

On the valuation of natural resource investments using optimal stochastic switching

J. Hinz ^a, T. Tarnopolskaya ^b and J. Yee ^c

^a*School of Maths, UTS*

^b*CSIRO, Sydney*

^c*School of Maths, UTS and CSIRO*

Email: jeremyee@outlook.com.au

Abstract: Complexity and technical difficulties associated with mineral extractions together with significant uncertainty surrounding mining operations require sophisticated stochastic control algorithms suitable for high-dimensional problems to optimise operating strategies under uncertainty.

Minerals extraction problems have attracted significant interest from researchers in the fields of stochastic control and real options since the 1980s. Numerous approaches have been suggested in the literature to treat such problems. Over the last decade, the least-squares Monte Carlo approach, originally motivated by financial applications, that approximates value function by linear combinations of a set of basic functions, has become very popular in the stochastic control problems in minerals applications. While such approach reduces computational complexity of the problems, it has two significant drawbacks. Specifically, (1) a choice of basis function is problematic; and (2) increasing the size of the basis may cause oscillations if the sample size is small.

Recently, a new algorithm known as Convex Stochastic Switching (CSS) has been suggested. The algorithm approximates a convex value function by piecewise affine linear convex functions in terms of sub-gradients. The linear dynamics of the state space and the Bellman recursion are utilised to estimate the value functions through operations on the sub-gradients. These value function approximations converge to the true values uniformly on compact sets with probability one. The method is completely free from the drawbacks of the least-squares Monte Carlo. This paper presents the first application of the new CSS algorithm to natural resource investment problem.

Keywords: *Convex switching systems, natural resource investments, real options*

1 INTRODUCTION

Optimal decision making in the area of natural resource valuation can often be formulated as an optimal switching problem. The literature devoted to this topic is already immense. A vast variety of methods have been suggested to address decision making for high-dimensional state spaces, most of them dealing with *state space discretization*, or with appropriate *function approximations* on the state space. Within the family of *state space discretization* methods, the *optimal quantization approach*¹ (see Pags et al. (2004)) has been successfully applied to solve optimal switching and stopping problems with applications to finance (see Bronstein et al. (2010)). In the framework of optimal quantization, an appropriate approximation of the controlled Markov process is achieved by a finite-state controlled Markov chain. This technique is based on an optimal choice of discrete distributions, approaching the continuous-space transition probabilities. The resulting discrete-space control problems are solved by traditional methods and provides nearly optimal solutions to the original problems. This approach is justified by its reliable asymptotic properties. Let us now focus on *function approximation* methods which are directly related to our approach. Within the family of these methods, the *least-squares Monte Carlo* technique represents a traditional way to approximate the value function by linear combinations of basis functions. Over the last decade, this approach has attracted remarkable attention. Following earlier works Carriere (1996); Tsitsiklis and Van Roy (2001); Tsitsiklis and Roy (1999), the contribution of Longstaff and Schwartz Longstaff and Schwartz (2001) became a source of extensive research, which focused on the theoretical justifications of this method. Convergence issues have been addressed in Clement et al. (2002) and later generalized in Stentoft (2004); Egloff (2005) and Egloff et al. (2007), extensions to multiple exercise rights were considered in Carmona and Touzi (2008), and recently studied in Belomestny and Schoenmakers (2010) where also connections to statistical learning theory and the theory of empirical processes was emphasized. Furthermore, advanced regression methods such as kernel methods Ormoneit and Glynn (2002a,b), local polynomial regression Fan and Gijbels (1996), and neural networks Bertsekas and Tsitsiklis (1996), have been suggested. However, most of the classical function approximation methods do not use specific structure of the value functions, unlike our approach which is based on certain convexity assumptions, frequently satisfied in practical applications.

This paper will use the method proposed by Hinz (2014) to treat the valuation problem under the framework of the so-called Convex Switching Systems (CSS). The CSS algorithm belongs to the class of value function approximation methods, such as the Least Squares Monte Carlo (LSMC) method popularized by Longstaff and Schwartz (2001). However, the LSMC methodology uses linear combinations of basis functions which makes implementation difficult since a choice of the basis may not be obvious. The growth of the sample size and basis dimension also cannot be chosen independently due to overfitting issues (see Stentoft (2004)). In contrast, the CSS algorithm approximates the value function using operations on subgradients and the growth of the sample size is decoupled from the approximating function space. In contrast to the popular LSMC method, the CSS algorithm ensures uniform convergence of value functions on compact sets, which ensures convergence of the prescribed policies to optimality. Furthermore, the present framework provides an efficient method for numerical solutions. The idea is based on approximations of *convex value functions* by piecewise affine linear functions. Representing such functions in terms of matrices, the entire numerical scheme boils down to a sequence of a few simple matrix transformations acting in an algorithmically appealing, dimension-independent way. Our methodology is directly applicable to high-dimensional problems and shows remarkable numerical efficiency and excellent precision. More importantly, we establish a sound and reliable diagnostics and quality assessment tool for a posterior justification of the numerical approximation (whose details are given in Hinz and Yap (2014)). The authors believe that such combination of efficient numerical schemes with a subsequent diagnostic check can be very useful in practical applications. and intend to publish an R package containing their implementation of CRAN next.

2 CONVEX SWITCHING SYSTEMS

This paper considers the following finite time horizon optimal decision problem. The system state consists of a discrete and a continuous component. More precisely, suppose that the state space $E = P \times \mathbb{R}^d$ is the product of a finite space P and the Euclidean space \mathbb{R}^d . Suppose that the first component $p \in P$ is deterministically driven by a finite set A of actions in terms of a function $\alpha : P \times A \rightarrow P$, $(p, a) \mapsto \alpha(p, a)$, where $\alpha(p, a) \in P$ is the new position if the previous position was p and action $a \in A$ was undertaken by the controller. Further, assume that the *continuous state component* $z \in \mathbb{R}^d$ evolves as a Markov process $(Z_t)_{t=0}^T$ on \mathbb{R}^d realized on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and governed by the transition kernels which satisfy

¹The author is grateful to an anonymous referee for pointing out this issue

$(z, dz') \mapsto \mathbb{P}(W_{t+1}z \in dz')$, $z \in \mathbb{R}^d$, $t = 0, \dots, T-1$, $a \in A$, with the so-called *disturbances* $(W_t)_{t=1}^T$ described by independent integrable $d \times d$ matrix-valued random variables. Intuitively, $(Z_t)_{t=0}^T$ can be thought as governed by the following uncontrolled linear dynamics $Z_{t+1} = W_{t+1}Z_t$ which is started at a given point $z_0 \in \mathbb{R}^d$. Let us now turn to the specification of control costs. Assume that taking an action $a \in A$ causes an immediate reward $r_t(p, z, a)$ which depends on the state $(p, z) \in E$ and on the action $a \in A$ through given reward functions $r_t : E \times A \rightarrow \mathbb{R}$ which may be time $t = 0, \dots, T$ dependent. 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 pre-specified *scrap function* $r_T : E \rightarrow \mathbb{R}$, is collected. The reward functions and scrap values are assumed to be convex and globally Lipschitz continuous in the continuous component z .

At each time $t = 0, \dots, T$ the *decision rule* π_t is given by a mapping $\pi_t : E \rightarrow A$, prescribing at time t an action $\pi_t(p, z) \in A$ for a given state $(p, z) \in E$. A sequence $\pi = (\pi_t)_{t=0}^T$ of decision rules is called *policy*. When controlling the system by the policy $\pi = (\pi_t)_{t=0}^T$, one recursively obtains actions and system states $a_t^\pi := \pi_t(p_t^\pi, Z_t)$, $p_{t+1}^\pi = \alpha(p_t^\pi, a_t^\pi)$, $Z_{t+1} = W_{t+1}Z_t$. The so-called *policy value* $v_0^\pi(p, z)$ is defined for $(p, z) \in E$ as the total expected reward $v_0^\pi(p, z) = \mathbb{E} \left[\sum_{t=0}^{T-1} r_t(p_t^\pi, Z_t, a_t^\pi) + r_T(p_T^\pi, Z_T) \right]$ under the assumptions that $p_0^\pi = p$ and $Z_0 = z$. A policy $\pi^* = (\pi_t^*)_{t=0}^T$ is called *optimal* if it maximizes the total expected reward over all policies $\pi \mapsto v_0^\pi(p, z)$. To obtain such policy, one introduces for $t = 0, \dots, T-1$ the so-called *Bellman operator*

$$\mathcal{T}_t v(p, z) = \max_{a \in A} \{ r_t(p, z, a) + \mathbb{E}(v(\alpha(p, a), W_{t+1}z)) \}, \quad (p, z) \in P \times \mathbb{R}^d \quad (2.1)$$

acting on all functions v where the stochastic kernel is defined. Consider the *Bellman recursion*, also referred to as *backward induction*:

$$v_T(p, z) = r_T(p, z), \quad v_t = \mathcal{T}_t v_{t+1} \quad \text{for } t = T-1, \dots, 0. \quad (2.2)$$

Under additional assumptions (specified by the notion of *Convex Switching System*) there exists a recursive solution $(v_t^*)_{t=0}^T$ to the Bellman recursion. These functions $(v_t^*)_{t=0}^T$ are called *value functions*, they determine an optimal policy $\pi^* = (\pi_t^*)_{t=0}^T$ via

$$\pi_t^*(p, z) = \operatorname{argmax}_{a \in A} \{ r_t(p, z, a) + \mathbb{E}(v_{t+1}^*(\alpha(p, a), W_{t+1}z)) \}, \quad t = T-1, \dots, 0 \quad (2.3)$$

3 CSS ALGORITHM

Let us now present an algorithm from Hinz (2014). The first step in obtaining a numerical solution to the backward induction (2.2) is an appropriate discretization of the Bellman operator (2.1) to

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

with the expectation replaced by its numerical counterpart which is defined in terms of an appropriate distribution sampling $(W_{t+1}(k))_{k=1}^n$ of each disturbance W_{t+1} with corresponding probability weighting $(\nu_{t+1}^n(k))_{k=1}^n$. In the resulting modified backward induction

$$v_T^{(n)} = r_T, \quad v_t^{(n)} = \mathcal{T}_t^n v_{t+1}^{(n)}, \quad t = T-1, \dots, 0 \quad (3.1)$$

the functions $(v_t^{(n)})_{t=0}^T$ need to be described by algorithmically tractable objects. Here we use convexity. Since the reward and scrap functions are convex in the continuous variable, the modified value functions (3.1) are also convex and we can approximate them by piecewise linear and convex functions. For this, we introduce the so-called *sub-gradient envelope* $\mathcal{S}_G f$ of a convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ on a grid $G \subset \mathbb{R}^d$ $\mathcal{S}_G f = \bigvee_{g \in G} (\nabla_g f)$ which is a maximum of the sub-gradients $\nabla_g f$ of f on all grid points $g \in G$. Using the sub-gradient envelope operator, we define the double-modified Bellman operator as

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

where the operator \mathcal{S}_{G^m} stands for the sub-gradient envelope on the grid $G^m = \{g^1, \dots, g^m\}$. The corresponding backward induction

$$v_T^{(m,n)}(p, \cdot) = \mathcal{S}_{G^m} r_T(p, \cdot), \quad p \in P \quad (3.2)$$

$$v_t^{(m,n)}(p, \cdot) = \mathcal{T}_t^{m,n} v_{t+1}^{(m,n)}(p, \cdot), \quad t = T-1, \dots, 0. \quad (3.3)$$

yields the so-called double-modified value functions $(v_t^{(n,m)})_{t=0}^T$ which enjoy excellent asymptotic and algorithmic properties. Namely, under appropriate additional assumptions, the double-modified value functions converge uniformly to the true value functions almost surely on compact sets. These assumptions include the convexity and global Lipschitz continuity of the rewards and scraps, the integrability of all disturbances and some restrictions on the distribution sampling and grid density.

Since the double-modified value functions $(v_t^{(m,n)})_{t=0}^T$ are piecewise linear and convex, they can be expressed in a compact and appealing form using matrix representations. Note that any piecewise convex function f can be described by a matrix where each linear functional is represented by a row in the matrix. To denote this relation, let us agree on the following notation: Given a function f and a matrix F , we write $f \sim F$ whenever $f(z) = \max(Fz)$ holds for all $z \in \mathbb{R}^d$. Let us emphasize that the sub-gradient envelope operation \mathcal{S}_G on a grid G is reflected in terms of a matrix representative by a specific row-rearrangement operator $f \sim F \Leftrightarrow \mathcal{S}_G f \sim \Upsilon_G[F]$ where the row-rearrangement operator Υ_G associated with the grid $G = \{g^1, \dots, g^m\} \subset \mathbb{R}^d$ acts matrix F with d columns as follows:

$$(\Upsilon_G F)_{i,\cdot} = F_{\arg\max(Fg^i),\cdot} \quad \text{for all } i = 1, \dots, m. \quad (3.4)$$

For piecewise convex functions, the result of maximization, summation, and composition with linear mapping, followed by sub-gradient envelope can be obtained using their matrix representatives. More precisely, if $f_1 \sim F_1$, $f_2 \sim F_2$ holds, then it follows that $\mathcal{S}_G(f_1 + f_2) \sim \Upsilon_G(F_1) + \Upsilon_G(F_2)$; $\mathcal{S}_G(f_1 \vee f_2) \sim \Upsilon_G(F_1 \sqcup F_2)$; $\mathcal{S}_G(f_1(W_{t+1}(k)\cdot)) \sim \Upsilon_G(F_1 W_{t+1}(k))$, where the operator \sqcup stands for binding matrices by rows.

Under the assumptions of global Lipschitz continuity and convexity for scraps and reward functions, the backward induction (3.2) and (3.3) can be expressed in terms of the matrix representatives $V_t^{m,n}(p)$ of the value functions $v_t^{(m,n)}(p, \cdot)$ for $p \in P$, $t = 0, \dots, T$. 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. Let us present the resulting algorithm:

Algorithm 1: Convex Switching System Algorithm

input :

- State discretization by a grid $G^{(m)} \subset \mathbb{R}^d$
- Disturbance distribution discretization $\{W_t(k)\}_{k=1}^n$ and their corresponding weights $\{\nu_t^n(k)\}_{k=1}^n$
- Matrix representatives for scrap functions: $R_T^m(p) \sim S_{G^{(m)}} r_t(p, \cdot)$
- Matrix representatives for reward functions: $R_t^m(p, a) \sim S_{G^{(m)}} r_t(p, \cdot, a)$, $p \in P$, $a \in A$

1 **for** $p \in P$ **do**
2 | $V_T^{m,n}(p) \leftarrow \Upsilon_{G^{(m)}} R_T^m(p)$;
3 **end**
4 **for** $t \in \{T-1, \dots, 0\}$ **do**
5 | **for** $p \in P$ **do**
6 | | $\tilde{V}_{t+1}^{m,n}(p) \leftarrow \sum_{k=1}^n \nu_{t+1}^n(k) \Upsilon_{G^{(m)}} (V_{t+1}^{m,n}(p) W_{t+1}(k))$;
7 | **end**
8 | **for** $p \in P$ **do**
9 | | $V_t^{m,n}(p) \leftarrow \Upsilon_{G^{(m)}} \sqcup_{a \in A} (R_t^m(p, a) + \tilde{V}_{t+1}^{m,n}(\alpha(p, a)))$;
10 | **end**
11 **end**

output:

- Matrix representatives of approximate value function: $(V_t^{m,n}(p))_{p \in P, t = 0, \dots, T}$
- Matrix representatives of approximate expected value function: $(\tilde{V}_t^{m,n}(p))_{p \in P, t = 1, \dots, T}$

If the disturbance matrices are identically distributed across time, a significant amount of computational effort can be saved by approximating the conditional expectation operator using a nearest neighbors algorithm, as described in Hinz and Yap (2014). Finally, let us turn to the optimality of the policy prescribed by the above algorithm. Our algorithm returns matrix representatives of approximate value functions and approximated expected value functions, which define a candidate for a nearly optimal strategy by

$$\pi_t^{m,n}(p, z) = \arg\max_{a \in A} \left(\max(R_t^m(p, a)z) + \max(\tilde{V}_{t+1}^{m,n}(\alpha(p, a))z) \right).$$

Given the pointwise convergence

$$\lim_{(m,n) \rightarrow \infty} \left(\max(R_t^m(p, a)z) + \max(\tilde{V}_{t+1}^{m,n}(\alpha(p, a))z) \right) = r_t(p, z, a) + \mathbb{E}(v_{t+1}^*(\alpha(p, z), W_{t+1}z)), \quad (3.5)$$

for all $p \in P$, $a \in A$, and $t = 0, \dots, T - 1$ we deduce that $\lim_{(m,n) \rightarrow \infty} \pi_t^{m,n}(p, z) = \pi^*(p, z)$ $p \in P, a \in A, t = 0, \dots, T - 1$ holds for each point $(p, z) \in E$ where the maximizing action in is uniquely determined in $A \rightarrow \mathbb{R}, a \mapsto r_t(p, z, a) + \mathbb{E}(v_{t+1}^*(\alpha(p, z), W_{t+1}z))$. The conditions for (3.5) to hold are given in Hinz (2014) and are satisfied under the assumptions of Convex Switching Systems. That is, using algorithm (1) with appropriate grid tightening $m \rightarrow \infty$ and disturbance sampling $n \rightarrow \infty$, we obtain a sequence $(\pi^{m,n})_{m,n \in \mathbb{N}}$ of policies which approximates an optimal policy π^* . Since the set A is finite, the convergence in (3.5) implies that algorithm (1) prescribes an optimal policy for a sufficiently dense grid and a sufficiently large disturbance sampling.

3.1 An efficient implementation

The computational performance of our approach suffers from the fact that most of the calculation time is being spent on matrix rearrangements required by the operator Υ . Namely, in order to calculate $\sum_{k=1}^n \nu_{t+1}(k) \Upsilon[V_{t+1}(p') \cdot W_{t+1}(k)]$ the 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. Let us omit the time index $t + 1$ to clarify the idea of efficiency improvement developed in Hinz and Yap (2014). This approach focuses on the two major computational problems: the rearrangement $\Upsilon[VW(k)]$ of large matrices $V \cdot W(k)$, and the summation of matrices $\Upsilon[V \cdot W(k)]$ over a large index range $k = 1, \dots, n$. The crucial point is that one can approximate the procedure in the first problem by replacing the row-rearrangement operation with an appropriate matrix multiplication. More precisely Hinz and Yap (2014) suggests

$$\text{constructing a matrix } Y(k) \text{ such that } Y(k)VW(k) \text{ approximates } \Upsilon[VW(k)]. \quad (3.6)$$

for $k = 1, \dots, n$. The realizations of the approximation (3.6) is technical and relies on efficient implementation of the so-called approximate next-neighbor search and hierarchical clustering, whose details are discussed in Hinz and Yap (2014). Given (3.6), we now have the following approximation:

$$\sum_{k=1}^n \nu(k) \Upsilon[VW(k)] \approx \sum_{k=1}^n \nu(k) Y(k) VW(k) \quad (3.7)$$

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. Assuming that disturbance matrix W is represented as the linear combination $W = \bar{W} + \sum_{j=1}^J \epsilon_j E(j)$ with non-random matrices \bar{W} and $(E(j))_{j=1}^J$, and random coefficients $(\epsilon_j)_{j=1}^J$, we obtain the following interchange of summations on the right-hand side of (3.7):

$$\sum_{k=1}^n \nu(k) Y(k) VW(k) = \left(\sum_{k=1}^n \nu(k) Y(k) \right) V \bar{W} + \sum_{j=1}^J \left(\sum_{k=1}^n \nu(k) \epsilon_j(k) Y(k) \right) V E(j).$$

By pre-computing the matrices in parenthesis one reduces the computations significantly, if $(W_t)_{t=1}^T$ are identically distributed.

4 NATURAL RESOURCES UNDER FLEXIBLE EXTRACTION

Consider a dynamic decision process where an agent aims to optimise the usage of a natural resource asset containing a fixed level of inventory. We assume that at any time $t = 0, \dots, T - 1$ an action $a \in A$ must be applied. The set A contains the following: open, close or abandon the asset. If the asset is opened, it returns a revenue depending on market conditions but the inventory level decreases. If the asset is closed, the inventory is preserved but a maintenance cost must be paid. An abandoned asset yields neither costs nor revenue. The

positions of this system carry the information about present inventory levels as well as the mode of the asset. We assume that an asset is automatically abandoned when the inventory is exhausted. Let us now specify the ingredients of our problem under the framework of Section 2. Assume that the set $P = \{0, 1, \dots, I\} \times \{1, 2\}$ stands for all possible system positions with the interpretation that the first component $p^{(1)}$ of the position $(p^{(1)}, p^{(2)}) \in P$ describes the inventory level; $p^{(1)} = I \in \mathbb{N}$ indicates that the inventory is full whereas $p^{(1)} = 0$ stands for an exhausted or abandoned asset. Further, assume that the second component $p^{(2)}$ of the position $(p^{(1)}, p^{(2)}) \in P$ indicates whether the asset is closed $p^{(2)} = 1$ or opened $p^{(2)} = 2$. Let us turn now to the action set $A = \{0, 1, 2\}$ and its impact on the evolution of the position. We assume that $a = 0, 1, 2$ stand for the actions to abandon, to close, and to open the asset, respectively. Define $\alpha : P \times A \rightarrow P$ in the following manner: abandoning the asset: $\alpha((p^{(1)}, p^{(2)}), 0) = (0, 1)$; closing the asset: $\alpha((p^{(1)}, p^{(2)}), 1) = (p^{(1)}, 1)$; and opening the asset: $\alpha((p^{(1)}, p^{(2)}), 2) = (\max\{p^{(1)} - 1, 0\}, 2)$.

The dynamics of the continuous state component will be described by a sampled geometric Brownian motion $Z_{t+1} = W_{t+1}Z_t$, $t = 0, \dots, T - 1$, $Z_0 \in \mathbb{R}_+$, where $(W_t)_{t=1}^T$ are independent random variables following a log-normal distribution. Finally, we determine the reward functions for our system. If the asset is abandoned or the resource is exhausted, we assume that there is neither costs nor revenue associated with asset $r_t((0, p^{(2)}), z, a) = 0$, $a \in A, z \in \mathbb{R}, t = 0, \dots, T - 1$. For the case where $p^{(1)} > 0$, we have $r_t((p^{(1)}, p^{(2)}), z, a) = h_t(z)1_{\{2\}}(a) + m_t1_{\{1\}}(a) + c_t|p^{(2)} - a|, t = 0, \dots, T - 1$ with the following interpretation: a revenue is collected which depends on z through a pre-specified convex function h_t if we decide to open the mine; a maintenance fee m_t is to be paid if the mine is chosen to be closed down; a switching fee c_t is incurred whenever the operational mode transitions from opened to closed or vice versa.

We now consider the problem studied by Brennan and Schwartz (1985). We assume that the decision maker does not sell to the buyer at a fixed price but instead trades directly in the spot commodity market. We assume that the cost of extraction is deterministic. Let denote by $\Delta > 0$ the physical time duration between subsequent decision times which are indexed by integers $\{0, 1, \dots, T\} \in \mathbb{N}$. Furthermore, assume that the inventory levels are equidistantly sampled with step size corresponding to the maximal extraction within one time step and are indexed by integers $\{0, \dots, I\} \subset \mathbb{N}$. To account for the time value, we introduce the variables $\rho > 0, \theta > 0$ and $\kappa > 0$ which stand for the interest rate property tax rate, and inflation rate, respectively. With these assumptions, the revenue from extraction at time $t \in \{0, \dots, T - 1\}$ is given by $h_t(z) = z \cdot H_1 \Delta e^{(-\rho-\theta)t\Delta} - H_2 \Delta e^{(-\rho-\theta+\kappa)t\Delta}$, $z \in \mathbb{R}$. Thereby, the gross revenue is represented by $z \cdot H_1 \Delta e^{(-\rho-\theta)t\Delta}$ where z stands for the resource price, $H_1 \Delta$ describes the extracted volume and the factor $e^{(-\rho-\theta)t\Delta}$ stands for discounting. Similarly, the expression $H_2 \Delta e^{(-\rho-\theta+\kappa)t\Delta}$ accounts for operating costs incurred from time t to $t + 1$ which are proportional to time duration Δ and are also adjusted due to discounting effects. Similarly, the maintenance and switching costs are given by $m_t = m \Delta e^{(-\rho-\theta+\kappa)t\Delta}$ and $c_t = c e^{(-\rho-\theta+\kappa)t\Delta}$ respectively for $t = 0, 1, \dots, T - 1$. We assume that the scrap value is zero and introduce the price of the resource to follow

$$Z_{t+1} = \exp\left(\left[\rho - y + \frac{\sigma^2}{2}\right] \Delta + \sigma\sqrt{\Delta}N_{t+1}\right) Z_t, \quad t = 0, \dots, T - 1 \tag{4.1}$$

with independent standard normally distributed random variables $(N_t)_{t=1}^{T-1}$ and $\rho = 0.1$ representing the interest rate, $y = 0.01$ the convenience yield, and $\sigma^2 = 0.08$ the volatility. For our numerical illustration, we assume $T = 30, \Delta = 0.25, \bar{I} = 15, H_1 = 5, H_2 = 2.5, \kappa = 0.08, c = 0.2$ and $\theta = 0.02$. The Table 4.1 shows

Price	B&S		CSS		Out Of Sample	
	Open	Closed	Open	Closed	Open	Closed
0.3	1.25	1.45	1.198	1.398	1.270(.042)	1.477(.042)
0.4	4.15	4.35	4.082	4.282	4.123(.057)	4.325(.057)
0.5	7.95	8.11	7.875	8.046	7.964(.067)	8.139(.072)
0.6	12.52	12.49	12.480	12.450	12.394(.076)	12.341(.081)
0.7	17.56	17.38	17.560	17.360	17.660(.089)	17.453(.089)
0.8	22.88	22.68	22.910	22.710	22.938(.101)	22.738(.101)
0.9	28.38	28.18	28.460	28.260	28.463(.114)	28.261(.114)
1.0	34.01	33.81	34.130	33.930	34.305(.131)	34.104(.131)

Table 4.1. The results from Brennan and Schwartz (1985) are listed in the second and third columns. The CSS estimates were computed using an equally spaced grid from 0 to 20 using 4,001 grid points and 10,000 disturbances. The out of sample values were computed using 100,000 anti-thetic paths.

that our algorithm (CSS) attains similar results to Brennan and Schwartz (1985) although we have assumed a lower operational flexibility (i.e. only able to switch four times per year) in difference to the continuous-time switching model from Brennan and Schwartz (1985) whose results are obtained from a numerical solution of a system of partial differential equations.

5 CONCLUSION

In this paper, we have shown that the CSS algorithm makes a suitable companion to LSMC and the classical PDE approach for the valuation of natural resource assets. The addition of another powerful tool to a practitioner's toolbox is always welcomed.

REFERENCES

- Belomestny, N., Kolodko, A., and J. Schoenmakers (2010). Regression methods for stochastic control problems and their convergence analysis. *SIAM J. Control Optim.* 48(5), 3562–3588.
- Bertsekas, D. P., and J. N. Tsitsiklis (1996). *Neuro-Dynamic Programming*. Athena Scientific.
- Brennan, M., and E. Schwartz (1985). Evaluating natural resource investments. *The Journal of Business* 58(2), 135–57.
- Bronstein, A. L., G. Pagès, and B. Wilbertz (2010). How to speed up the quantization tree algorithm with an application to swing options. *Quantitative Finance* 10(9), 995–1007.
- Carmona, R., and N. Touzi (2008). Optimal multiple stopping and valuation of swing options. *Mathematical Finance* 18, 239 – 268.
- Carriere, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19, 19–30.
- Clement, E., D. Lamberton, and P. Protter (2002). An analysis of the Longstaff-Schwartz algorithm for American option pricing. *Finance and Stochastics* 6(4), 449–471.
- Egloff, D. (2005). Monte Carlo algorithms for optimal stopping and statistical learning. *Appl. Probab.* 15, 1396–1432.
- Egloff, D., M. Kohler, and N. Todorovic (2007). A dynamic look-ahead Monte Carlo algorithm. *Appl. Appl. Probab.* 17, 1138–1171.
- Fan, J., and I. Gijbels (1996). *Local Polynomial Modelling and Its Applications*. Chapman and Hall.
- Hinz, J. (2014). Optimal stochastic switching under convexity assumptions. *SIAM Journal on Control and Optimization* 52(1), 164–188.
- Hinz, J., and N. Yap (2014). Algorithms for optimal control of stochastic switching systems. *Preprint*.
- Longstaff, F., and E. Schwartz (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* (14), 113–147.
- Ormoneit, D., and P. Glynn (2002a). Kernel-based reinforcement learning. *Machine Learning* 49, 161–178.
- Ormoneit, D., and P. Glynn (2002b). Kernel-based reinforcement learning in average-cost problems. *IEEE Transactions in Automatic Control* 47, 1624–1636.
- Pagès, G., H. Pham, and J. Printems (2004). Optimal quantization methods and applications to numerical problems in finance. *Handbook of computational and numerical methods in finance*, 253–297.
- Stentoft, L. (2004). Convergence of the least squares Monte Carlo approach to American option valuation. *Management Science* 50(9), 576–611.
- Tsitsiklis, J., and B. V. Roy (1999). Optimal stopping of Markov processes: Hilbert space, theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control* 44(10), 1840–1851.
- Tsitsiklis, J. N., and B. Van Roy (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12(4), 694–703.