

ALGORITHMS FOR OPTIMAL CONTROL OF STOCHASTIC SWITCHING SYSTEMS*

J. HINZ[†] AND N. YAP[†]

Abstract. Optimal control problems of switching type with linear state dynamics are ubiquitous in applications of stochastic optimization. For high-dimensional problems of this type, solutions which utilize some convexity related properties are useful. For such problems, we present novel algorithmic solutions which require minimal assumptions while demonstrating remarkable computational efficiency. Furthermore, we devise procedures of the primal-dual kind to assess the distance to optimality of these approximate solutions.

Key words. optimal control, American options, pathwise stochastic control, duality

DOI. 10.1137/S0040585X97T987910

1. Introduction. When making decisions under uncertainty, the major difficulty is to determine how to update estimates and decisions in order to achieve optimality over a given time period. These kinds of questions are often framed within the realm of *Markov decision theory*, which can be viewed as discrete-time *optimal stochastic control*.

The theoretical underpinnings of Markov decision theory are now well understood. Rigorous mathematical treatments are available in textbook form (see [2], [4], [12], and [23]). However, practical applications remain persistently challenging despite the rich arsenal of theoretical tools that are currently available. In this context, *approximate dynamic programming* (see [22]) grew from attempts to provide simultaneously practically implementable heuristics and theoretical insights as to why they perform well in practice.

In order to control a large system, a practical approach to dealing with the high dimensionality of the state space is to first achieve a finite discretization of it. Alternatively, one can rely on an efficient approximation of functions on this space. In this spirit, function-based methods suggest approximating value functions on the state space. One such method is the least squares Monte Carlo approach, which suggests an approximation by a suitably parameterized set of basis functions. As these parameters are computed by performing successive regressions, this method is placed within the regression-based method family.

Following [7], [28], [29], the contribution [18] became the source of subsequent research focused on its theoretical justification. Convergence issues are addressed in [8] and later generalized in [27], [9], and [10], and extensions to multiple exercise rights were considered in [6] and studied in [3], where the connection to statistical learning theory and the theory of empirical processes is emphasized. For an overview of the applications of Monte Carlo methods in financial engineering, we refer the interested reader to Glasserman's book [13] and to the literature cited therein. Beyond financial applications, function approximation methods have also been used to capture local behavior of value functions, and advanced regression methods; e.g., kernel methods [20],

*Received by the editors December 24, 2014; revised June 15, 2015. This work was supported by Australian Research Council's Discovery Projects funding scheme, project DP13010335. Originally published in the Russian journal *Teoriya Veroyatnostei i ee Primeneniya*, 60 (2015), No. 4, pp. 770–800.

<http://www.siam.org/journals/tvp/60-4/T98791.html>

[†]University of Technology Sydney, Ultimo NSW 2007, Australia (juri.hinz@uts.edu.au).

[21], local polynomial regression [11], and neural networks [5] have been established.

One of the main advantages of regression-based methods is that they reduce computations to simple linear algebraic operations in low dimension. However, due to the successive iterations required by the implementation of the dynamic programming principle, computed solutions often exhibit instability. The thrust of this paper is to present certain numerical aspects of a novel function-based method, which utilizes some convexity assumptions to establish stable and fast solutions in terms of algebraic matrix operations. Although our approach requires us to focus on a rather specific problem structure, it covers a large number of important practical applications.

The goal of the present work is to extend and apply a concept suggested in [15] to overcome its rather restrictive assumptions. With this extension, our methodology applies to a broad problem class. Furthermore, we develop a technique to assess the quality of our numerical solutions in terms of estimation of their distance to optimality. This is realized by a computation of the so-called confidence intervals (or *bounds*) when evaluating value functions at state space variables. We present a novel and numerically efficient algorithm for estimation of the value functions from above and below.

The remainder of the paper shall be structured as follows. After introducing a general framework in section 2, section 3 presents the notion of a *convex switching system* and discusses solutions to this stochastic problem class. In section 4, we review and analyze the numerical scheme of [15] that provides fast and stable solutions to convex switching problems. Section 5 represents a first step to relaxing the requirement of convexity. A remarkable generalization is achieved in section 6 where a method has been devised that allows us to bypass any convexity requirement while yielding significant improvements in computation time. Another major contribution of the paper is presented in section 7, where we suggest an adaptation of the approach [25] to obtain recursive schemes for upper bound estimates of an approximate solution. Section 8 provides two numerical examples.

2. Markov decision theory. We begin by reviewing the classical framework of finite-horizon Markov decision theory, where we closely follow Chapter 2 of [2] and tailor it to suit our purposes. Consider a system on the finite time horizon $0, \dots, T$ whose state varies in a measurable space (E, \mathcal{E}) and is affected by elements from a set A of possible actions. For each $a \in A$, we assume that $K_t^a(x, dx')$ is a stochastic transition kernel on (E, \mathcal{E}) . Consider a fixed sequence $(X_t)_{t=0}^T$ of random variables which can be thought of as coordinate projections acting on the product $E^{\{0, \dots, T\}}$ of copies of (E, \mathcal{E}) . A mapping $\pi_t: E \mapsto A$ which describes the action that the controller of the system takes at time t is called a *decision rule*. A sequence of decision rules $\pi = (\pi_t)_{t=0}^{T-1}$ is called a *policy*. For each initial point $x_0 \in E$ and each policy $\pi = (\pi_t)_{t=0}^{T-1}$, there exists a probability measure $\mathbf{P}^{x_0, \pi}$ for which $\mathbf{P}^{x_0, \pi}(X_0 = x_0) = 1$ and where

$$(2.1) \quad \mathbf{P}^{x_0, \pi}(X_{t+1} \in B \mid X_0, \dots, X_t) = K_t^{\pi_t(X_t)}(X_t, B)$$

holds for each measurable $B \in \mathcal{E}$ and $t = 0, \dots, T-1$. That is, given that the system is in state X_t at time t , the action $a = \pi_t(X_t)$ is used to pick the transition probability $K_t^{a=\pi_t(X_t)}(X_t, \cdot)$ which assigns the random evolution of the state from X_t to X_{t+1} with the distribution $K_t^{\pi_t(X_t)}(X_t, \cdot)$. For the sake of notational convenience, we use \mathcal{K}_t^a to denote the one-step transition operator associated with the transition kernel K_t^a when the action $a \in A$ is chosen. In other words, for each action $a \in A$ the

operator \mathcal{K}_t^a acts on functions φ by

$$(2.2) \quad (\mathcal{K}_t^a \varphi)(x) = \int_E \varphi(x') K_t^a(x, dx'), \quad x \in E,$$

whenever the above integrals are well defined.

At each time t , we are given the t -step reward function $r_t: E \times A \mapsto \mathbf{R}$, where $r_t(x, a)$ represents the reward for applying an action $a \in A$ when the state of the system is $x \in E$ at time t . At the end of the time horizon, at time T , it is assumed that no action can be taken. Here, if the system is in a state x , a scrap value $r_T(x)$, which is described by a prespecified scrap function $r_T: E \rightarrow \mathbf{R}$, is collected. Given an initial point x_0 , our goal is to maximize the expected finite-horizon total reward, in other words, to find the argument $\pi^* = (\pi_t^*)_{t=0}^{T-1}$ such that

$$(2.3) \quad \pi^* = \operatorname{argmax}_{\pi \in \mathcal{A}} \mathbf{E}^{x_0, \pi} \left(\sum_{t=0}^{T-1} r_t(X_t, \pi_t(X_t)) + r_T(X_T) \right),$$

where \mathcal{A} is the set of all policies, and $\mathbf{E}^{x_0, \pi}$ denotes the expectation over the controlled Markov chain defined by (2.1). The maximization (2.3) is well defined under diverse additional assumptions (see [2, p. 199]).

The calculation of the optimal policy is addressed in the following setting. We introduce for $t = 0, \dots, T-1$ the Bellman operator

$$(2.4) \quad \mathcal{T}_t v(x) = \sup_{a \in A} (r_t(x, a) + \mathcal{K}_t^a v(x)), \quad x \in E,$$

which acts on each measurable function $v: E \rightarrow \mathbf{R}$, where the integrals $\mathcal{K}_t^a v$ for all $a \in A$ exist. Further, consider the Bellman recursion

$$(2.5) \quad v_T = r_T, \quad v_t = \mathcal{T}_t v_{t+1} \quad \text{for } t = T-1, \dots, 0.$$

Under appropriate assumptions, there exists a recursive solution $(v_t^*)_{t=0}^T$ to the Bellman recursion, which gives the so-called value functions and determines an optimal policy π^* via

$$\pi_t^*(x) = \operatorname{argmax}_{a \in A} (r_t(x, a) + \mathcal{K}_t^a v_{t+1}^*(x)), \quad x \in E,$$

for all $t = 0, \dots, T-1$.

3. Convex switching systems. For the remainder of this work, we concentrate on Markov decision problems which satisfy specific additional assumptions under which the solutions to the Bellman recursion exist. This enables us to focus on finding numerical approximations. To guarantee that these assumptions are satisfied in practical examples, a variety of sufficient conditions have been developed (see [2]).

Consider a Markov decision model whose state evolution consists of one discrete and one continuous component. To be more specific, we assume that the state space $E = P \times \mathbf{R}^d$ is the product of a finite space P and the Euclidean space \mathbf{R}^d . We suppose that the first component $p \in P$ is deterministically driven by a finite set A of actions in terms of a function

$$\alpha: P \times A \rightarrow P, \quad (p, a) \rightarrow \alpha(p, a),$$

where $\alpha(p, a) \in P$ is the new value of the discrete component of the state if its previous discrete component value was p and the action $a \in A$ was taken by the

controller. Furthermore, we assume that the continuous state component evolves as an uncontrolled Markov process $(Z_t)_{t=0}^T$ on \mathbf{R}^d whose evolution is driven by random linear transformations

$$Z_{t+1} = W_{t+1}Z_t$$

with prespecified independent and integrable disturbance matrices $(W_t)_{t=1}^T$. Finally, let us assume that the reward functions

$$r_t(p, z, a), \quad t = 0, \dots, T-1, \quad p \in P, \quad a \in A,$$

and scrap functions

$$r_T(p, z), \quad p \in P,$$

are convex and globally Lipschitz continuous in the continuous component of the state space $z \in \mathbf{R}^d$. In this setting, the transition operators are given by

$$(3.1) \quad \mathcal{K}_t^a v(p, z) = \mathbf{E}(v(\alpha(p, a), W_{t+1}z)), \quad t = 0, \dots, T-1, \quad a \in A,$$

and the Bellman operators are

$$(3.2) \quad \mathcal{T}_t v(p, z) = \sup_{a \in A} \left(r_t(p, z, a) + \mathbf{E}(v(\alpha(p, a), W_{t+1}z)) \right)$$

for all $p \in P$, $z \in \mathbf{R}^d$, and $t = 0, \dots, T-1$. Markov decision problems satisfying these assumptions are referred to as *convex switching systems* in what follows.

4. Algorithmic solutions. For such systems, the backward induction described by (2.5) solves our control problem (see [15]). However, by inspecting the Bellman operator (3.2) we see that solving the Bellman recursion results in a number of problems, the most pressing of which is that one requires a pointwise solution for each $z \in \mathbf{R}^d$. In [15], a method was presented that targeted a solution in a “functional” form. We now provide a detailed account and justification of their approach.

First, by approximating the expectation in the Bellman operator in (3.2) by finite summation, we obtain the *modified* Bellman operator \mathcal{T}_t^n that acts on a given value function according to

$$(4.1) \quad \mathcal{T}_t^n v(p, z) = \max_{a \in A} \left(r_t(p, z, a) + \sum_{k=1}^n \nu_{t+1}(k) v(\alpha(p, a), W_{t+1}(k)z) \right),$$

where $(W_{t+1}(k))_{k=1}^n$ represents appropriate realizations of disturbances with the corresponding probability weights $(\nu_{t+1}(k))_{k=1}^n$. By replacing the true Bellman operator (3.2) in the backward induction of (2.5) by its modified counterpart that is given by (4.1), we obtain the modified induction

$$(4.2) \quad v_T^n = r_T, \quad v_t^n = \mathcal{T}_t^n v_{t+1}^n \quad \text{for } t = T-1, \dots, 0.$$

Although the integration is now replaced by a finite sum, determining $(v_t^n)_{t=0}^T$ is still algorithmically intractable as the calculation must be performed at each point $z \in \mathbf{R}^d$. At this point, we turn to the important observation that since the scrap and reward functions, $r_t(p, z, a)$, $t = 0, \dots, T-1$, and $r_T(p, z)$, used in (4.1) and (4.2), are

convex in the continuous component, the resulting value functions $(v_t^n)_{t=0}^T$ must also be convex in the same component.

We now suggest an approximation of these functions $(v_t^n)_{t=0}^T$ in terms of maxima over a finite selection of their subgradients. Before we can begin to explain the advantage of such a piecewise linear approximation, we need to first establish a few concepts.

First, let us refer to a countable subset $G \subset \mathbf{R}^d$ as a *grid*. For a grid G , the *subgradient envelope* $\mathcal{S}_G f$ of a convex function f is defined to be the maximum of subgradients $\nabla_g f$ of f at each grid point $g \in G$, and so

$$\mathcal{S}_G f = \vee_{g \in G} \nabla_g f.$$

Given a family $\{(W_t(k))_{t=1}^T : k = 1, \dots, n\}$ of trajectories of disturbances that increases with $n \in \mathbf{N}$ and a family of grids $(G^m)_{m \in \mathbf{N}}$ whose tightness increases with $m \in \mathbf{N}$, we define for each $n, m \in \mathbf{N}$ the *double modified* Bellman operators $\mathcal{T}_t^{m,n}$ for $t = 0, \dots, T-1$

$$(\mathcal{T}_t^{m,n} v)(p, \cdot) = \mathcal{S}_{G^m} \max_{a \in A} \left(r_t(p, \cdot, a) + \sum_{k=1}^n \nu_{t+1}(k) v(\alpha(p, a), W_{t+1}(k) \cdot) \right).$$

Using these operators, the double modified value functions $(v_t^{m,n})_{t=0}^T$ are obtained from the backward induction which starts with

$$(4.3) \quad v_T^{m,n}(p, \cdot) = \mathcal{S}_{G^m} r_T(p, \cdot), \quad p \in P,$$

and recursively determines functions

$$(4.4) \quad v_t^{m,n} = \mathcal{T}_t^{m,n} v_{t+1}^{m,n}, \quad t = T-1, \dots, 0.$$

Obviously, this approach involves two approximation parameters $n \in \mathbf{N}$ and $m \in \mathbf{N}$, which correspond to the sampled disturbances $(W_{t+1}(k))_{k=1}^n$ with their weights $(\nu_{t+1}(k))_{k=1}^n \subset \mathbf{R}_+$ and the grid tightening $(G^m)_{m \in \mathbf{N}}$. Under appropriate assumptions, this scheme enjoys excellent convergence properties (see [15]). However, we shall now focus solely on its algorithmic aspect.

Since the double-modified backward induction (4.3) and (4.4) returns value functions $(v_t^{m,n})_{t=0}^T$ which are piecewise linear and convex (in the continuous component), we now address an appropriate representation of such functions in terms of matrices in order to rewrite the backward induction algorithm (4.3) and (4.4) in a compact matrix form.

A matrix L with d columns is called a *matrix representative* of a piecewise linear convex function $l: \mathbf{R}^d \rightarrow \mathbf{R}$ if it holds that $l(z) = \max(Lz)$ for all $z \in \mathbf{R}^d$. We shall use the expression $l \sim L$ if a piecewise linear convex function l possesses a matrix representative L .

It turns out that the formation of a subgradient envelope can be directly described in terms of matrix representatives. Namely, if l possesses a matrix representative L , then its subgradient envelope \mathcal{S}_G on the grid $G = \{g^1, \dots, g^m\}$ possesses a matrix representative $\Upsilon_G[L]$ where the row-rearrangement operator Υ_G is defined by

$$\Upsilon_G[L]_{i,\cdot} = L_{\arg \max(Lg^i)} \quad \text{for } i = 1, \dots, m.$$

In other words, when Υ_G is applied to a matrix L with d columns, the result $\Upsilon_G[L]$ of the row-rearrangement yields an $m \times d$ matrix whose i th row is the row of L at

which the maximum in Lg^i at the i th grid point is attained. As mentioned above, the relation between the subgradient envelope of a function and its matrix representative is thus given in terms of the row-rearrangement operator Υ_G :

$$l \sim L \implies \mathcal{S}_G l \sim \Upsilon_G[L].$$

A similar relation holds for the summation of piecewise linear and convex functions, followed by subgradient envelope. Namely, it corresponds to a straight summation of their matrix representatives, after row-rearrangement:

$$l \sim L, \quad f \sim F \implies \mathcal{S}_G(f + l) \sim \Upsilon_G[L] + \Upsilon_G[F].$$

Similarly, maximization of piecewise linear and convex functions, followed by subgradient envelope, is realized on matrix level by binding by rows of matrix representatives, followed by the row-rearrangement:

$$l \sim L, \quad f \sim F \implies \mathcal{S}_G(l \vee f) \sim \Upsilon_G[L \sqcup F].$$

Here, the binding-by-row operation $L \sqcup F$ performs a row concatenation of the two matrices L and F . Observe that the rows of the matrix L representing a subgradient envelope $\mathcal{S}_G f$ of a convex piecewise linear function f can always be arranged such that $L = \Upsilon_G(L)$ holds. We say that a subgradient representative L is in the *normal form* if it holds that $L = \Upsilon_G(L)$. In what follows, we shall assume that the matrix representatives are given in the normal form.

Let us also introduce an equivalent but algorithmically more convenient procedure of maximization on the level of matrix representatives for later use. Given a grid $G = \{g^1, \dots, g^m\} \subset \mathbf{R}^d$ and $m \times d$ matrices $F(a)$, $a \in A$, which are given in the normal form $F(a) = \Upsilon_G(F(a))$ for all $a \in A$, introduce

$$F := \bigsqcup_{a \in A} F(a)$$

to denote an $m \times d$ matrix F whose i th row

$$F_i = F_i(a(i)), \quad i = 1, \dots, m,$$

equals the i th row of the matrix $F(a(i))$ where the maximum at the i th grid point g^i is reached, i.e.,

$$a(i) = \operatorname{argmax}\{F_i(a)g^i : a \in A\}, \quad i = 1, \dots, m.$$

This maximization is used to obtain a subgradient envelope of the maximum over a family f^a , $a \in A$, of piecewise linear and convex functions in terms of the matrix representatives (given in the normal form) of their subgradient envelopes:

$$\mathcal{S}_G f^a \sim F(a), \quad a \in A \implies \mathcal{S}_G \left(\bigvee_{a \in A} f^a \right) \sim \bigsqcup_{a \in A} F(a).$$

Finally, we emphasize that determining the subgradient envelope of the composition of a function with a linear mapping corresponds to a simple matrix product followed by a row-rearrangement. In other words, for each $d \times d$ matrix W , it holds that

$$l \sim L \implies \mathcal{S}_G(l(W \cdot)) \sim \Upsilon_G[LW].$$

Since the double-modified backward induction involves maximization, summations, and compositions with linear mappings applied to piecewise linear convex functions, it can be rewritten in terms of matrix operations.

Precomputations. Given a grid $G^m = \{g^1, \dots, g^m\}$, implement the row-rearrangement operator $\Upsilon = \Upsilon_{G^m}$ and the row maximization operator $\sqcup_{a \in A}$. Determine a distribution sampling $(W_t(k))_{k=1}^n$ of each disturbance W_t with the corresponding weights $(\nu_t(k))_{k=1}^n$ for $t = 1, \dots, T$. Given reward functions $(r_t)_{t=0}^{T-1}$ and scrap value r_T , determine the normal form of the matrix representatives of their subgradient envelopes

$$\mathcal{S}_{G^m} r_t(p, \cdot, a) \sim R_t(p, a), \quad \mathcal{S}_{G^m} r_T(p, \cdot) \sim R_T(p)$$

for $t = 0, \dots, T-1$, $p \in P$, and $a \in A$. Introduce matrix representatives $V_t(p)$ for $t = 0, \dots, T$, $p \in P$, of each value function by

$$v_t^{n,m}(p, \cdot) \sim V_t(p) \quad \text{for } t = 0, \dots, T, \quad p \in P,$$

which are obtained via the following matrix of the backward induction.

Initialization. Start with the matrices

$$V_T(p) = R_T(p) \quad \text{for all } p \in P.$$

Recursion. For $t = T-1, \dots, 0$ calculate for $p \in P$

$$(4.5) \quad V_t(p) = \bigsqcup_{a \in A} \left(R_t(p, a) + \sum_{k=1}^n \nu_{t+1}(k) \Upsilon \left[V_{t+1}(\alpha(p, a)) W_{t+1}(k) \right] \right).$$

5. Nonconvex extension. In this section, we demonstrate that for nonconvex reward and scrap functions the above algorithm can be adapted if the functions are representable as a difference of two convex functions. More precisely, assume that for all $t = 0, \dots, T-1$, and $p \in P$, $a \in A$, it holds that

$$(5.1) \quad r_t(p, \cdot, a) = \check{r}_t(p, \cdot, a) - \hat{r}_t(p, \cdot, a)$$

and

$$(5.2) \quad r_T(p, \cdot) = \check{r}_T(p, \cdot) - \hat{r}_T(p, \cdot)$$

with convex functions $\check{r}_t(p, \cdot, a)$, $\hat{r}_t(p, \cdot, a)$, $\check{r}_T(p, \cdot)$, and $\hat{r}_T(p, \cdot)$ for $p \in P$. Given such a representation, the idea is to decompose the backward induction into parallel procedures that operate on convex functions. Suppose that at the step t , the value function v_{t+1} can be represented as a difference $v_{t+1} = \check{v}_{t+1} - \hat{v}_{t+1}$ of convex functions $\check{v}_{t+1}(p, \cdot)$ and $\hat{v}_{t+1}(p, \cdot)$ for $p \in P$. With this, we have

$$\begin{aligned} \mathcal{T}_t v(p, z) &= \sup_{a \in A} (r_t(p, z, a) + \mathcal{K}_t^a v_{t+1}(p, z)) \\ &= \sup_{a \in A} \left([\check{r}_t(p, z, a) + \mathcal{K}_t^a \check{v}_{t+1}(p, z)] - [\hat{r}_t(p, z, a) + \mathcal{K}_t^a \hat{v}_{t+1}(p, z)] \right) \end{aligned}$$

showing that before maximization in $a \in A$, the result is obtained as a difference of two convex functions. However, a direct application of convex function maximization (i.e., the use of the row maximization operator \sqcup) is not compatible with this decomposition. Therefore, we require a way to express the maximum over differences

of convex functions as a difference of two convex functions. The following simple observation helps here.

Consider for each $a \in A$ the difference $\check{f}_a - \widehat{f}_a$ of two convex functions \check{f}_a and \widehat{f}_a , and let $\widehat{f} := \sum_{a \in A} \widehat{f}_a$. Then for each $a \in A$ the functions $\check{f}_a - \widehat{f}_a + \widehat{f}$ and \widehat{f} are convex and yield the desired decomposition

$$(5.3) \quad \max_{a \in A} (\check{f}_a - \widehat{f}_a) = \max_{a \in A} (\check{f}_a - \widehat{f}_a + \widehat{f}) - \widehat{f}.$$

Having this approach in mind, we propose the following algorithm.

Precalculation. Decompose the reward $(r_t)_{t=0}^{T-1}$ and scrap r_T functions into a difference of convex functions as in (5.1) and (5.2) with their (normal form) matrix representatives

$$(5.4) \quad \begin{aligned} \mathcal{S}_{G^m} \check{r}_t(p, \cdot, a) &\sim \check{R}_t(p, a), & \mathcal{S}_{G^m} \widehat{r}_t(p, \cdot, a) &\sim \widehat{R}_t(p, a), \\ \mathcal{S}_{G^m} \check{r}_T(p, \cdot) &\sim \check{R}_T(p), & \mathcal{S}_{G^m} \widehat{r}_T(p, \cdot) &\sim \widehat{R}_T(p) \end{aligned}$$

for all $t = 0, \dots, T - 1, p \in P$, and $a \in A$. Introduce the approximate value functions $(v_t^{n,m})_{t=0}^T$ which possess the decomposition

$$(5.5) \quad v_t^{n,m} = \check{v}_t^{n,m} - \widehat{v}_t^{n,m},$$

where $\check{v}_t^{n,m}(p, \cdot)$ and $\widehat{v}_t^{n,m}(p, \cdot)$ are piecewise linear convex functions with matrix representatives

$$(5.6) \quad \check{v}_t^{n,m}(p, \cdot) \sim \check{V}_t(p) \quad \text{and} \quad \widehat{v}_t^{n,m}(p, \cdot) \sim \widehat{V}_t(p)$$

for $t = 0, \dots, T, p \in P$.

Initialization. Start with the matrices

$$\check{V}_T(p) = \check{R}_T(p) \quad \text{and} \quad \widehat{V}_T(p) = \widehat{R}_T(p) \quad \text{for all } p \in P.$$

Recursion. For $t = T - 1, \dots, 1$, calculate

$$(5.7) \quad \begin{aligned} \check{\Psi}_t(p, a) &= \check{R}_t(p, a) + \sum_{k=1}^n \nu_{t+1}(k) \Upsilon \left[\check{V}_{t+1}(\alpha(p, a)) \cdot W_{t+1}(k) \right], \\ \widehat{\Psi}_t(p, a) &= \widehat{R}_t(p, a) + \sum_{k=1}^n \nu_{t+1}(k) \Upsilon \left[\widehat{V}_{t+1}(\alpha(p, a)) \cdot W_{t+1}(k) \right], \end{aligned}$$

and determine

$$(5.8) \quad \begin{aligned} \widehat{V}_t(p) &= \sum_{a \in A} \widehat{\Psi}_t(p, a), \\ \check{V}_t(p) &= \bigsqcup_{a \in A} (\check{\Psi}_t(p, a) - \widehat{\Psi}_t(p, a) + \widehat{V}_t(p)) \end{aligned}$$

for all $p \in P$.

6. An efficient approximation. Although numerical experiments indicate stable and reliable results, it seems that the computational performance suffers from the fact that most of the calculation time is being spent on matrix rearrangements required by the operator Υ . We see from (4.5) that in order to calculate

$$(6.1) \quad \sum_{k=1}^n \nu_{t+1}(k) \Upsilon \left[V_{t+1}(\alpha(p, a)) W_{t+1}(k) \right]$$

at each step of the recursion, row-rearrangement must be performed n times, once for each disturbance matrix multiplication. This task becomes increasingly demanding for larger values of the disturbance sampling sizes n , particularly in high dimensions. Before we proceed, let us omit the time index $t + 1$ in (6.1) to ease notation. We then focus on the two major sources of computational effort in evaluation of this expression, namely,

$$(6.2) \quad \begin{array}{l} \text{the rearrangement } \Upsilon[VW(k)] \text{ of} \\ \text{large matrices } VW(k) \end{array}$$

and

$$(6.3) \quad \begin{array}{l} \text{the summation of matrices } \Upsilon[VW(k)] \text{ over} \\ \text{a large index range } k = 1, \dots, n. \end{array}$$

The remainder of this section will be divided into two parts. In section 6.1, we present a method that approximates (6.1) and addresses both problems simultaneously. The improvement in computational effort makes it feasible to obtain *approximate* solutions for large grids and distribution sample sizes. Furthermore, we will see that unlike (6.1), this approximation does not require $V = V_{t+1}$ to be convex. In section 6.2, we derive a suitable first-order approximation that provides an efficient way of evaluating functions without having to decompose them into convex components. By combining this approximation with the method in section 6.1, we obtain an efficient algorithm where we are no longer encumbered by the requirement of convexity.

6.1. Estimating the conditional expectation. The crucial point is that one can approximate the procedure in (6.2) by replacing the row-rearrangement operation with an appropriate matrix multiplication. More precisely, for $k = 1, \dots, n$ we

$$(6.4) \quad \begin{array}{l} \text{construct a matrix } Y(k) \text{ such that} \\ Y(k)VW(k) \text{ approximates } \Upsilon[VW(k)]. \end{array}$$

Before we justify the approximation (6.4), let us first see how it can be used to address the computational problem associated with (6.3). Given (6.4), we now have the following approximation to (6.1):

$$(6.5) \quad \sum_{k=1}^n \nu(k) \Upsilon[VW(k)] \approx \sum_{k=1}^n \nu(k) Y(k) VW(k),$$

and this in turn requires an efficient calculation of sums of matrices. In practical examples, the distribution sample size n and the grid size m (row number of V) will typically be orders of magnitude of the dimension d of the disturbance matrices $W(k)$. For instance, to achieve an acceptable level of numerical convergence in typical applications, the sample size n and the grid size m must be chosen in the range of several thousands, whereas the state size dimension d is typically of several dozens. This insight shows that a significant reduction in computational effort can be achieved by an additive decomposition of the disturbance realizations. Assume that disturbance matrix W is represented as the linear combination

$$(6.6) \quad W = \overline{W} + \sum_{j=1}^J \varepsilon_j E(j)$$

with nonrandom matrices \overline{W} and $(E(j))_{j=1}^J$, and random coefficients $(\varepsilon_j)_{j=1}^J$. With this decomposition, each realization $W(k)$ of the disturbance matrix W is obtained as

$$(6.7) \quad W(k) = \overline{W} + \sum_{j=1}^J \varepsilon_j(k)E(j), \quad k = 1, \dots, n.$$

Utilizing this, we obtain the following interchange of summations on the right-hand side of (6.5):

$$\begin{aligned} \sum_{k=1}^n \nu(k)Y(k)VW(k) &= \left(\sum_{k=1}^n \nu(k)Y(k) \right) V\overline{W} \\ &\quad + \sum_{j=1}^J \left(\sum_{k=1}^n \nu(k)\varepsilon_j(k)Y(k) \right) VE(j). \end{aligned}$$

If one precomputes the following matrices:

$$(6.8) \quad D_0 = \sum_{k=1}^n \nu(k)Y(k), \quad D_j = \sum_{k=1}^n \nu(k)\varepsilon_j(k)Y(k), \quad j = 1, \dots, J,$$

we then obtain a significant simplification to (6.5),

$$(6.9) \quad \sum_{k=1}^n \nu(k)Y(k)VW(k) = D_0V\overline{W} + \sum_{j=1}^J D_jVE(j),$$

which only involves a low number of matrix summations and multiplications. We shall denote this efficient calculation of the conditional expectation by \mathcal{E} , where

$$(6.10) \quad \mathcal{E}(V) := D_0V\overline{W} + \sum_{j=1}^J D_jVE(j).$$

We now address the justification of the approximation in (6.4). Suppose that the grid $\{g^1, \dots, g^m\}$ is represented by the matrix G , where each row i contains a row vector, representing the grid point g^i . Thus, G will consist of m rows with $G_{i,\cdot} = g^i$ for $i = 1, \dots, m$. Now let $\tilde{L} = \Upsilon[L]$ be the result of the application of Υ to a matrix L . The matrix \tilde{L} is then characterized by the following requirements:

$$(6.11) \quad \begin{aligned} \tilde{L} = \Upsilon[L] \text{ consists of } m \text{ rows which are obtained} \\ \text{from the rows of } L \text{ by an arrangement,} \end{aligned}$$

such that

$$(6.12) \quad \tilde{L}_{i,\cdot} G_{i,\cdot}^\top \geq L_{j,\cdot} G_{i,\cdot}^\top \quad \text{for all } i, j = 1, \dots, m.$$

According to requirement (6.11), we therefore assume that

$$(6.13) \quad \begin{aligned} \Upsilon[VW(k)] \text{ consists of } m \text{ rows which are obtained} \\ \text{from the rows of } VW(k) \text{ by row-rearrangement.} \end{aligned}$$

Since any row rearrangement can be achieved by a left-multiplication with appropriate matrix, there will always exist a permutation matrix $Y_V(k)$ such that

$$(6.14) \quad Y_V(k) V W(k) = \Upsilon[V W(k)].$$

Computing each $Y_V(k)$ requires great effort since it is dependent not only on $W(k)$, but also on each V . We suggest determining a reasonable surrogate $Y(k)$ for $Y_V(k)$ which depends only on $W(k)$ and not on V . Since $Y_V(k)$ must satisfy (6.14), we observe with (6.12) in mind that

$$(6.15) \quad (Y_V(k)V)_{i,\cdot} (W(k)G)_{i,\cdot}^\top \geq V_{j,\cdot} (W(k)G)_{i,\cdot}^\top \quad \text{for } i, j = 1, \dots, m.$$

Now, for each $i = 1, \dots, m$ consider the row $(W(k)G)_{i,\cdot}$ and determine the closest row $G_{h_k(i),\cdot}$ in the original grid matrix by

$$(6.16) \quad h_k(i) = \operatorname{argmin} \left\{ j = 1, \dots, m : \|(W(k)G)_{i,\cdot} - G_{j,\cdot}\| \right\}, \quad i = 1, \dots, m.$$

With this proximity function $h_k: \{1, \dots, m\} \rightarrow \{1, \dots, m\}$, we may consider, in place of the relation (6.15), the condition

$$(6.17) \quad (Y(k)V)_{i,\cdot} G_{h_k(i),\cdot}^\top \geq V_{j,\cdot} G_{h_k(i),\cdot}^\top \quad \text{for } i, j = 1, \dots, m,$$

with an appropriate permutation matrix $Y(k)$. While (6.17) is clearly not equivalent to (6.15), it does provide a reasonable approximation when the grid is sufficiently dense. Now define $Y(k)$ to be such that

$$(6.18) \quad Y(k)_{i,j} = \begin{cases} 1 & \text{if } j = h_k(i), \\ 0 & \text{otherwise,} \end{cases}$$

and observe that with this permutation matrix $Y(k)$, the assertion

$$V_{h_k(i),\cdot} G_{h_k(i),\cdot}^\top \geq V_{j,\cdot} G_{h_k(i),\cdot}^\top \quad \text{for } i, j = 1, \dots, m$$

holds if V is in the normal form $\Upsilon[V] = V$. That is, the required approximation (6.4) is determined by (6.18).

The precalculations involved in the approximation of (6.1) (i.e., computing D_0, \dots, D_j) are computationally demanding. Thus, a gain in computation performance can only be realized if disturbances $(W_t)_{t=1}^T$ are identically distributed whereby the precalculations need only be done once. In this case, the ideas presented in this section will be encapsulated in the following algorithm.

Precalculations. Determine a sampling $(W(k))_{k=1}^n$ from the target distribution. For each disturbance $W(k)$, find the corresponding permutation matrix $Y(k)$ as in (6.18) using the proximity function (6.16). Use these matrices and the components of the decomposition described in (6.7) of each $W(k)$ to compute the matrices (6.8).

Continuation. Execute the algorithm (5.4)–(5.8) but replace (5.7) with

$$\begin{aligned} \check{\Psi}_t(p, a) &= \check{R}_t(p, a) + \mathcal{E} \left[\check{V}_{t+1}(\alpha(p, a)) \right], \\ \widehat{\Psi}_t(p, a) &= \widehat{R}_t(p, a) + \mathcal{E} \left[\widehat{V}_{t+1}(\alpha(p, a)) \right] \end{aligned}$$

by substituting the conditional expectations with its efficient counterpart (6.10).

6.2. A direct approach. So far, we have worked with parallel procedures on convex functions. However, an important point to note is that in no part of the efficient conditional expectation procedure was the convexity of the target function required. With this in mind, we shall now present a further simplification to this algorithm based on a first-order approximation. Previously, we considered the convex decomposition $f = \hat{f} - \check{f}$ of a nonconvex function f , where the two convex piecewise linear functions \hat{f} and \check{f} with respective matrix representatives \hat{F} and \check{F} , are given in their normal form with respect to the grid G . The value $f(z)$ at point z is then calculated as

$$f(z) = \max(\hat{F}z) - \max(\check{F}z).$$

However, if only the matrix difference $\hat{F} - \check{F}$ is known, then it is possible to use a first-order approximation

$$f(z) \approx (\hat{F} - \check{F})_{h(z),z},$$

where h is the so-called host function of the underlying grid $G = \{g^1, \dots, g^m\}$ with $m \in \mathbf{N}$,

$$h(z) = \operatorname{argmin}\{\|z - g^i\| : i \in \{1, \dots, m\}\},$$

which returns to each argument $z \in \mathbf{R}^d$ the so-called host—the index of the point on the grid with the smallest distance to z . The first-order approximation uses the difference $\hat{F} - \check{F}$ directly and, unlike convex decomposition, does not require a separate calculation of convex and concave parts. If one decides to use this first-order approximation to access the functions, then there is no need to trace convex and concave parts separately. This gives a significant simplification and results in the following direct algorithm.

Precomputations. Determine the operator \mathcal{E} as in (6.10), under the assumptions required therefore. Determine for $p \in P$, $a \in A$, and $t = 0, \dots, T - 1$ the matrices

$$(6.19) \quad R_t(p, a) = \hat{R}_t(p, a) - \check{R}_t(p, a), \quad R_T(p) = \hat{R}_T(p) - \check{R}_T(p),$$

which are obtained as in (5.4). Introduce the approximate value functions, and their convex decomposition and representatives as in (5.5) and (5.6). The matrices

$$V_t(p) = \check{V}_t(p) - \hat{V}_t(p) \quad \text{for } t = 0, \dots, T, \quad p \in P$$

are obtained via the following scheme.

Initialization. Start with the matrices

$$V_T(p) = \check{R}_T(p) - \hat{R}_T(p) \quad \text{for all } p \in P.$$

Recursion. For $t = T - 1, \dots, 1$ calculate for $p \in P$

$$(6.20) \quad V_t(p) = \bigsqcup_{a \in A} \left(R_t(p, a) + \mathcal{E}(V_{t+1}(p, a)) \right).$$

Remark. Unlike in the convex decomposition case (5.4)–(5.8), the *direct* algorithm (6.19)–(6.20) merely returns the difference $V_t(p) = \check{V}_t(p) - \hat{V}_t(p)$. That is, access to the approximate value functions is provided via

$$v_t^{m,n}(p, z) \approx V_t(p)_{h(z),z}$$

using the host function h of the grid G^m . In particular, we suggest an approximation of the optimal policy $\pi^{m,n} = (\pi_t^{m,n})_{t=0}^{T-1}$ as

$$(6.21) \quad \pi_t^{m,n}(p, z) = \operatorname{argmax}_{a \in A} \left((R_t(p, a) + \mathcal{E}(V_{t+1}(p, a)))_{h(z), \cdot} \cdot z \right)$$

for $t = 0, \dots, T-1$, $z \in \mathbf{R}^d$, $p \in P$. To obtain an efficient implementation of a host function h , a tree-like structure on the grid can be used which may be established using hierarchical clustering methods.

7. Solution diagnostics. In section 6.1, we derived a heuristic method to obtain an efficient approximation to the conditional expectation in the Bellman recursion. In section 6.2, we saw that by combining this with a first-order approximation, we were then able to obtain the approximation (6.21) for a given grid to the optimal policy in (2.3). In order to address the distance to optimality of this approximate solution, we first need to outline an appropriate measure for this distance.

Suppose we are given an arbitrary policy $\pi = (\pi_t)_{t=0}^{T-1}$. For such a policy one can define an associated set of *policy values* $(v_t^\pi(p, z))_{t=0}^T$ that follow the recursion

$$(7.1) \quad v_T^\pi(p, z) = r_T(p, z),$$

$$(7.2) \quad v_t^\pi(p, z) = r_t(p, z, \pi_t(p, z)) + \mathbf{E} \left(v_{t+1}^\pi(\alpha(p, \pi_t(p, z)), W_{t+1} z) \right)$$

for $t = T-1, \dots, 0$. Let us consider a switching system, which starts in a given initial position $p_0^\pi = p_0 \in P$, and state $Z_0 = z_0 \in \mathbf{R}^d$. At any time t , the actions and new positions are determined recursively, following policy $\pi = (\pi_t)_{t=0}^{T-1}$ as

$$a_t^\pi := \pi_t(p_t^\pi, Z_t), \quad p_{t+1}^\pi := \alpha(p_t^\pi, a_t^\pi), \quad t = 0, \dots, T-1.$$

These values define a *policy run* $\mathcal{V}_0^\pi(p_0^\pi, z_0)$, where

$$\mathcal{V}_0^\pi(p_0^\pi, Z_0) = \sum_{s=0}^{T-1} r_t(p_s^\pi, Z_s, a_s^\pi) + r_T(p_T^\pi, Z_T).$$

According to the definition, $v_0^\pi(p_0, z_0)$ is the expected value of the policy run

$$v_0^\pi(p_0, z_0) = \mathbf{E}(\mathcal{V}_0^\pi(p_0, z_0)), \quad p \in P, \quad z \in \mathbf{R}^d.$$

In practice, one can use a Monte Carlo procedure to estimate this value since given a sequence $(\omega_k)_{k \in \mathbf{N}}$ of independent random draws,

$$(7.3) \quad v_0^\pi(p_0, z_0) = \mathbf{E}(\mathcal{V}_0^\pi(p_0, z_0)) = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathcal{V}_0^\pi(p_0, z_0)(\omega_k)$$

holds true due to the strong law of large numbers.

Such a Monte Carlo procedure may be useful for estimating the performance of a given strategy π , and it is motivated by (2.3). However, it does not clarify how far the value $v_0^\pi(p_0, z_0)$ is from what is theoretically possible $v_0^{\pi^*}(p_0, z_0)$.

In the remainder of this section, we suggest a sound solution to this question in terms of a *diagnostic method*. Given a starting point (p_0, z_0) , we explain how the gap

$$(7.4) \quad [v_0^\pi(p_0, z_0), v_0^{\pi^*}(p_0, z_0)]$$

between a given strategy π and the optimal strategy π^* can be assessed. Our methodology is based on a finite sample $\{\omega_1, \dots, \omega_K\}$ of trajectory realizations and utilizes a built-in variance reduction technique to derive *tight confidence bounds* for upper and lower estimates of the interval (7.4).

Let us focus on the upper bound first. Consider a sequence $\varphi = (\varphi_t)_{t=1}^T$ of random mappings

$$(7.5) \quad \varphi_t: P \times \mathbf{R}^d \times A \times \Omega \rightarrow \mathbf{R}, \quad (p, z, a, \omega) \mapsto \varphi_t(p, z, a)(\omega),$$

which for $t = 1, \dots, T$ satisfy

$$(7.6) \quad \mathbf{E}(\varphi_t(p, z, a)) = 0, \quad p \in P, \quad z \in \mathbf{R}^d, \quad a \in A,$$

such that the σ -algebras

$$(7.7) \quad \sigma(\varphi_t(p, z, a), W_t; a \in A, p \in P, z \in \mathbf{R}^d), \quad t = 1, \dots, T,$$

are independent. Given these mappings $\varphi = (\varphi_t)_{t=1}^T$, we now introduce the random functions $(\bar{v}_t^\varphi)_{t=0}^T$,

$$\bar{v}_t^\varphi: P \times \mathbf{R}^d \times \Omega \rightarrow \mathbf{R}, \quad t = 0, \dots, T,$$

which are recursively defined for $t = T, \dots, 1$ via

$$(7.8) \quad \bar{v}_T^\varphi(p, z) = r_T(p, z),$$

$$(7.9) \quad \bar{v}_t^\varphi(p, z) = \max_{a \in A} \left(r_t(p, z, a) + \varphi_{t+1}(p, z, a) + \bar{v}_{t+1}^\varphi(\alpha(p, a), W_{t+1}z) \right).$$

Using $(\bar{v}_t^\varphi)_{t=0}^T$, the following theorem holds.

THEOREM 7.1. (i) *For each policy $\pi = (\pi_t)_{t=0}^{T-1}$, it holds that the policy values $(v_t^\pi)_{t=0}^T$ are dominated from above:*

$$(7.10) \quad v_t^\pi(p, z) \leq \mathbf{E}(\bar{v}_t^\varphi(p, z)) \quad \text{for all } t = 0, \dots, T, \quad p \in P, \quad z \in \mathbf{R}^d.$$

(ii) *Given the policy values $(v_t^{\pi^*})_{t=0}^T$ associated with the optimal policy $\pi^* = (\pi_t^*)_{t=0}^{T-1}$, let $(\varphi_t^*)_{t=1}^T$ be defined by*

$$(7.11) \quad \varphi_{t+1}^*(p, z, a) = \mathbf{E} \left(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) \right) - v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)$$

for all $p \in P, z \in \mathbf{R}^d, a \in A$, and $t = 0, \dots, T - 1$. It then holds that the mappings $(\varphi_t^*)_{t=1}^T$ satisfy (7.5)–(7.7) and that (7.10) holds with equality

$$(7.12) \quad v_t^{\pi^*}(p, z) = \bar{v}_t^{\varphi^*}(p, z) \quad \text{for all } t = 0, \dots, T, \quad p \in P, \quad z \in \mathbf{R}^d.$$

Proof. (i) The value $(v_t^\pi)_{t=0}^T$ of the policy $\pi = (\pi_t)_{t=0}^{T-1}$ satisfies the recursion (7.2). Using this recursion and (7.6) we obtain

$$(7.13) \quad \begin{aligned} v_t^\pi(p, z) &= \mathbf{E} \left(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z)) \right) \\ &\quad + \mathbf{E} \left(v_{t+1}^\pi(\alpha(p, \pi_t(p, z)), W_{t+1}z) \right). \end{aligned}$$

Now, let us prove the assertion (7.10) by induction. For $t = T$, the inequality (7.10) holds with equality because of the initialization

$$(7.14) \quad v_T^\pi = r_T = \bar{v}_T^\varphi$$

in (7.1) and (7.8). Given the induction assumption

$$v_{t+1}^\pi(p, z) \leq \mathbf{E}(\bar{v}_{t+1}^\varphi(p, z)) \quad \text{for all } p \in P, z \in \mathbf{R}^d,$$

we use (7.7) to conclude that

$$v_{t+1}^\pi(\alpha(p, \pi_t(p, z)), W_{t+1}z) \leq \mathbf{E}(\bar{v}_{t+1}^\varphi(\alpha(p, \pi_t(p, z)), W_{t+1}z) | W_{t+1})$$

holds for all $p \in P, z \in \mathbf{R}^d$. Using this, we obtain in (7.13) an estimate

$$\begin{aligned} v_t^\pi(p, z) &\leq \mathbf{E}\left(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z))\right) \\ &\quad + \mathbf{E}\left(\mathbf{E}(\bar{v}_{t+1}^\varphi(\alpha(p, \pi_t(p, z)), W_{t+1}z) | W_{t+1})\right), \end{aligned}$$

from which the assertion follows that

$$\begin{aligned} v_t^\pi(p, z) &= \mathbf{E}\left(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z))\right) \\ &\quad + \bar{v}_{t+1}^\varphi(\alpha(p, \pi_t(p, z)), W_{t+1}z) \\ &\leq \mathbf{E}\left(\max_{a \in A} \left[r_t(p, z, a) + \varphi_{t+1}(p, z, a) + \bar{v}_{t+1}^\varphi(\alpha(p, a), W_{t+1}z)\right]\right) \\ &\leq \mathbf{E}(\bar{v}_t^\varphi(p, z)), \end{aligned}$$

where the last step results from the recursion (7.9).

(ii) Now suppose that π^* is an optimal policy and define $\varphi^* = (\varphi_t^*)_{t=1}^T$ as in (7.11), which satisfies the assumption (7.6). Furthermore, the independence (7.7) holds since for $t = 0, \dots, T-1$ the random component in $\varphi_{t+1}^*(p, z, a)$ is

$$v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z),$$

which is in turn a function of W_{t+1} . Let us now verify the assertion (7.12). By the induction started in (7.8) we can assume that $v_{t+1}^{\pi^*} = \bar{v}_{t+1}^{\varphi^*}$ holds. Using this, we deduce for all $p \in P, z \in \mathbf{R}^d$, and $a \in A$ the assertion

$$\begin{aligned} (7.15) \quad &r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + \bar{v}_{t+1}^{\varphi^*}(\alpha(p, a), W_{t+1}z) \\ &= r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z). \end{aligned}$$

On the other hand, using (7.11), we infer that for all $p \in P, z \in \mathbf{R}^d$, and $a \in A$ it holds that

$$\begin{aligned} (7.16) \quad &r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) \\ &= r_t(p, z, a) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)) \\ &\quad - v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) \\ &= r_t(p, z, a) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)). \end{aligned}$$

Now, in the recursion (7.9) we replace (7.15) by (7.16) to obtain the desired result (7.12):

$$\begin{aligned} \bar{v}_t^{\varphi^*}(p, z) &= \max_{a \in A} \left(r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + \bar{v}_{t+1}^{\varphi^*}(\alpha(p, a), W_{t+1}z) \right) \\ &= \max_{a \in A} \left(r_t(p, z, a) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)) \right) \\ &= v_t^{\pi^*}(p, z), \quad p \in P, \quad z \in \mathbf{R}^d. \end{aligned}$$

Let us elaborate on a practical application of this technique. Suppose that we attempt to assess the distance to optimality of an approximate policy $\tilde{\pi}$, obtained by a numerical procedures described previously. According to (i) of Theorem 7.1, arbitrary $(\varphi_t)_{t=1}^T$ satisfying (7.6) and (7.7) yields an upper bound

$$(7.17) \quad v_0^{\tilde{\pi}}(p, z) \leq v_0^{\pi^*}(p, z) \leq \mathbf{E}(\bar{v}_0^\varphi(p, z)), \quad p \in P, \quad z \in \mathbf{R}^d.$$

Note that the expectation $\mathbf{E}(\bar{v}_0^\varphi(p, z))$ will be estimated via a Monte Carlo procedure. Thus, we obtain the following estimation procedure.

Upper bound estimation.

(1) Given a switching system, implement $(\varphi_t)_{t=1}^T$ which satisfies (7.5), (7.6), and (7.7).

(2) Choose a number $K \in \mathbf{N}$ of Monte Carlo trials and obtain for $k = 1, \dots, K$ independent realizations $(W_t(\omega_k))_{t=1}^T$ of disturbances. (3) Starting at $z_0^k := z_0 \in \mathbf{R}^d$, define for $k = 1, \dots, K$ the trajectories $(z_t^k)_{t=0}^T$ recursively,

$$z_{t+1}^k = W_{t+1}(\omega_k)z_t^k, \quad t = 0, \dots, T - 1,$$

and determine realizations

$$\varphi_{t+1}(p, z_t^k, a)(\omega_k), \quad t = 0, \dots, T - 1, \quad k = 1, \dots, K.$$

(4) For each $k = 1, \dots, K$ initialize the recursion at $t = T$ as

$$\bar{v}_T^\varphi(p, z_T^k) = r_T(p, z_T^k) \quad \text{for all } p \in P$$

and continue for $t = T - 1, \dots, 0$ by

$$\bar{v}_t^\varphi(p, z_t^k) = \max_{a \in A} (r_t(p, z_t^k, a) + \varphi_{t+1}(p, z_t^k, a)(\omega_k) + \bar{v}_{t+1}^\varphi(\alpha(p, a), z_{t+1}^k)).$$

Store the value as $\bar{v}_0^\varphi(p, z_0^k)$ for $k = 1, \dots, K$.

(5) Determine the sample mean $\frac{1}{K} \sum_{k=1}^K \bar{v}_0^\varphi(p, z_0^k)$ and its upper confidence bound to estimate $v_0^{\pi^*}(p, z_0)$ from above.

To obtain a tight upper bound, $(\varphi_t)_{t=1}^T$ must be chosen accordingly. Thereby, assertion (ii) of Theorem 7.1 suggests an appropriate choice. Namely, in the hypothetical case that the optimal policy value functions $(v_t^{\pi^*})_{t=0}^T$ are known, the $(\varphi_t^*)_{t=1}^T$ that is obtained via (7.11) will give an exact and nonrandom upper bound. In practice, this situation is not feasible, since an optimal strategy π^* is not known. Instead, we suggest using an approximate value function $(\tilde{\varphi}_t)_{t=0}^T$, returned by one of the algorithms described in this work. That is, following (7.11), a reasonable candidate for $t = 0, \dots, T - 1$ could be given as

$$(7.18) \quad \varphi_{t+1}(p, z, a) = \mathbf{E}(\tilde{v}_{t+1}(\alpha(p, a), W_{t+1}z)) - \tilde{v}_{t+1}(\alpha(p, a), W_{t+1}z).$$

Note that this choice involves an exact calculation of $\mathbf{E}(\tilde{v}_{t+1}(\alpha(p, a), W_{t+1}z))$, which is not possible in practice. For this reason, we suggest a modification. We introduce φ_{t+1} as in (7.18), with the expectation replaced by an arithmetic mean over a number I of independent copies $(W_{t+1}^{(i)})_{i=1}^I$ of W_{t+1} . That is, given independent random variables W_{t+1} and $W_{t+1}^{(i)}$ for $i = 1, \dots, I$ and $t = 0, \dots, T - 1$ such that the distribution of $W_{t+1}^{(i)}$ equals that of W_{t+1} , we define for all $t = 0, \dots, T - 1$, $a \in A$, $p \in P$, and $z \in \mathbf{R}^d$,

$$(7.19) \quad \varphi_{t+1}(p, z, a) = \frac{1}{I} \sum_{i=1}^I \tilde{v}_{t+1}(\alpha(p, a), W_{t+1}^{(i)}z) - \tilde{v}_{t+1}(\alpha(p, a), W_{t+1}z).$$

With this definition, $(\varphi_t)_{t=1}^T$ satisfies conditions (7.6) and (7.7), and we thus obtain a valid and computable upper bound.

Let us turn now to the estimation of the lower boundary of the interval (7.4). Given a strategy $\pi = (\pi_t)_{t=0}^{T-1}$, the value $v_0^\pi(p_0, z_0)$ can in principle be approached as in (7.3) from test runs of the strategy in a series of independent back-testing experiments. However, it turns out that a slight adaptation of the upper bound technique provides far better results, due to a built-in variance reduction technique. Similarly to part (ii) of the previous theorem, which indicates that the variance of the Monte Carlo trials reduces, if approximate solution is close to the optimal one, we establish a recursive procedure with a control variate built in. The idea is simple: given a nearly optimal policy $\pi = (\pi_t)_{t=0}^{T-1}$ we alter the recursion (7.8), (7.9), replacing the maximization by an exact choice of the action according to the policy $\pi = (\pi_t)_{t=0}^{T-1}$.

Given a sequence $\varphi = (\varphi_t)_{t=1}^T$ satisfying (7.6) and (7.7) we introduce the random functions $(\underline{v}_t^{\pi, \varphi})_{t=0}^T$,

$$\underline{v}_t^{\pi, \varphi} : P \times \mathbf{R}^d \times \Omega \rightarrow \mathbf{R}, \quad t = 0, \dots, T,$$

which are recursively defined for $t = T, \dots, 1$ via

$$(7.20) \quad \underline{v}_T^{\pi, \varphi}(p, z) = r_T(p, z),$$

$$(7.21) \quad \underline{v}_t^{\pi, \varphi}(p, z) = r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z)) \\ + \underline{v}_{t+1}^{\pi, \varphi}(\alpha(p, \pi_t(p, z)), W_{t+1}z).$$

The following theorem holds for $(\underline{v}_t^{\pi, \varphi})_{t=0}^T$.

THEOREM 7.2. (i) *Given $\varphi = (\varphi_t)_{t=1}^T$ as in (7.5) satisfying (7.6) and (7.7) and a policy $\pi = (\pi_t)_{t=0}^{T-1}$, introduce $(\underline{v}_t^{\pi, \varphi})_{t=0}^T$ by (7.20), (7.21). It holds that*

$$(7.22) \quad v_t^\pi(p, z) = \mathbf{E}(\underline{v}_t^{\pi, \varphi}(p, z)) \quad \text{for all } t = 0, \dots, T, \quad p \in P, \quad z \in \mathbf{R}^d.$$

(ii) *Given the value $(v_t^{\pi^*})_{t=0}^T$ of the optimal policy $\pi^* = (\pi_t^*)_{t=0}^{T-1}$, define $(\varphi_t^*)_{t=1}^T$ as in (7.11). Then the mappings $(\varphi_t^*)_{t=1}^T$ satisfy (7.5)–(7.7) such that (7.22) holds with equality as*

$$(7.23) \quad v_t^{\pi^*}(p, z) = \underline{v}_t^{\pi^*, \varphi^*}(p, z) \quad \text{for all } t = 0, \dots, T, \quad p \in P, \quad z \in \mathbf{R}^d.$$

Proof. (i) The value $(v_t^\pi)_{t=0}^T$ of the policy $\pi = (\pi_t)_{t=0}^{T-1}$ satisfies the recursion (7.2). Using this recursion and (7.6) we obtain

$$(7.24) \quad v_t^\pi(p, z) = \mathbf{E}(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z))) \\ + \mathbf{E}\left(v_{t+1}^\pi(\alpha(p, \pi_t(p, z)), W_{t+1}z)\right).$$

Now, let us prove the assertion (7.22) by induction. For $t = T$ the inequality (7.22) holds with equality because of the initialization

$$(7.25) \quad v_T^\pi = r_T = \underline{v}_T^{\pi, \varphi}$$

in (7.1) and (7.20). Given the induction assumption

$$v_{t+1}^\pi(p, z) = \mathbf{E}(\underline{v}_{t+1}^{\pi, \varphi}(p, z)) \quad \text{for all } p \in P, \quad z \in \mathbf{R}^d,$$

we use (7.7) to conclude that

$$v_{t+1}^\pi(\alpha(p, \pi_t(p, z)), W_{t+1}z) = \mathbf{E}\left(\underline{v}_{t+1}^{\pi, \varphi}(\alpha(p, \pi_t(p, z)), W_{t+1}z) \mid W_{t+1}\right)$$

holds for all $p \in P, z \in \mathbf{R}^d$. Using this, we obtain in (7.24) the equality

$$\begin{aligned} v_t^\pi(p, z) &= \mathbf{E}\left(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z))\right) \\ &\quad + \mathbf{E}\left(\mathbf{E}\left(\underline{v}_{t+1}^{\pi, \varphi}(\alpha(p, \pi_t(p, z))W_{t+1}z) \mid W_{t+1}\right)\right) \\ &= \mathbf{E}\left(r_t(p, z, \pi_t(p, z)) + \varphi_{t+1}(p, z, \pi_t(p, z)) + \underline{v}_{t+1}^{\pi, \varphi}(\alpha(p, \pi_t(p, z))W_{t+1}z)\right). \end{aligned}$$

By using the recursion (7.21), the assertion (7.22) follows.

(ii) Let us now verify the assertion (7.23). By induction which is started as in (7.20) we can assume that $v_{t+1}^{\pi^*} = \underline{v}_{t+1}^{\pi^*, \varphi^*}$ holds. Using this, we deduce for all $p \in P, z \in \mathbf{R}^d$, and $a \in A$ the assertion

$$(7.26) \quad \begin{aligned} r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + \underline{v}_{t+1}^{\pi^*, \varphi^*}(\alpha(p, a), W_{t+1}z) \\ = r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z). \end{aligned}$$

On the other hand, using (7.11), we infer that for all $p \in P, z \in \mathbf{R}^d$, and $a \in A$ it holds that

$$(7.27) \quad \begin{aligned} r_t(p, z, a) + \varphi_{t+1}^*(p, z, a) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) \\ = r_t(p, z, a) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)) \\ - v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) + v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z) \\ = r_t(p, z, a) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, a), W_{t+1}z)). \end{aligned}$$

Now, in recursion (7.21) we replace (7.26) by (7.27) to obtain the desired claim (7.12):

$$\begin{aligned} \underline{v}_t^{\pi^*, \varphi^*}(p, z) &= r_t(p, z, \pi_t^*(p, z)) + \varphi_{t+1}^*(p, z, \pi_t^*(p, z)) + \underline{v}_{t+1}^{\pi^*, \varphi^*}(\alpha(p, \pi_t^*(p, z)), W_{t+1}z) \\ &= r_t(p, z, \pi_t^*(p, z)) + \mathbf{E}(v_{t+1}^{\pi^*}(\alpha(p, \pi_t^*(p, z)), W_{t+1}z)) \\ &= v_t^{\pi^*}(p, z), \quad p \in P, \quad z \in \mathbf{R}^d, \quad a \in A. \end{aligned}$$

Theorem 2 is proved.

The practical implementation of the lower bound estimation is based on the same realization of $(\varphi_t)_{t=1}^T$ as in (7.19) using independent copies of disturbances. Let us summarize this procedure as follows.

Lower bound estimation.

(1) Given approximate value functions $(\tilde{v}_t)_{t=0}^T$ and a corresponding strategy $\tilde{\pi} = (\tilde{\pi}_t)_{t=0}^{T-1}$, choose $\varphi = (\varphi_t)_{t=0}^{T-1}$ as in (7.19).

(2) Given $K \in \mathbf{N}$ Monte Carlo trials, obtain for $k = 1, \dots, K$ independent realizations $(W_t(\omega_k))_{t=1}^T$ of disturbances.

(3) Starting at $z_0^k := z_0 \in \mathbf{R}^d$, define for $k = 1, \dots, K$ trajectories $(z_t^k)_{t=0}^T$ recursively,

$$z_{t+1}^k = W_{t+1}(\omega_k)z_t^k, \quad t = 0, \dots, T - 1,$$

and determine realizations

$$\varphi_{t+1}(p, z_t^k, a)(\omega_k), \quad t = 0, \dots, T - 1, \quad k = 1, \dots, K.$$

(4) For each $k = 1, \dots, K$ initialize the recursion at $t = T$ as

$$\underline{v}_T^{\tilde{\pi}, \varphi}(p, z_T^k) = r_T(p, z_T^k) \quad \text{for all } p \in P$$

and continue for $t = T - 1, \dots, 0$ and for all $p \in P$ by

$$(7.28) \quad \begin{aligned} \underline{v}_t^{\tilde{\pi}, \varphi}(p, z_t^k) &= r_t(p, z_t^k, \tilde{\pi}_t(p, z_t^k)) + \varphi_{t+1}(p, z_t^k, \tilde{\pi}_t(p, z_t^k))(\omega_k) \\ &\quad + \underline{v}_{t+1}^{\tilde{\pi}, \varphi}(\alpha(p, \tilde{\pi}_t(p, z_t^k)), z_{t+1}^k). \end{aligned}$$

Store the value as $\underline{v}_0^{\tilde{\pi}, \varphi}(p, z_0^k)$ for $k = 1, \dots, K, p \in P$.

(5) Calculate the sample mean $\frac{1}{K} \sum_{k=1}^K \underline{v}_0^{\tilde{\pi}, \varphi}(p, z_0^k)$ and use its lower confidence bound to estimate $v_0^*(p, z_0)$ for $p \in P$ from below.

8. Examples. In the literature, the estimation of a complementary upper bound for the optimal stopping problem relies heavily on the concept of martingale duality and has been addressed in [14], [24], and [16]. From a computational perspective, achieving a sufficiently tight upper bound is equivalent to constructing a “good” martingale, and a tractable algorithm to do so was given by [1]. Upper bound methods have since been extended to the more general class of optimal multiple stopping problems by [19], [26], and [17]. Finally, the paper [25] generalizes this technique to a wider class of discrete-time stochastic control problems. The combination of upper and lower bound methods is known as *primal-dual simulation*. In this section, we compare our technique to results achieved using a standard least-squares regression method.

We will now perform value function approximations using the method outlined in section 6 and the associated diagnostics established in section 7 on two examples of Markov decision problems, an optimal stopping and an optimal multiple exercise problem. Optimal stopping problems are an important subclass of Markov decision problems (see Chapters 10 and 11 of [2]), whose upper bound estimation using duality is well studied. As an illustration of our approach, we obtain in section 8.1 bounds on the price of the *Bermudan* put option, a practically important discrete-time optimal stopping problem. In section 8.2, we use these methods to obtain numerical solutions to an optimal multiple stopping problem—the *swing* option (see [6]). A swing option allows the holder to buy a fixed quantity of the underlying instrument at a predetermined price more than once before the maturity of the option. However, there is a limit to the maximum number of times this can be done. Swing options are predominant in commodity markets, particularly in the energy sector.

For both applications, we will consider the evolution of the discounted asset price $(S_t)_{t=0}^T$ in discrete time, with respect to a risk-neutral measure. The dynamics $(S_t)_{t=0}^T$ of the discounted price depends on the asset type. For the Bermudan put, the discounted price process $(S_t)_{t=0}^T$ is modeled as a martingale in the risk-neutral measure. For the swing option, we suppose that the price process $(S_t)_{t=0}^T$ is modeled by the exponential of an Ornstein–Uhlenbeck process to explain the mean-reverting price property naturally expected for commodity prices.

The logarithm $(Z_t)_{t=0}^T$ of the price forms the continuous component of our state dynamics. In practice, a further transformation of the state space is usually required before linear state dynamics can be achieved. In most cases, an augmentation with 1 via

$$Z_t = \begin{bmatrix} \tilde{Z}_t \\ 1 \end{bmatrix}, \quad t = 0, \dots, T,$$

is needed to represent the evolution of the continuous state component. In this representation, the system state follows a multiplicative dynamic

$$Z_{t+1} = W_{t+1} Z_t, \quad t = 0, \quad T - 1,$$

with independent and identically distributed matrix-valued random variables $(W_t)_{t=1}^T$. The entries of these disturbance matrices reflect the underlying process model.

The grid choice is a key ingredient in the algorithm. For multivariate state processes, a convenient way of grid construction is by simulation of appropriate trajectories. Thus, we create a grid of a desired size by simulating and storing a sufficient number of paths of $(Z_t)_{t=0}^{k_p T}$ of an appropriate length $k_p T \in \mathbf{N}$. In our examples, we have used a number of steps that is twice that of the time horizon ($k_p = 2$). The distribution of disturbances is approximated by a discrete distribution. For this, a sample of $(W(k))_{k=1}^n$ of independent realizations was generated and stored. All required steps from section 6 and the Monte Carlo simulation for diagnostics refer to this discrete distribution approximation. For bound computations, we use confidence intervals based on K simulated trajectories. More precisely, we quote the intervals as

$$(8.1) \quad \left[\underline{\mu} - \Phi^{-1}\left(1 - \frac{x}{2}\right) \frac{\underline{\sigma}}{\sqrt{K}}, \quad \bar{\mu} + \Phi^{-1}\left(1 - \frac{x}{2}\right) \frac{\bar{\sigma}}{\sqrt{K}} \right],$$

where $1 - x$ denotes the confidence level and $(\underline{\mu}, \underline{\sigma})$ and $(\bar{\mu}, \bar{\sigma})$ denote the sample mean and sample standard deviation of $(\underline{z}_0^{\tilde{\pi}, \varphi}(p, z_0^k))_{k=1}^K$ and $(\bar{\sigma}_0^{\varphi}(p, z_0^k))_{k=1}^K$, respectively.

8.1. The Bermudan put option. This option allows the holder to sell the underlying asset at a prespecified strike price on a discrete set of exercise dates up to and including the expiry date of the option. The fair price of a Bermudan put is given by the supremum

$$\sup_{\tau} \mathbf{E} [(Ke^{-\rho\tau} - S_{\tau})^+],$$

where τ runs through all $\{0, \dots, T\}$ -valued stopping times. First, let us express this control problem as a switching system. We use the position set $P = \{1, 2\}$ to indicate whether the option has been exercised ($p = 1$) or not ($p = 2$). The action set $A = \{1, 2\}$ represents the choice between exercising ($a = 1$) or not exercising ($a = 2$). The control α of the discrete component of the state space

$$(\alpha(p, a))_{p,a=1}^2 \sim \begin{bmatrix} \alpha(1, 1) & \alpha(1, 2) \\ \alpha(2, 1) & \alpha(2, 2) \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$$

ensures that $p = 1$ is absorbing. The continuous state space component follows

$$(8.2) \quad \tilde{Z}_{t+1} = \tilde{Z}_t + \gamma + \beta \varepsilon_{t=1}^T, \quad \tilde{Z}_0 = \log S_0,$$

where $(\varepsilon)_{t=1}^T$ are independent standard normally distributed random variables. We set the parameters as $\gamma = -\frac{1}{2} \sigma^2 \Delta$ and $\beta = \sigma \sqrt{\Delta}$, where $\Delta > 0$ is the time duration (in years) from time point t to $t + 1$, and $\sigma > 0$ represents the volatility of the process that is measured on a yearly scale. Given this price process, the disturbances are given as

$$W_{t+1} = \begin{bmatrix} 1 & \gamma + \beta \varepsilon_{t+1} \\ 0 & 1 \end{bmatrix}, \quad t = 0, \dots, T - 1.$$

For the two-dimensional space \mathbf{R}^2 with the evolution of the continuous component as above, let us now determine all reward and scrap functions. Consider a realization of the continuous component $(z^{(1)}, z^{(2)}) \in \mathbf{R}^2$ at the current time $t = 0, \dots, T - 1$; then, given a Bermudan put option, the action $a \in \{1, 2\}$ leads to the reward

$$(8.3) \quad r_t(p, (z^{(1)}, z^{(2)}), a) = (Ke^{-\rho t} - e^{z^{(1)}})^+(p - \alpha(p, a))$$

TABLE 8.1
*Bermudan put option with $K = 40$.**

S_0	σ	Maturity	Confidence interval	LSM mean	LSM se
36	0.2	1	[4.4763, 4.4768]	4.472	.0100
36	0.2	2	[4.8296, 4.8312]	4.821	.0120
36	0.4	1	[7.0989, 7.0992]	7.091	.0200
36	0.4	2	[8.4965, 8.4968]	8.488	.0240
38	0.2	1	[3.2481, 3.2489]	3.244	.0090
38	0.2	2	[3.7355, 3.7370]	3.735	.0110
38	0.4	1	[6.1451, 6.1452]	6.139	.0190
38	0.4	2	[7.6580, 7.6583]	7.669	.0220
40	0.2	1	[2.3119, 2.3129]	2.313	.0090
40	0.2	2	[2.8765, 2.8776]	2.879	.0100
40	0.4	1	[5.3093, 5.3094]	5.308	.0180
40	0.4	2	[6.9075, 6.9077]	6.921	.0220
42	0.2	1	[1.6150, 1.6158]	1.617	.0070
42	0.2	2	[2.2053, 2.2060]	2.206	.0100
42	0.4	1	[4.5797, 4.5798]	4.588	.0170
42	0.4	2	[6.2351, 6.2354]	6.243	.0210
44	0.2	1	[1.1081, 1.1087]	1.118	.0070
44	0.2	2	[1.6836, 1.6843]	1.675	.0090
44	0.4	1	[3.9449, 3.9450]	3.957	.0170
44	0.4	2	[5.6324, 5.6326]	5.622	.0210

*These results were produced using a grid size of $m = 1024$ and disturbances of size $n = 4096$. Diagnostics is based on 1024 sample paths, and 99% confidence bounds are calculated by setting $x = 0.01$ in (8.1). For comparison, the means and standard errors obtained by least-squares Monte Carlo are given in the last two columns *LSM mean* and *LSM se*, respectively; they are cited from [18], where numbers were quoted with three decimal points.

for all $p \in P$ and $a \in A$. At final time T , we suppose that the put option is exercised automatically, which gives, for all $p \in P$, the scrap value

$$(8.4) \quad r_T(p, (z^{(1)}, z^{(2)})) = (Ke^{-\rho T} - e^{z^{(1)}})^+(p - \alpha(p, 1)).$$

Note that the reward and scrap functions are not convex in the continuous component $z = (z^{(1)}, z^{(2)})$. Hence we decompose them into the difference of two convex functions

$$(8.5) \quad r_t(p, \cdot, a) = \check{r}_t(p, \cdot, a) - \hat{r}_t(p, \cdot, a),$$

$$(8.6) \quad r_T(p, \cdot) = \check{r}_T(p, \cdot) - \hat{r}_T(p, \cdot),$$

given by

$$\check{r}_t(p, (z^{(1)}, z^{(2)}), a) = (e^{z^{(1)}} - Ke^{-\rho t})^+(p - \alpha(p, a)),$$

$$\hat{r}_t(p, (z^{(1)}, z^{(2)}), a) = (e^{z^{(1)}} - Ke^{-\rho t})(p - \alpha(p, a)),$$

$$\check{r}_T(p, (z^{(1)}, z^{(2)})) = (e^{z^{(1)}} - Ke^{-\rho T})^+(p - \alpha(p, 1)),$$

$$\hat{r}_T(p, (z^{(1)}, z^{(2)})) = (e^{z^{(1)}} - Ke^{-\rho T})(p - \alpha(p, 1))$$

for all $p \in P, a \in A$, and $(z^{(1)}, z^{(2)}) \in \mathbf{R}^2$.

We compare our results with the low-biased estimates given in the literature for the Bermudan put, where the risk-free rate is 0.06 and the strike is set at 40. The results are given in Table 8.1 for different combinations of initial prices, volatilities, and maturities.

8.2. The swing option. We consider a specific case of the swing option, referred to as a *unit-time refraction period* condition. This condition limits the holder to exercise one right at any time. Given the discounted asset price $(S_t)_{t=0}^T$, the price of a swing option with N rights is given by the supremum

$$\sup_{0 \leq \tau_1 < \dots < \tau_N \leq T} \mathbf{E} \left[\sum_{k=1}^N (S_{\tau_k} - Ke^{-\rho\tau_k})^+ \right],$$

over all stopping times τ_1, \dots, τ_N with values in $\{0, \dots, T\}$. In order to represent this control problem as a switching system, we use the position set $P = \{1, \dots, N + 1\}$ to represent the number of rights remaining. That is, $p \in P$ stands for the situation when there are $p - 1$ rights remaining to be exercised. The action set $A = \{1, 2\}$ is the same as in the case of the Bermudan put with control matrix α given by

$$(\alpha(p, a))_{p,a} \sim \begin{bmatrix} \alpha(1, 1) & \alpha(1, 2) \\ \alpha(2, 1) & \alpha(2, 2) \\ \alpha(3, 1) & \alpha(3, 2) \\ \dots & \dots \\ \alpha(N + 1, 1) & \alpha(N + 1, 2) \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 2 & 3 \\ \dots & \dots \\ N & N + 1 \end{bmatrix}.$$

Having modeled the discounted stock price process as an exponential mean-reverting process with reversion parameter $\kappa \in [0, 1[$, long run mean $\mu > 0$, and volatility $\sigma > 0$, we obtain the logarithm of the discounted price process as

$$(8.7) \quad \tilde{Z}_{t+1} = (1 - \kappa)(\tilde{Z}_t - \mu) + \mu + \sigma\varepsilon_{t+1}, \quad \tilde{Z}_0 = \log S_0$$

with the disturbance matrix

$$W_{t+1} = \begin{bmatrix} (1 - \kappa) & \kappa\mu + \sigma\varepsilon_{t+1} \\ 0 & 1 \end{bmatrix}, \quad t = 0, \dots, T - 1.$$

The reward and scrap values are given by

$$(8.8) \quad r_t(p, (z^{(1)}, z^{(2)}), a) = (e^{z^{(1)}} - Ke^{-\rho t})^+(p - \alpha(p, a)), \quad t = 0, \dots, T - 1,$$

and

$$(8.9) \quad r_T(p, (z^{(1)}, z^{(2)})) = (e^{z^{(1)}} - Ke^{-\rho T})^+(p - \alpha(p, 1)),$$

respectively, for all $p \in P$ and $a \in A$.

In Table 8.2 we compare our results to those given in [19] with bounds on the swing option price where the underlying process is assumed to follow the dynamics (8.7) with parameters

$$\rho = 0, \quad \kappa = 0.9, \quad \mu = 0, \quad \sigma = 0.5, \quad S_0 = 1, \quad K = 0, \quad T = 1000.$$

TABLE 8.2
Swing option numerical results.[†]

	CSS	MH
Position (rights + 1)	Confidence interval	Confidence interval
2	[4.737, 4.761]	[4.773, 4.794]
3	[9.005, 9.031]	[9.016, 9.091]
4	[13.001, 13.026]	[12.959, 13.100]
5	[16.805, 16.830]	[16.773, 16.906]
6	[20.465, 20.491]	[20.439, 20.580]
11	[37.339, 37.363]	[37.305, 37.540]
16	[52.694, 52.718]	[52.670, 53.009]
21	[67.070, 67.095]	[67.050, 67.525]
31	[93.811, 93.835]	[93.662, 94.519]
41	[118.639, 118.663]	[118.353, 119.625]
51	[142.059, 142.084]	[141.703, 143.360]
61	[164.368, 164.392]	[163.960, 166.037]
71	[185.757, 185.781]	[185.335, 187.729]
81	[206.362, 206.386]	[205.844, 208.702]
91	[226.284, 226.308]	[225.676, 228.985]
101	[245.601, 245.625]	[244.910, 248.651]

[†]Results were produced using a grid size of $m = 1024$ and disturbances of size $n = 1024$. Diagnostics is based on 1024 sample paths, and 99% confidence bounds are calculated by setting $x = 0.01$. The columns under MH denote the results from [19].

9. Conclusion. In this work we present a novel class of algorithms to solve stochastic switching problems whose processes follow linear state space dynamics. Our methodology is directly applicable to high-dimensional problems and shows remarkable numerical efficiency and excellent precision. More importantly, we adapt the primal-dual approach to estimate the distance to optimality of approximate solutions using Monte Carlo techniques. With this, we establish a sound and reliable diagnostics and quality assessment tool for a posteriori justification of the numerical approximation. We believe that such a combination of efficient numerical schemes with a subsequent diagnostic check can be very useful in practical applications. This approach may help in the development and justification of further approximate methods. In this context, natural extensions of the present scheme (say, from linear to piecewise linear dynamics) can be examined in detail. We address this promising direction in future research.

REFERENCES

- [1] L. ANDERSEN AND M. BROADIE, *A primal-dual simulation algorithm for pricing multidimensional American options*, Management Sci., 50 (2004), pp. 1222–1234.
- [2] N. BÄUERLE AND U. RIEDER, *Markov Decision Processes with Applications to Finance*, Springer, Heidelberg, 2011.
- [3] D. BELOMESTNY, A. KOLODKO, AND J. SCHOENMAKERS, *Regression methods for stochastic control problems and their convergence analysis*, SIAM J. Control Optim., 48 (2010), pp. 3562–3588.
- [4] D. P. BERTSEKAS, *Dynamic Programming and Optimal Control*, Athena Scientific, Nashua, 2005.
- [5] D. P. BERTSEKAS AND J. N. TSITSIKLIS, *Neuro-Dynamic Programming*, Athena Scientific, Nashua, 1996.
- [6] R. CARMONA AND N. TOUZI, *Optimal multiple stopping and valuation of swing options*, Math. Finance, 18 (2008), pp. 239–268.
- [7] J. F. CARRIERE, *Valuation of the early-exercise price for options using simulations and non-parametric regression*, Insurance Math. Econom., 19 (1996), pp. 19–30.

- [8] E. CLEMENT, D. LAMBERTON, AND P. PROTTER, *An analysis of the Longstaff-Schwartz algorithm for American option pricing*, Finance Stoch., 6 (2002), pp. 449–471.
- [9] D. EGLOFF, *Monte Carlo algorithms for optimal stopping and statistical learning*, Appl. Probab., 15 (2005), pp. 1396–1432.
- [10] D. EGLOFF, M. KOHLER, AND N. TODOROVIC, *A dynamic look-ahead Monte Carlo algorithm*, Ann. Appl. Probab., 17 (2007), pp. 1138–1171.
- [11] J. FAN AND I. GJIBELS, *Local Polynomial Modelling and Its Applications*, Chapman and Hall, London, 1996.
- [12] E. A. FEINBERG AND A. SHWARTZ, *Handbook of Markov Decision Processes*. Kluwer, Norwell, MA, 2002.
- [13] P. GLASSERMAN, *Monte Carlo Methods in Financial Engineering*, Springer, New York, 2003.
- [14] M. HAUGH AND L. KOGAN, *Pricing American options: A duality approach*, Oper. Res., 52 (2004), pp. 258–270.
- [15] J. HINZ, *Optimal stochastic switching under convexity assumptions*, SIAM J. Control Optim., 52 (2014), pp. 164–188.
- [16] F. JAMSHIDIAN, *Numeraire-Invariant Option Pricing & American, Bermudan, and Trigger Stream Rollover*, UT Publications, University of Twente, 2004.
- [17] M. S. JOSHI AND N. YAP, *The Multiplicative Dual for Multiple-Exercise Options*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2430558, 2014.
- [18] F. LONGSTAFF AND E. SCHWARTZ, *Valuing American options by simulation: A simple least-squares approach*, Rev. Financial Stud., 14 (2001), pp. 113–147.
- [19] N. MEINSHAUSEN AND B. M. HAMBLY, *Monte Carlo methods for the valuation of multiple-exercise options*, Math. Finance, 14 (2004), pp. 557–583.
- [20] D. ORMONEIT AND P. GLYNN, *Kernel-based reinforcement learning*, Machine Learning, 49 (2002), pp. 161–178.
- [21] D. ORMONEIT AND P. GLYNN, *Kernel-based reinforcement learning in average-cost problems*, IEEE Trans. Automat. Control, 47 (2002), pp. 1624–1636.
- [22] W. B. POWELL, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley, New York, 2007.
- [23] M. L. PUTERMAN, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley, New York, 1994.
- [24] L. C. G. ROGERS, *Monte Carlo valuation of American options*, Math. Finance, 12 (2002), pp. 271–286.
- [25] L. C. G. ROGERS, *Pathwise stochastic optimal control*, SIAM J. Control Optim., 46 (2007), pp. 1116–1132.
- [26] J. SCHOENMAKERS, *A pure martingale dual for multiple stopping*, Finance Stoch., 16 (2012), pp. 319–334.
- [27] L. STENTOFT, *Convergence of the least squares Monte Carlo approach to American option valuation*, Management Sci., 50 (2001), pp. 576–611.
- [28] J. N. TSITSIKLIS AND B. VAN ROY, *Regression methods for pricing complex American-style options*, IEEE Trans. Neural Networks, 12 (2001), pp. 694–703.
- [29] J. N. TSITSIKLIS AND B. VAN ROY, *Optimal stopping of Markov processes: Hilbert space, theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives*, IEEE Trans. Automat. Control, 44 (1999), pp. 1840–1851.