\infty} \left\{ w_1(\xi, C_1) - (v_0^R)'(\hat{x}) \right\} d\xi = v_1^R(\hat{x}). \quad (34)$$

Given v_1^L and v_1^R , these equations can be solved for $W_1(0)$ and $W_1'(0) = C_1$; these values determine in turn the solution of (27) if $\varepsilon = 0$.

In general the function $w_j = W_j'$ satisfies a linear differential equation of the form

$$\sigma^2(\hat{x})w_j' + \mathcal{H}_p(\hat{x}, W_1')w_j = R_j(\xi, \hat{v}, W_1, W_1', W_1'', \dots, W_{j-1}, W_{j-1}', W_{j-1}''). \quad (35)$$

Writing the solution of this equation as $w_j = w_j(\xi, C_j)$ the limit (30) can be written as

$$W_j(0) - \int_{-\infty}^0 \left\{ w_j(\xi, C_j) - \sum_{\ell=1}^j \frac{1}{(\ell-1)!} (v_{j-\ell}^L)^{(\ell)}(\hat{x}) \xi^{\ell-1} \right\} d\xi = v_j^L(\hat{x}), \quad (36)$$

$$W_j(0) + \int_0^{+\infty} \left\{ w_j(\xi, C_j) - \sum_{\ell=1}^j \frac{1}{(\ell-1)!} (v_{j-\ell}^R)^{(\ell)}(\hat{x}) \xi^{\ell-1} \right\} d\xi = v_j^R(\hat{x}). \quad (37)$$

As before, given $v_j^L(\hat{x})$ and $v_j^R(\hat{x})$ the values $W_j(0)$ and C_j can be solved from these equations.

4.4. Step 4: composite expansion

The solution consists of two parts V^{out} and V^{in} , which now have to be combined to form a composite expansion. This is done by adding the expansions and subtracting the common part $\Gamma^{i,k}$, which in outer coordinates reads as

$$\Gamma^{i,k}\left(\frac{x-\hat{x}}{\varepsilon}\right) = \sum_{j=0}^k \varepsilon^j \sum_{\ell=0}^{k-j} \frac{1}{\ell!} \left(v_j^i\right)^{(\ell)}(\hat{x})(x-\hat{x})^\ell \quad (38)$$

Thus the k 'th order asymptotic expansion is determined by

$$V^k(x, \varepsilon) = V^{\text{out},i,k}(x) + V^{\text{in},k}\left(\frac{x-\hat{x}}{\varepsilon}\right) - \Gamma^{i,k}\left(\frac{x-\hat{x}}{\varepsilon}\right), \quad (39)$$

where $i = L$ if $x < \hat{x}$ and $i = R$ if $x > \hat{x}$.

These asymptotic expansions can be computed very efficiently, as they require only a couple of numerical integrations of ordinary differential equations.

5. Regime shifts and stochastic bifurcations

An indifference point of a deterministic optimal control problem is an initial state at which there are two optimal controls yielding different long run outcomes. The optimal vector field is multi-valued at such a point, and it belongs to two separate basins of attractions or regimes of the optimal vector field [44,25,26].

Controlled stochastic diffusion processes with nowhere vanishing variance have unique ergodic stationary distributions. This precludes defining regimes in terms of different asymptotic behaviour, and an alternative idea has to be found. One, which goes back to an idea of Zeeman [46], is to define a regime as the basin of attraction of an attractor relative to the gradient dynamics of the stationary distribution. Another would be to substitute the drift vector field for the gradient vector field in the previous definition. For one-dimensional diffusions with constant variance, that is, with the diffusion coefficient $\sigma(x)$ independent of x , both notions coincide.

When considering general diffusions with varying variance, the following difficulty appears. The stationary distributions of these diffusions are not invariant under coordinate transformations; therefore, if regimes are defined in terms of these distributions, they also would depend on the particular coordinate system. This is clearly undesirable. However, there is a functional geometrical invariant, the transformation invariant function introduced by Wagenmakers et al. [45] and denoted by I in the following, which reduces to the stationary distribution if the variance is constant. Defining regimes in terms of the gradient dynamics of I rather than the stationary distribution overcomes the difficulty [10,11].

Regimes in a one-dimensional problem are then defined as those open interval of state space that have local minima of the transformation invariant function I as their boundary points. These points will be denoted as 'regime switching thresholds'. The dynamics of a system that features regime switching thresholds is characterised by persistent state fluctuations around one of the maximisers of I until a number of large shocks forces the system to pass a threshold between two regimes; this is a noise-induced transition [22]. After that the system fluctuates around another local maximiser until the system shifts regimes again after a number of large shocks. The study of stochastic regime switching dynamics has recently received a new impetus by proposals of 'early warning indicators' of impending shifts [34].

If the optimal control problem depends on one or more system parameters, we can study qualitative changes or bifurcations as the value of these parameters change: the most important of these changes is obviously the genesis of a regime switching threshold. The remainder of this section discusses the function I , the definition of regimes and of stochastic bifurcations in more detail.

5.1. Constant diffusion processes

Let $x(t)$ be a one-dimensional stochastic process, with a constant diffusion coefficient $\sigma(x) = \sigma$, satisfying the stochastic differential equation

$$dx = f(x)dt + \sigma dw, \quad (40)$$

where w is a Wiener process. The probability distribution function $\varphi_t(x)$ satisfies the Fokker–Plank (or Kolmogorov forward) equation

$$\frac{\partial \varphi_t}{\partial t} = -\frac{\partial}{\partial x}(f\varphi_t) + \frac{1}{2} \frac{\partial^2}{\partial x^2}(\sigma^2 \varphi_t). \quad (41)$$

It is assumed that given an initial distribution φ_0 of $x(0)$, φ_t exists for all $t \geq 0$. The stationary distribution satisfies moreover $\frac{\partial \varphi_t}{\partial t} = 0$ and is given by

$$\varphi(x) = C \exp \left\{ \frac{1}{\sigma^2} \int_{x_0}^x f(\tau) d\tau \right\}, \quad (42)$$

where C is such that $\int_X \varphi(x) dx = 1$. It is easily shown that $\varphi(x)$ satisfies

$$\frac{\sigma^2}{\varphi} \frac{d\varphi}{dx} = f. \quad (43)$$

which implies that zeros of the drift function f correspond to critical points of φ , i.e. modes and antimodes of the stationary distribution. The modes of φ indicate the most frequent states of the system. A mode acts somewhat like an attractor: if the system slightly deviates from the mode, the drift f forces it to come back, as $f = 0$ and $f' < 0$ at such a point. On the contrary, a local minimum of φ , separating two modes, acts as a repeller, as $f = 0$ and $f' > 0$ at such a point. A naive definition of a regime of a one-dimensional diffusion process would then be ‘an interval of state space that is bounded by antimodes of the stationary distribution’, see Fig. 2. Equivalently, a regime can be defined as the basin of attraction of an attracting point of the (negative) gradient dynamics $\dot{x} = -\varphi'(x)$, or again as a basin of the drift dynamics $\dot{x} = f(x)$. By virtue of (43) these latter two definitions are equivalent.

5.2. Transformation invariant function

The drawback of the naive definitions given in Section 5.1 is that they depend on the particular coordinate system in which the stochastic process is described. To develop the motivation for a more general definition, consider a stochastic process $y(t)$, with positive but potentially variable diffusion coefficient $\sigma(y)$, given as

$$dy = f(y)dt + \sigma(y)dw. \quad (44)$$

We make use of the fact that any such process can be transformed into a constant diffusion process by a suitable coordinate transformation $z = r(y)$. Itô's rule implies

$$dz = \tilde{f}(z)dt + \tilde{\sigma}(z)dw \quad (45)$$

where for $y = r^{-1}(z)$ we have

$$\tilde{f}(z) = r'(y)f(y) + \frac{1}{2}\sigma^2(y)r''(y) \quad \text{and} \quad \tilde{\sigma}(z) = r'(y)\sigma(y). \quad (46)$$

Let $\psi(y)$ and $\varphi(z)$ denote the stationary probability distributions of the processes $y(t)$ and $z(t)$ respectively. Then

$$\varphi(z) = \frac{\psi(r^{-1}(z))}{r'(r^{-1}(z))}. \quad (47)$$

Eq. (47) shows that the stationary density function ψ is not invariant under coordinate transformations; more precisely, critical points of ψ are not mapped to critical points of φ . To see this in an extreme situation, choose r such that

$$r'(y) = \psi(y) \quad (48)$$

for all y . Then (47) implies that φ is constant on the interval $(0,1)$, and all points in that interval are critical. This non-invariance and the limited suitability of the invariant density as a basis of a bifurcation theory has already been pointed out by Zeeman [46].

Fig. 2. Correspondence between the deterministic drift function $f(x)$ and the probability density function $\varphi(x)$ for a process with a constant diffusion σ . Stochastic stable steady states correspond to modes of $\varphi(x)$, which are zeros of $f(x)$ where it changes its sign from positive to negative. Whereas stochastic regime switching thresholds (RST) correspond to antimodes of $\varphi(x)$, which are zeros of $f(x)$ where it changes its sign from negative to positive.

But we may transform the general process to a constant diffusion process by choosing r such that

$$r'(y) = \frac{1}{\sigma(y)} \quad \text{for all } y, \quad (49)$$

as (46) implies that $\bar{\sigma}(z) = 1$. From (47) we then infer that the invariant density of the constant diffusion process z is mapped by this transformation to the function

$$I(y) = \sigma(y)\psi(y). \quad (50)$$

In [45] the authors show that I is indeed invariant under all diffeomorphic transformations of the coordinates. This is an easy corollary of (45) and (47), which together imply that

$$\bar{\sigma}(r(y))\varphi(r(y)) = r'(y)\sigma(y)\frac{\psi(y)}{r'(y)} = \sigma(y)\psi(y). \quad (51)$$

5.3. Regime switching thresholds

The transformation invariant function I allows us to define notions like ‘stochastic attractor’, ‘regime’ and ‘regime switching threshold’ in such a way that they are independent of the coordinates used, and that they are an intrinsic property of the underlying stochastic process. While it is clear how to define these for systems with one-dimensional state spaces, a natural extension to higher dimensional systems has not yet been proposed.

In particular, we define a *stochastic attractor* of the stochastic process (44) to be a local maximiser of the invariant I , and a *regime switching threshold* to be a local minimiser of I . A subinterval of state space that is bounded by two regime switching thresholds is then called a *regime*.

Stochastic attractors generalise attractors of deterministic dynamical systems; regimes generalise basins of attraction; and regime switching thresholds generalise the boundary points of these basins.

5.4. Stochastic bifurcations

The main purpose of a bifurcation theory is to extract qualitative information about a family of objects. Especially when these objects are solutions to nonlinear equations, for which there are only very complex representations or no closed form representations at all, bifurcation theory provides a classification of these object according to qualitative characteristics. In a theoretic context like that of economic analysis this is often the kind of information which is ultimately required.

The abstract framework of bifurcation theory is that of an equivalence relation defined on a topological space. The objects are the points of the space, and equivalence of two objects is interpreted as their being qualitatively equal or having the same ‘form’ or ‘shape’. Objects whose equivalence classes are open sets are called (structurally) stable or robust: when a robust object is changed a little, that is, moved to a nearby point in the space, it remains in the same class; the form of the object does not change. Bifurcating objects are those that are not structurally stable: they trace out the boundaries of the stable equivalence classes. Clearly, a bifurcation theory is not useful if there are no structurally stable elements, or, equivalently, if there are too many inequivalent objects.

In the present context, the objects are stationary diffusion processes (44) with positive diffusion coefficient and ergodic invariant probability densities. A process is represented by its transformation invariant function I , which if drift and diffusion coefficient are both smooth, is a smooth function itself. Two processes, represented by I_1 and I_2 , are equivalent if there are diffeomorphic transformations r of state space and s of $\mathbb{R}$, such that

$$I_1 = s \circ I_2 \circ r. \quad (52)$$

The stable equivalence classes of this equivalence relation are the so-called ‘Morse functions’: a function is *Morse*, if all critical points are nondegenerate, and if all critical values are distinct [16].

5.4.1. Stochastic saddle node bifurcation

The simplest bifurcations of this theory are those which cannot be avoided in a one-parameter family $I = I_\mu$ of systems; these are accordingly called codimension one bifurcations. For families of Morse functions, there are two types of codimension one bifurcations: the stochastic saddle-node bifurcation, which occurs for $\mu = \mu_0$, if there is $x = x_0$ such that

$$I'_{\mu_0}(x_0) = 0, \quad I''_{\mu_0}(x_0) = 0, \quad I'''_{\mu_0}(x_0) \neq 0, \quad \left. \frac{\partial I'_\mu}{\partial \mu}(x_0) \right|_{\mu=\mu_0} \neq 0. \quad (53)$$

Note that these conditions imply that the negative gradient system $\dot{x} = -I'_\mu(x)$ goes through a saddle-node bifurcation at $\mu = \mu_0$, cf. [17]. As in such a bifurcation a repelling equilibrium is generated, it follows that the number of regimes changes when a stochastic saddle-node bifurcation occurs.

The second type of codimension one bifurcations of Morse functions occurs if there are two distinct critical points having the same critical value. This cannot be interpreted directly in terms of stochastic processes, and we shall disregard this

possible bifurcation in the following. When doing so, we implicitly change the definition of equivalent systems in the situation that the associated invariants are Morse functions: two of these processes are equivalent, if their associated invariants have the same number of local maximisers and local minimisers respectively.

5.4.2. Stochastic stability-exchange bifurcation

There is however one other kind of codimension one bifurcation which we consider. The escape time of a stochastic process from a given regime is the average time to reach a regime switching threshold, when starting from the mode of the regime. Regimes with higher escape times are ‘more stable’, as the process spends there, on average, more time. Moreover, escape times do not depend on the representation of the process: they are geometric invariants. A stability-exchange bifurcation is then given by the condition that the escape times of two neighbouring regimes are equal.

6. The stochastic lake model

In this section the methods developed in Sections 3 and 4 are applied to analyse the stochastic lake problem. Its deterministic counterpart, introduced in [31], is a model example of an optimal management problem with conflicting intertemporal interests and non-convex feedbacks. It has been studied extensively [44,25,20,27,19,33,9].

In the model a social planner controls usage of fertilisers by farmers; the runoff of these fields pollute a neighbouring lake. The planner tries to maximise a welfare functional which models the conflicting interests of farmers and tourists: farmers benefit from farming and hence indirectly from polluting the lake, while for instance fishermen and the tourist industry suffers.

Denote by $x(t)$ the pollution level and by $u(t)$ the inflow of pollution due to farming. We assume that there are constants $m, M > 0$ such that

$$m \leq u(t) \leq M$$

for all t , where $M \gg 1$ is an upper limit of the amount fertilisers that land can bear, and $0 < m \ll 1$ assures sustainability of the agricultural sector. Following [31], we model the pollution dynamics as

$$dx = \left(u - bx + \frac{x^2}{1+x^2} \right) dt + \sqrt{2\epsilon x^2} dw, \quad x(0) = x_0; \quad (54)$$

here $b + \sqrt{2\epsilon}dw/dt$ is the stochastic sedimentation rate of phosphorus, and the term $x^2/(1+x^2)$ is a consequence of the ecological dynamics within the lake.

The stochastic lake problem is to maximise the social benefit functional

$$B(x_0, u) = \mathbb{E}_{x_0} \int_0^\infty (\log u(t) - cx^2(t)) e^{-\rho t} dt \quad (55)$$

subject to the lake dynamics. The term $\log u(t)$ in the benefit functional (55) represents instantaneous profits derived from using fertilisers proportional to the intensity $u(t)$, whereas $cx^2(t)$ represents pollution damage. The parameter $c > 0$ models the relative costs of pollution.

The current value Hamiltonian takes the form

$$H(x, p, u) = \log u - cx^2 + p \left(u - bx + \frac{x^2}{x^2 + 1} \right).$$

If $p \in [-1/M, -1/m]$, maximising over u yields the control rule

$$U(x, p) = -\frac{1}{p}$$

and the maximised current value Hamiltonian

$$\mathcal{H}(x, p) = -\log(-p) - cx^2 - 1 + p \left(-bx + \frac{x^2}{x^2 + 1} \right). \quad (56)$$

It is easy to check that the functions in the model satisfy the conditions of Theorem 2.1 if the control u takes values in a compact interval $[m, M]$. Then Theorem 2.1 implies that the value function $V(x, \epsilon)$ solves (8), which here takes the form

$$\epsilon x^2 V_{xx} + \max_u \left[\log u - cx^2 + V_x \left(u - bx + \frac{x^2}{1+x^2} \right) \right] - \rho V = 0. \quad (57)$$

Whenever $V_x \in [-1/m, -1/M]$, the optimal policy takes the form

$$u^*(x, \epsilon) = -\frac{1}{V_x(x, \epsilon)} \quad (58)$$

and $V(x, \varepsilon)$ solves

$$\varepsilon x^2 V_{xx} - \log(-V_x) - cx^2 - 1 + V_x \left(-bx + \frac{x^2}{1+x^2} \right) - \rho V = 0. \quad (59)$$

As in Section 4, write the outer expansion of V in regime i as

$$V(x, \varepsilon) = v_0^i(x) + \varepsilon v_1^i(x) + \dots$$

Then the deterministic costate rule $p_0^i = p_0^i(x) = (v_0^i)'(x)$ satisfies Eq. (16), which takes the form

$$\left(-\frac{1}{p_0^i} - bx + \frac{x^2}{x^2+1} \right) (p_0^i)' = \rho p_0^i. \quad (60)$$

From it and the control rule, we obtain for the deterministic case the vector field in regime i under optimal management

$$\dot{x} = f^0(x) = -\frac{1}{p_0^i(x)} - bx + \frac{x^2}{x^2+1}. \quad (61)$$

The first stochastic correction v_1 of the value function satisfies (15c), which reads as

$$\left(-\frac{1}{p_0^i(x)} - bx + \frac{x^2}{x^2+1} \right) (v_1^i)' = \rho v_1^i - x^2 (p_0^i)'(x). \quad (62)$$

Throughout the section, approximate solutions to (59) are obtained by the asymptotic expansion method discussed in Sections 2 and 3. In particular, first order approximations have been used throughout; a robustness check showed that for small values of ε , the results hardly change if second order approximations are taken into account. The values $W_j(0)$ and C_j have been determined by using the MATLAB boundary value routines `bvp4c`, `bvp5c` and `bvp6c`.

6.1. Single and multiple regimes

6.1.1. Single regime

For the choices $b = 0.65$, $c = 0.7$ and $\rho = 0.03$, the dynamics

$$\dot{x} = f^0(x) = f(x, u^*(x, 0)) = u^*(x, 0) - bx + \frac{x^2}{x^2+1}$$

of the deterministic lake problem under optimal management $u = u^*(x, 0)$ has a single globally attracting equilibrium. In Fig. 3(a) are shown deterministic optimal policy, the stochastic optimal policy for $\varepsilon = 0.003$, the corresponding unimodal invariant function I and the drift

$$f_\varepsilon(x) = f(x, u(x, \varepsilon)) = u^*(x, \varepsilon) - bx + \frac{x^2}{x^2+1}$$

under optimal stochastic management $u = u^*(x, \varepsilon)$.

Fig. 3. The graph of the optimal policy function for the deterministic lake problem (solid) and the stochastic lake problem (dashed). Left: single regime; $(b, c, \rho, \varepsilon) = (0.65, 0.7, 0.03, 0.003)$; right: two regimes; $(b, c, \rho, \varepsilon) = (0.65, 0.509, 0.03, 0.003)$. Also shown are the invariant function I (blue) and the optimal drift (grey). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6.1.2. Multiple regimes

In the situation that c is lowered to $c = 0.5136$, the other parameters being held fixed, the optimally managed deterministic lake problem exhibits an indifference point. As illustrated in Fig. 3(a), at this point the optimal policy is discontinuous. The point separates two basins of attraction of the deterministic optimal vector field f^0 , see Fig. 3(a).

The stochastic optimal control and the associated invariant function are shown for $\varepsilon = 0.003$. The geometric invariant of the optimally controlled process is bimodal. The antimode, a single regime switching threshold as defined in Section 5.3, separates two regimes, to be denoted ‘clean’ and ‘turbid’.

Fig. 4(a) shows the deterministic and stochastic value functions for this situation. Note that in the presence of noise, the outcome deteriorates for small pollution levels, but improves for large levels. Fig. 4(b) shows a time series of the optimally controlled stochastic lake dynamics, illustrating the occasional occurrence of regime shifts.

6.2. Bifurcation analysis of the stochastic lake model

In this section we explore the stochastic bifurcations of the lake model as the noise intensity ε is varied, in particular the stochastic saddle-node bifurcation defined in Section 5.4.

For fixed parameter values $(b, c, \rho) = (0.65, 0.4902, 0.03)$, Fig. 5(a) shows the bifurcation diagram of the stochastic lake model with respect to the noise intensity ε . The diagram shows the location of the modes and the antimodes of the invariant $I(x)$ as a function of ε . For small values of ε the stationary distribution of the pollution is bimodal with the regime switching threshold separating the two modes. As the noise intensity increases the threshold and the ‘clean regime’ disappear in a stochastic saddle-node bifurcation around $\varepsilon_{SN} \approx 0.0048$.

Fig. 5(b) shows the diagram when c is increased to $c = 0.5304$. The figure is qualitatively different compared to the previous one. Higher pollution costs yield more restrictive policies, and as a result uniqueness of the ‘clean regime’ for large noise intensity ε .

Fig. 6(a) shows a collection of transformation invariant functions for different values of ε and fixed costs of pollution $c = 0.49$. For small ε the functions are bimodal (see Fig. 5(a)), the most of the mass is concentrated in the ‘turbid regime’, so the lake is polluted most of the time. As the noise in the system increases the mass is shifted to the ‘clean regime’. For large noise the long run distribution is unimodal with a peak at the ‘turbid’ steady state. For a better presentation of the results Figs. 5(a) and 6(a) are combined in Fig. 7.

6.3. Escape times

Recall that the escape time is by definition the expected time taken by the system to reach the regime switching threshold, i.e. to escape from the current regime, when starting from the mode of the regime. As the system spends most time in the regime with the largest escape time, this quantity is a measure of the stochastic stability of a regime.

If x_m is the mode of the current regime, $\hat{x}$ the regime switching point, and the evolution equation of the system equals

$$dy = f(y)dt + \sigma(y)dw, \quad (63)$$

Fig. 4. (a) The deterministic (solid) and the stochastic (dashed) value functions for the lake model. (b) Times series of optimally controlled stochastic lake with two regimes. The system fluctuates around one of the modes of I until kicked by a large shock to the other regime. The regimes are separated by the stochastic regime switching threshold (thickly dotted line). Parameters: $(b, c, \rho, \varepsilon) = (0.65, 0.5, 0.03, 0.003)$.

Fig. 5. Bifurcation diagrams of the long run distribution in the stochastic lake model with respect to the noise level ε . Solid lines indicate modes of the transformation invariant function $I(x)$, the dashed lines indicates the antimode of $I(x)$ – the regime switching threshold. For small noise levels the long run distribution is bimodal. As the noise parameter ε increases one of the stable states and the regime switching point merge in a saddle-node bifurcation and disappear, the long run distribution becomes unimodal. Parameters: $(b, \rho) = (0.65, 0.03)$.

Fig. 6. Transformation invariant functions and optimal control policies for different values of the noise intensity ε . Parameters: $(b, c, \rho) = (0.65, 0.49, 0.03)$.

with $\sigma(y) > 0$ for all y and w a Wiener process, then the escape time $T_{x_m \rightarrow \hat{x}}$ equals (cf. [15], p. 132)

$$T_{x_m \rightarrow \hat{x}} = \int_{x_m}^{\hat{x}} \frac{1}{\omega(y)} \int_{-\infty}^y \frac{\omega(z)}{\sigma(z)} dz dy \quad (64)$$

where

$$\omega(x) = \exp \left(\int_0^x \frac{f(\xi)}{\sigma(\xi)} d\xi \right). \quad (65)$$

Based on the approximation of value function and optimal control based on the asymptotic expansion derived in Sections 2 and 3, we can numerically compute the drift function $f(x)$ and the integrals $\omega(x)$ and $T_{x_m \rightarrow \hat{x}}$.

Fig. 8 shows the escape time from the turbid and clean regimes of the lake system with respect to the noise intensity for different values of costs of pollution c . For small pollution costs, the system is predominantly in the turbid regime; for large costs, predominantly in the clean regime. The intermediate figure shows that if the noise level increases, the curves of the escape times cross: an optimally controlled noisy shallow lake system spends relatively more time in the clean regime if the noise level is high, and relatively more in the turbid regime if the level is low.

Fig. 7. Bifurcation diagram with respect to the noise intensity ε . Parameters: $(b, c, \rho) = (0.65, 0.49, 0.03)$.

Fig. 8. Escape time (in log) from clean (solid) and turbid (dashed) regimes with respect to the noise intensity ε . Parameters: $(b, \rho) = (0.65, 0.003)$.

Fig. 9. Escape time from clean (solid) and turbid (dashed) regimes with respect to the costs of pollution (a) and the discount rate (b). The model parameters are $b = 0.56$, $\varepsilon = 10^{-3}$.

Fig. 10. Bifurcation diagram with respect to the noise intensity ε and costs of pollution c .

Fig. 9(a) shows the escape times from the turbid and clean regime respectively as a function of the cost of pollution c . The range of this parameter is chosen such that the optimal dynamics of the corresponding deterministic system exhibits multiple steady states, implying the presence of multiple regimes in the stochastic system. The figure shows that the escape time from the clean regime increases dramatically as pollution costs increase.

Figs. 9(b), depicting escape time as a function of the discount rate, show a similar pattern: small variations in the discount rate entail huge variations in the escape time from the clean regime. To see this in more detail, define the planning horizon T_ρ of a discount rate ρ by the equation $e^{-\rho T_\rho} = \frac{1}{2}$. As the discount rate in Fig. 9(b) increases from $\rho_1 = 0.025$ to $\rho_2 = 0.05$, which corresponds to a decrease from $T_{\rho_1} = 27.7$ to $T_{\rho_2} = 13.9$, the escape time from the clean regime decreases from $T_{1,\text{clean} \rightarrow \text{turbid}} = 2.4 \cdot 10^{17}$ to $T_{2,\text{clean} \rightarrow \text{turbid}} = 6.6 \cdot 10^7$, and the escape time from the turbid regime increases in a similar manner. Put differently, for small noise levels the stochastic stability of a regime depends strongly on length of the planning horizon.

Fig. 10 shows a bifurcation diagram with respect to two parameters: the pollution costs c and the noise intensity ε . Solid curves correspond to saddle-node bifurcations depicted in Fig. 5. For each value of ε there exist two critical values $\underline{c}$ and $\bar{c}$ of the pollution costs for which saddle-node bifurcations occurs. For high pollution costs $c > \bar{c}$ the clean regime is the unique regime of the lake system. For low pollution costs $c < \underline{c}$ the turbid regime is the only regime of the lake system. For intermediate values $\underline{c} < c < \bar{c}$ the lake system can be in either of the regimes. Moreover, there is a critical level of costs $\hat{c} \in (\underline{c}, \bar{c})$ that corresponds to a stability-exchange bifurcation. For $c \in (\underline{c}, \hat{c})$ the turbid regime is 'more stable', whereas for $c \in (\hat{c}, \bar{c})$ the clean regime is 'more stable'.

7. Concluding remarks

We have constructed asymptotic expansions for solutions of the small-noise Hamilton–Jacobi–Bellman equation. From these expansions, we have derived approximate solutions of stochastic control problems whose state dynamics feature different regimes and tipping points. This is motivated by the desire to perform bifurcation analysis on stochastic control problems, where many systems have to be analysed quickly for different parameter values. As our approximations can be calculated very efficiently, they are well suited for this.

We introduced the concept of regime switching threshold as a natural generalisation of the deterministic notion of indifference point. This concept is naturally embedded in the larger theory of stochastic bifurcations in the spirit of [10,11], of which the present article gives one of the first applications.

Systems with several regimes and tipping points occur in many economic applications, of which the climate pollution problem is probably the most well-known. To illustrate the generality of the methods, we presented an analysis of the stochastic shallow lake pollution management problem.

We expect to be able to lift our approach also to higher dimensional problems; there ordinary differential equations will be replaced by partial differential equations, and indifference points will be replaced by indifference surfaces, or more generally, indifference manifolds. The structure of the asymptotic expansion should however be similar.

Acknowledgements

This research was financially supported by the Austrian Science Fund (FWF) under Grant N P 23084-N13 (A MATLAB Package for Analyzing Optimal Control Problems). Part of this research was supported by Research Priority Area Behavioral Economics of the University of Amsterdam. Support by NWO under a MaGW-VIDI grant is gratefully acknowledged.

References

- [1] Arnold L. Random dynamical systems. Springer; 1998.
- [2] Brock WA, Starrett D. Managing systems with non-convex positive feedback. *Environ Resour Econ* 2003;26(4):575–602.
- [3] Bultmann R, Feichtinger G, Tragler G. Stochastic skiba sets: an example from models of illicit drug consumption. In: Lirkov I, Margenov S, Waśniewski J, editors. *Large-Scale Scientific Computing*. Springer; 2010. p. 239–46.
- [4] Caulkins JP, Hartl RF, Tragler G, Feichtinger G. Why politics makes strange bedfellows: dynamic model with DNS curves. *J Optim Theory Appl* 2001;111(2):237–54 [ISSN 0022-3239].
- [5] Crepin AS. Using fast and slow processes to manage resources with thresholds. *Environ Resour Econ* 2007;36(2):191–213 [ISSN 0924-6460].
- [6] Dasgupta P, Mäler K-G, editors. *The economics of non-convex ecosystems*. Kluwer; 2004.
- [7] Dawid H, Deissenberg Ch. On the efficiency-effects of private (dis-)trust in the government. *J Econ Behav Organiz* 2005;57(4):530–50.
- [8] Dechert WD, Nishimura K. A complete characterization of optimal growth paths in an aggregated model with a non-concave production function. *J Econ Theory* 1983;31(2):332–54 [ISSN 0022-0531].
- [9] Dechert WD, O'Donnell SI. The stochastic lake game: a numerical solution. *J Econ Dyn Control*, 0165-1889 2006;30(9–10):1569–87.
- [10] Diks CGH, Wagener FOO. A bifurcation theory for a class of discrete time Markovian stochastic systems. *Phys D* 2008;237(24):3297–306.
- [11] Diks CGH, Wagener FOO. Phenomenological and ratio bifurcations of a class of discrete time stochastic processes. *Indagationes Math* 2011;22:207–21.
- [12] Eckhaus W. Matched asymptotic expansions and singular perturbations. Number 6 in *North Holland mathematical studies*. North Holland, 1973.
- [13] Fleming WH. Stochastic control for small noise intensities. *SIAM J Control Optim* 1971;9:473–517.
- [14] Fleming WH, Souganidis PE. Asymptotic series and the method of vanishing viscosity. *Indiana Univ Math J* 1986;35(2):425–47.
- [15] Gardiner C. *Handbook of stochastic methods: for physics, chemistry & the natural sciences*. Springer series in synergetics. Berlin: Springer; 2010.
- [16] Golubitsky M, Guillemin V. *Stable mappings and their singularities*. Graduate text in mathematics, vol. 14. New York: Springer; 1973.
- [17] Guckenheimer J, Holmes P. *Nonlinear oscillations, dynamical systems and bifurcations of vector fields*. New York: Springer; 1983, ISBN 0387908196.
- [18] Haunschmied JL, Kort PM, Hartl RF, Feichtinger G. A DNS-curve in a two-state capital accumulation model: a numerical analysis. *J Econ Dyn Control* 2003;27(4):701–16 [ISSN 0165-1889].
- [19] Heijdra BJ, Heijnen P. Environmental Abatement and the Macroeconomy in the Presence of Ecological Thresholds. *Environ Res Econ* 2013;55:47–70.
- [20] Heijnen P, Wagener FOO. Avoiding an ecological regime shift is sound economic policy. *J Econ Dyn Control* 2013;37:1322–41.
- [21] Holmes MH. *Introduction to perturbation methods*. New York: Springer; 1995.
- [22] Horsthemke W, Lefever R. *Noise-induced transitions*. Springer series in synergetics. Berlin: Springer; 2006.
- [23] Judd KL. *Numerical methods in economics*. Cambridge: The MIT Press; 1998, ISBN 0262100711.
- [24] Kamihigashi T. Necessity of transversality conditions for infinite horizon problems. *Econometrica* 2001;69(4):995–1012.
- [25] Kiseleva T, Wagener FOO. Bifurcations of optimal vector fields in the shallow lake model. *J Econ Dyn Control* 2010;34(5):825–43 [ISSN 0165-1889].
- [26] Kiseleva T, Wagener FOO. Bifurcations of optimal vector fields. *Math Oper Res*. 2014. <http://dx.doi.org/10.1287/moor.2014.0655> [Epub ahead of print].
- [27] Kossioris G, Plexousakis M, Xepapadeas A, de Zeeuw A, Mäler KG. Feedback Nash equilibria for non-linear differential games in pollution control. *J Econ Dyn Control* 2008;32(4):1312–31 [ISSN 0165-1889].
- [28] Kushner HJ. Optimal discounted stochastic control for diffusion processes. *SIAM J Control Optim* 1967;5:520–31.
- [29] Lenton TM. Early warning of climate tipping points. *Nat Clim Change* 2011;1(4):201–9.
- [30] Lenton TM, Held H, Kriegler E, Hall JW, Lucht W, Rahmstorf S, Schellnhuber HJ. Tipping elements in the earth's climate system. *Proc Natl Acad Sci* 2008;105(6):1786.
- [31] Mäler KG, Xepapadeas A, de Zeeuw A. The economics of shallow lakes. *Environ Resour Econ* 2003;26:603–24.
- [32] O'Malley Jr RE. On singular perturbations, especially matching. *SIAM Rev* 1994;36(3):413–4.
- [33] Salerno G, Beard R, McDonald S. Rent seeking behavior and optimal taxation of pollution in shallow lakes. MPRA Working paper. 2007.
- [34] Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, Held H, Van Nes EH, Rietkerk M, Sugihara G. Early-warning signals for critical transitions. *Nature* 2009;461(7260):53–9.
- [35] Sethi SP. Nearest feasible paths in optimal control problems: theory, examples, and counterexamples. *J Optim Theory Appl* 1977;23(4):563–79 [ISSN 0022-3239].
- [36] Sethi SP, Thompson GL. *Optimal control theory: applications to management science and economics*. Boston: Kluwer Academic Publishers; 2000.
- [37] Skiba AK. Optimal growth with a convex–concave production function. *Econometrica* 1978;46(3):527–39 [ISSN 0012-9682].
- [38] Tahvonen O, Salo S. Nonconvexities in optimal pollution accumulation. *J Environ Econ Manage* 1996;31(2):160–77 [ISSN 0095-0696].
- [39] Tahvonen O, Withagen C. Optimality of irreversible pollution accumulation. *J Econ Dyn Control* 1996;20(9–10):1775–95.
- [40] Treadway AB. On rational entrepreneurial behaviour and the demand for investment. *Rev Econ Stud* 1969;36(2):227–39.
- [41] Van Dyke M. Nineteenth-century roots of the boundary-layer idea. *SIAM Rev* 1994;36(3):415–24.
- [42] Vasilieva AB. On the development of singular perturbation theory at Moscow state university and elsewhere. *SIAM Rev* 1994;36(3):440–52.
- [43] Verhulst F. *Methods and applications of singular perturbations: boundary layers and multiple timescale dynamics*. New York: Springer; 2005.
- [44] Wagener FOO. Skiba points and heteroclinic bifurcations, with applications to the shallow lake system. *J Econ Dyn Control* 2003;27(9):1533–61 [ISSN 0165-1889].
- [45] Wagenmakers EJ, Molenaar P, Grasman RPPP, Hartelman PAI, van der Maas HLJ. Transformation invariant stochastic catastrophe theory. *Phys D Nonlinear Phenom* 2005;211(3–4):263–76 [ISSN 0167-278].
- [46] Zeeman EC. Stability of dynamical systems. *Nonlinearity* 1988;1:115–55.