European Journal of Operational Research 256 (2017) 196–204

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

# Comparison of least squares Monte Carlo methods with applications to energy real options

Selvaprabu Nadarajah[a,*], François Margot[b], Nicola Secomandi[b]

[a] College of Business Administration, University of Illinois at Chicago, 601 South Morgan Street, Chicago, IL 60607, USA
[b] Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890, USA

ARTICLE INFO

*Article history:*
Received 15 October 2014
Accepted 9 June 2016
Available online 23 June 2016

*Keywords:*
Energy
Real options
Least-squares Monte Carlo
Approximate dynamic programming
Information relaxation and duality

ABSTRACT

Least squares Monte Carlo (LSM) is a state-of-the-art approximate dynamic programming approach used in financial engineering and real options to value and manage options with early or multiple exercise opportunities. It is also applicable to capacity investment and inventory/production management problems with demand/supply forecast updates arising in operations and hydropower-reservoir management. LSM has two variants, referred to as regress-now/later (LSMN/L), which compute continuation/value function approximations (C/VFAs). We provide novel numerical evidence for the relative performance of these methods applied to energy swing and storage options, two typical real options, using a common price evolution model. LSMN/L estimate C/VFAs that yield equally accurate (near optimal) and precise lower and dual (upper) bounds on the value of these real options. Estimating the LSMN/L C/VFAs and their associated lower bounds takes similar computational effort. In contrast, the estimation of a dual bound using the LSML VFA instead of the LSMN CFA takes seconds rather than minutes or hours. This finding suggests the use of LSML in lieu of LSMN when estimating dual bounds on the value of early or multiple exercise options, as well as of related capacity investment and inventory/production policies.



## 1. Introduction

The valuation and management of options with early or multiple exercise opportunities is a fundamental problem in financial engineering and real options (Detemple, 2006; Dixit & Pindyck, 1994; Glasserman, 2004; Guthrie, 2009; Shreve, 2004; Trigeorgis, 1996). A variety of standard and customized stock, interest rate, commodity, energy, and weather options are traded on both organized exchanges and over-the-counter markets (Hull, 2014). When these options give their holders the ability to exercise them one or more times before expiration, the optimization of an exercise policy in the presence of market uncertainty is a key aspect of the valuation and management of these options (Detemple, 2006; Glasserman, 2004; Shreve, 2004). Real options are models of operational flexibility embedded in managerial activities performed in the face of market or operational uncertainty (Dixit & Pindyck, 1994; Guthrie, 2009; Trigeorgis, 1996). Typical applications are the timing of capacity/technology/product investment or divestment decisions, and the switching among inputs or outputs or production modes of manufacturing processes. Common sources of uncertainty include the market value of a completed development project or the levels of input and output prices. The optimization of capacity investment and inventory/production management policies under supply or demand forecast uncertainty (Graves, Meal, Dasu, & Qiu, 1986; Heath & Jackson, 1994; Iida & Zipkin, 2006), possibly combined with market uncertainty (Goel & Gutierrez, 2011; Kouvelis, Chambers, & Wong, 2006), is a critical concern in both operations management (Porteus, 2002; Zipkin, 2000) and hydropower-reservoir management (Nandalal & Bogardi, 2007; Zhao, Cao, & Yang, 2011; Zhao, Zhao, Yang, & Wang, 2013). The modeling of these operational problems shares salient features with the modeling of options with early or multiple exercise opportunities.

Our focus in this paper is on real options, in particular energy real options. Applications include process innovations (Khansa & Liginlal, 2009), manufacturing flexibility (Fontes, 2008; Triantis & Hodder, 1990), capital budgeting (Gamba, 2003), renewable energy investments (Boomsma, Meade, & Fleten, 2012), and commodity and energy acquisition, disposal, processing, production,

* Corresponding author. Tel.: +1 312 355 2774, +1 4126921061.
*E-mail addresses:* selvan@uic.edu (S. Nadarajah), fmargot@andrew.cmu.edu (F. Margot), ns7@andrew.cmu.edu (N. Secomandi).

http://dx.doi.org/10.1016/j.ejor.2016.06.020
0377-2217/

storage, and swing assets (Adkins and Paxson, 2011; Barbieri and Garman, 1996; Boogert and De Jong, 2008; 2011/12; Brandão, Dyer, and Hahn, 2005; Carmona and Ludkovski, 2010; Chandramouli and Haugh, 2012; Cortazar, Gravet, and Urzua, 2008; Devalkar, Anupindi, and Sinha, 2011; Enders, Scheller-Wolf, and Secomandi, 2010; Felix and Weber, 2012; Hahn and Dyer, 2008; 2011; Jaillet, Ronn, and Tompaidis, 2004; Lai, Margot, and Secomandi, 2010; Maragos, 2002; Secomandi, 2010; Smith, 2005; Smith and McCardle, 1998; 1999; Thompson, 2012; 2013; Thompson, Davison, and Rasmussen, 2004; Wang and Dyer, 2010, Arvesen, Medbø, Fleten, Tomasgard, and Westgaard, 2013; Denault, Simonato, and Stentoft, 2013; Mazières and Boogert, 2013; Wu, Wang, and Qin, 2012, Secomandi and Seppi, 2014, Chapters 5–7, Bäuerle & Riess, 2016; Gyurkó, Hambly, & Witte, 2015; Nadarajah, Margot, & Secomandi, 2015).

The modeling of early and multiple exercise options and related operational problems generally gives rise to intractable Markov decision problems (MDPs) with states containing both endogenous and exogenous components. The endogenous part of the MDP state represents the status of the option or the operational system. It is determined by exercise or operational decisions and is low dimensional in several of the discussed applications. The exogenous part of the MDP state is a term structure, such as a commodity or energy forward (futures) curve, a yield curve, or a demand/supply forecast curve. The stochastic dynamics of this term structure are assumed unaffected by option-exercise or operational decisions, and are represented using high dimensional models that share a common mathematical structure (Blanco, Soronow, & Stefiszyn, 2002; Clewlow & Strickland, 2000; Cortazar & Schwartz, 1994; Graves et al., 1986; Heath & Jackson, 1994; Ho & Lee, 1986; Veronesi, 2010). The MDP intractability is thus typically due to two curses of dimensionality: (i) The high dimensionality of the exogenous part of the state space and (ii) the inability to exactly compute expectations of future exogenous state components (Powell, 2011, Section 4.1).

The least squares Monte Carlo (LSM) approach, pioneered by Carriere (1996), Longstaff and Schwartz (2001), and Tsitsiklis and Van Roy (2001), is a prominent approximate dynamic programming (ADP) methodology (Powell, 2011, p. 307) for the valuation and management of early and multiple exercise options (Arvesen et al., 2013; Bäuerle and Riess, 2016; Boogert and De Jong, 2008; 2011/12; Boomsma et al., 2012; Carmona & Ludkovski, 2010; Carriere, 1996; Cortazar et al., 2008; Denault et al., 2013; Desai, Farias, & Moallemi, 2012; Gyurkó et al., 2015; Longstaff & Schwartz, 2001; Smith, 2005; Tsitsiklis & Van Roy, 2001). Similar techniques have been developed for inventory and hydropower-reservoir management problems (Iida & Zipkin, 2006; Wang, Atasu, & Kurtuluş, 2012; Zhao et al., 2011; Zhao et al., 2013), and can be applied to capacity investment and production management settings.

The standard LSM method, known as regress-now LSM (LSMN), addresses the two stated curses of dimensionality by approximating via a linear combination of basis functions the *continuation* function of the MDP formulated as a stochastic dynamic program (SDP). The weights of the basis functions are fitted through least-squares regression on Monte Carlo samples of the exogenous part of the state in a backward recursive fashion. Although convenient for lower bound estimation based on its continuation function approximation (CFA), this method requires executing potentially time consuming nested simulations and optimization when estimating a dual (upper) bound (Glasserman, 2004, Section 8.7, Brown, Smith, & Sun, 2010). A nonstandard LSM variant, proposed by Glasserman and Yu (2004) and known as regress-later LSM (LSML), uses a linear combination of basis functions to approximate the SDP *value* function rather than its continuation function. In this case specifying a value function approximation (VFA) by choosing basis functions of which expectations can be computed in essentially closed form allows avoiding the nested simulations and optimizations that must be performed when estimating a dual bound based on a CFA. Such basis functions include polynomials of term structure elements and prices of call and put options on such elements (Boogert and De Jong, 2008; 2011/12; Boomsma et al., 2012; Broadie and Cao, 2008; Cortazar et al., 2008; Desai et al., 2012; Longstaff and Schwartz, 2001; Tsitsiklis and Van Roy, 2001, Gyurkó et al., 2015).

Despite its appeal, LSML has gone largely unnoticed in the literature. Broadie and Cao (2008) exemplify in passing its application to estimate lower bounds on the prices of multiple exercise (specifically Bermudan max) options. We are not aware of research that compares the performance of LSML and LSMN, or even uses LSML, to value and manage real options.

In this paper we compare LSML and LSMN applied to energy swing and storage options modeled using a common futures term structure model. We use realistic instances to numerically contrast the performance of LSMN/L when obtaining C/VFAs, the computational effort required to estimate lower and dual bounds based on these C/VFAs, and the quality of these bounds. The LSMN/L C/VFAs lead to similarly accurate (near optimal) and precise lower and dual bound estimates. LSMN/L exert comparable effort to obtain their respective C/VFAs and estimate their associated lower bounds. In contrast, estimating the dual bounds using the LSML VFAs instead of the LSMN CFAs takes seconds rather than minutes or hours. This difference is attributable to the nested simulations and the optimizations that must be carried out when using the LSMN CFA to estimate these bounds, but are instead avoided when employing the LSML VFA.

Our findings suggest the use of LSML rather than LSMN when estimating a dual bound on the value of energy swing and storage options. In particular, it may be useful to include the estimation of LSML-based dual bounds as a feature in commercial software packages that use LSMN to obtain lower bounds on the value of operating policies for these options (see, e.g., EnergyQuants, 2016; KYOS, 2013; Matlab, 2015). Potentially, the relevance of our research extends to other options with early or multiple exercise opportunities and capacity investment and inventory/production management models with demand/supply forecast updates.

In Section 2 we formulate our MDP, apply it to energy swing and storage options, and discuss the two curses of dimensionality that arise when attempting to solve this MDP. In Section 3 we present LSMN and LSML and conceptually contrast these methods. In Section 4 we explain how to use their C/VFAs to estimate lower and dual bounds on a real option value. We conduct our numerical study in Section 5. We conclude in Section 6. An online supplement includes Supporting material.

## 2. MDP, energy applications, and curses of dimensionality

We describe our MDP in Section 2.1. We apply this model to energy swing and storage options in Section 2.2. In Section 2.3 we discuss the two curses of dimensionality that typically make solving this MDP intractable.

### 2.1. MDP

There are $N$ exercise dates, each denoted as $T_i$, $i \in \mathcal{I} := \{0, \ldots, N-1\}$. The set $\mathcal{I}$ is the stage set. The state of our MDP at stage $i$ is partitioned into *endogenous* and *exogenous* components. The endogenous component is the scalar $x_i$. It belongs to the finite set $\mathcal{X}_i$ that represents information about the number of remaining exercise rights at stage $i$. The *exogenous* component is the vector $F_i \in \mathbb{R}^{N-i}$ that represents the option underlying term structure $(F_{i,i}, F_{i,i+1}, \ldots, F_{i,N-1})$, where each $F_{i,j}$, $j \geq i$, is the element of the term structure associated with date $T_j$ at time $T_i$. We define

$F_N := 0$. In commodity and energy applications, $F_i$ is a forward curve, $F_{i,i}$ is the time $T_i$ spot price, and $F_{i,j}$ is the date $T_i$ futures price with maturity at time $T_j$.

At stage $i$ and state $(x_i, F_i)$, the decision maker chooses an exercise action $a$ from the finite set $\mathcal{A}_i(x_i)$ and receives the reward $r_i(a, F_i) : \mathcal{A}_i(x_i) \times \mathbb{R}^{N-i} \to \mathbb{R}$. Subsequently, the transition function $f_i(x_i, a)$ specifies the stage $i+1$ endogenous part of the state $f_i(x_i, a)$. The exogenous part of the state evolves from $F_i$ to $F_{i+1}$ according to a known stochastic process independently of $x_i$ and $a$.

A policy $\pi$ is the collection of decision functions $\{A_0^\pi, \ldots, A_{N-1}^\pi\}$, where $A_i^\pi(x_i, F_i) : \mathcal{X}_i \times \mathbb{R}^{N-i} \to \mathcal{A}_i(x_i)$ for each $i \in \mathcal{I}$. We let $\Pi$ be the set of all feasible policies. Let $\mathbb{E}$ denote expectation under the risk-neutral probability measure (Shreve, 2004) for the term structure stochastic process. We denote by $\delta \in (0, 1]$ the risk-free discount factor from each time $T_i$ back to time $T_{i-1}$, $i \in \mathcal{I} \setminus \{0\}$. That is, the discount factor is constant across stages, an assumption that can be relaxed in a straightforward manner. Let $(x_0, F_0)$ be the time $T_0$ state. Computing the real option value and an optimal exercise policy entails solving

$$\max_{\pi \in \Pi} \sum_{i \in \mathcal{I}} \delta^i \mathbb{E}\Big[r_i(A_i^\pi(x_i^\pi, F_i), F_i) | x_0, F_0\Big], \tag{1}$$

where $x_i^\pi$ is the random endogenous part of the state at stage $i$ when using policy $\pi$.

To simplify our notation, for the most part in the rest of this paper we omit the sets that index a tuple. For example, we write $(i, x_i, F_i, a)$ in lieu of $(i, x_i, F_i, a) \in \mathcal{I} \times \mathcal{X}_i \times \mathbb{R}^{N-i} \times \mathcal{A}_i(x_i)$. In addition, we write $(\cdot)_{-(i)}$ to indicate that $i$ is excluded from $\mathcal{I}$ in the tuple ground set.

### 2.2. Energy applications

We present the two applications that are the focus of our numerical study in Section 5.

#### 2.2.1. Energy swing option

Swing options are common in the energy industry (Barbieri & Garman, 1996; Jaillet et al., 2004). We focus on a purchase swing option. This option could be used, for example, by a producer of ethylene that requires an amount $q_i$ of crude oil as input to a thermal cracking process at each time $T_i$. The contract has two parts: A *purchase* part that involves buying the quantity $q_i$ at price $\mathsf{K}_i$ on date $T_i$; and a *swing* part that endows the producer with $n \leq N$ *swing rights* to increase or decrease each purchase amount $q_i$ by a fixed constant $Q_i \in (0, q_i]$ at price $\mathsf{K}_i$ at each stage $i$. At most one swing right can be exercised at a given stage $i$.

The incentive to exercise this swing option at stage $i$ stems from the producer's ability to transact in the spot market at the prevailing spot price $F_{i,i}$. If $\mathsf{K}_i > F_{i,i}$ the producer has the incentive to acquire an amount $q_i - Q_i$ from the swing contract at price $\mathsf{K}_i$ and purchase a quantity $Q_i$ from the spot market at price $F_{i,i}$. This combined trade results in a gain of $Q_i(\mathsf{K}_i - F_{i,i})$ relative to procuring $q_i$ at price $\mathsf{K}_i$. Similarly, if $\mathsf{K}_i < F_{i,i}$, the producer has the incentive to buy an amount $q_i + Q_i$ from the swing contract at price $\mathsf{K}_i$ and sell a quantity $Q_i$ into the spot market at price $F_{i,i}$, for a gain of $Q_i(F_{i,i} - \mathsf{K}_i)$.

The value of the purchase part of this contract is $-\sum_i \delta^i F_{0,i} q_i$. The valuation of the swing part of this contract can be modeled using our MDP by defining the endogenous state variable $x_i$ to be the number of swing rights available at stage $i$. The set $\mathcal{X}_i$ is thus $\{\max\{0, n-i\}, \ldots, n\}$. The feasible action set $\mathcal{A}_i(x_i)$ is $\{0, 1\}$ if $x_i > 0$, and $\{0\}$ if $x_i = 0$. The endogenous state transition function $f_i(x_i, a)$ equals $x_i - a$, because the number of exercise rights remains the same if the no-exercise action is taken, and decreases by one if the exercise action is chosen. The stage $i$ reward function $r_i(a, F_i)$ is $Q_i \cdot |\mathsf{K}_i - F_{i,i}| \cdot a$.

#### 2.2.2. Energy storage option

Consider a finite-term lease contract on a portion of the space and capacity of an energy (e.g., natural gas) storage facility (see Maragos, 2002 for details). At each of a given number of dates, the contract owner can buy energy from the wholesale spot market and inject it into this facility, withdraw previously purchased and injected energy and sell it into the wholesale spot market, or do nothing.

The endogenous state $x_i$ is the inventory in storage at stage $i$. The maximum amount of inventory in storage is denoted by the scalar $\bar{x}$. At each stage $i$, the feasible inventory interval is $[0, \bar{x}]$. At stage $i$, the action $a$ indicates the inventory change between times $T_i$ and $T_{i+1}$. A positive action is an energy withdrawal-and-sell decision, a negative action is an energy purchase-and-inject decision, and zero is the do-nothing decision. The inventory transition function $f_i(x_i, a)$ is $x_i - a$. The storage contract withdrawal and injection capacities are the scalars $\underline{a}$ and $\bar{a}$, which satisfy $0 \leq \underline{a}$, and $\bar{a} \leq \bar{x}$. The sets of feasible injections, withdrawals, and overall actions are $[\max\{-\bar{a}, (x_i - \bar{x})\}, 0]$, $[0, \min\{x_i, \underline{a}\}]$, and the union of the latter two sets, respectively. These feasible inventory and action sets are intervals. By Lemma 1 in Secomandi, Lai, Margot, Scheller-Wolf, and Seppi (2015), they can be optimally discretized if $\underline{a}$, $\bar{a}$, and $\bar{x}$ are integer multiples of some common scalar, which we assume to be the case in this paper. The sets $\mathcal{X}_i$ and $\mathcal{A}_i(x_i)$ are thus defined accordingly (see, Secomandi et al., 2015). Let the coefficients $\alpha^{\mathrm{W}} \in (0, 1]$ and $\alpha^{\mathrm{I}} \geq 1$ model energy losses associated with energy withdrawals and injections, respectively, and the coefficients $\varsigma^{\mathrm{W}}$ and $\varsigma^{\mathrm{I}}$ represent withdrawal and injection marginal costs, respectively. The immediate reward function $r_i(a, F_i)$ is $(\alpha^{\mathrm{I}} F_{i,i} + \varsigma^{\mathrm{I}})a$ if $a \in \mathbb{R}_-$, and $(\alpha^{\mathrm{W}} F_{i,i} - \varsigma^{\mathrm{W}})a$ if $a \in \mathbb{R}_+$ (it thus equals zero if $a = 0$).

### 2.3. Curses of dimensionality

In theory, an optimal policy to the MDP (1) can be obtained by stochastic dynamic programming. We formulate two SDPs for this purpose: The *value function* SDP and the *continuation function* SDP. The LSM methods discussed in Section 3 approximate these SDPs.

The value function SDP for each $(i, x_i, F_i)$ is

$$V_i(x_i, F_i) = \max_{a \in \mathcal{A}_i(x_i)} r_i(a, F_i) + \delta \mathbb{E}[V_{i+1}(f_i(x_i, a), F_{i+1}) | F_i], \tag{2}$$

with boundary conditions $V_N(x_N, F_N) := 0$ for each $x_N \in \mathcal{X}_N$, and $V_i(x_i, F_i)$ the optimal value function at stage $i$ and state $(x_i, F_i)$. The real option value is $V_0(x_0, F_0)$.

The continuation function is defined as

$$C_i(x_{i+1}, F_i) := \delta \mathbb{E}[V_{i+1}(x_{i+1}, F_{i+1}) | F_i] \tag{3}$$

for each $(i, x_{i+1}, F_i)$ with $i < N-1$, and $C_{N-1}(x_N, F_{N-1}) := 0$ for each $(x_N, F_{N-1})$. The continuation function SDP for each $(i, x_{i+1}, F_i)$ is

$$C_i(x_{i+1}, F_i) = \delta \mathbb{E}\Big[\max_{a \in \mathcal{A}_{i+1}(x_{i+1})} r_{i+1}(a, F_{i+1}) + C_{i+1}(f_{i+1}(x_{i+1}, a), F_{i+1}) | F_i\Big]. \tag{4}$$

In this case, the real option value $V_0(x_0, F_0)$ is $\max_{a \in \mathcal{A}_0(x_0)} r_0(a, F_0) + C_0(f_0(x_0, a), F_0)$. The SDP (4) can be derived by replacing $V_{i+1}(x_{i+1}, F_{i+1})$ in (3) with the right hand side of (2) expressed for $i+1$.

Solving the SDPs (2) and (4) is typically intractable due to two curses of dimensionality: (i) The high dimensionality of the exogenous components (term structures) of their states; (ii) the inability to evaluate exactly the expectations in these SDPs when using common term structure models (see Section 3).

## 3. LSMN and LSML

LSMN and LSML, presented in Sections 3.1 and 3.2, address the curses of dimensionality discussed in Section 2.3 by computing

CFAs and VFAs, respectively, which can be used to estimate lower and dual bounds and obtain an exercise policy in the manner discussed in Section 4. In Section 3.3 we briefly contrast these methods.

### 3.1. LSMN

We define a stage $i$ CFA for each endogenous part of the stage $i+1$ state as a linear combination of a number $B_i$ of basis functions of the exogenous part of the stage $i$ state. Such CFAs are commonly used in the ADP literature (e.g., see Glasserman, 2004, p. 430; Bertsekas, 2007, Section 6.1.1; Powell, 2011, p. 326). Let $\psi_{i,b}$ denote the $b$-th CFA basis function at stage $i$. The weight associated with this basis function for the stage $i+1$ endogenous state $x_{i+1}$ is $\theta_{i,x_{i+1},b}$. The corresponding CFA is $\sum_b \psi_{i,b}(F_i)\theta_{i,x_{i+1},b}$. Determining this CFA reduces to estimating the vector of weights $\theta_{i,x_{i+1}} := (\theta_{i,x_{i+1},1}, \ldots, \theta_{i,x_{i+1},B_i})$.

**Algorithm 1:** LSMN.

**Inputs:** Number of regression sample paths $P$ and set of basis function vectors $\{\psi_i, i_{-(N-1)}\}$.
**Initialization:** Generate the set of $P$ term structure regression sample paths $\{F_i^p, (i,p)\}$; $\hat{\theta}_{N-1,x_N} := 0$ for each $x_N$.
**For** each $i = N-2$ to 0 **do**: **For** each $x_{i+1}$ **do**:

(i) **For** each $p$ **do**: Compute the CFA estimate

$$c_i(x_{i+1}, F_i^p) := \delta\left[\max_a r_{i+1}(a, F_{i+1}^p) + \sum_b \psi_{i+1,b}(F_{i+1}^p)\hat{\theta}_{i+1,f_{i+1}(x_{i+1},a),b}\right].$$

(ii) Perform a two-norm regression on the CFA estimates in set $\{c_i(x_{i+1}, F_i^p), p\}$ to determine the vector of CFA weights $\hat{\theta}_{i,x_{i+1}}$.

Algorithm 1 summarizes LSMN by extending the presentation of this method by Tsitsiklis and Van Roy (2001) for American options. For each stage $i_{-(N-1)}$, we define the vector of CFA basis functions as $\psi_i := (\psi_{i,1}, \ldots, \psi_{i,B_i})$. LSMN begins by generating $P$ term structure Monte Carlo regression sample paths and initializing the stage $N-1$ weight vector $\hat{\theta}_{N-1,x_N}$ for each $x_N$ to a vector of zeros. At each stage $i$, starting from stage $N-2$ and moving backward to stage 0, and each stage $i+1$ endogenous state $x_{i+1}$ this method performs steps (i) and (ii). In step (i) LSMN computes estimates $c_i(x_{i+1}, F_i^p)$ of the CFA by replacing the stage $i+1$ continuation function $C_{i+1}(f_{i+1}(x_{i+1}, a), F_{i+1})$ in the SDP (4) with the known stage $i+1$ CFA $\sum_b \psi_{i+1,b}(F_{i+1})\theta_{i+1,f_{i+1}(x_{i+1},a),b}$, and approximating the expectation

$$\delta\mathbb{E}\left[\max_a r_{i+1}(a, F_{i+1}) + \sum_b \psi_{i+1,b}(F_{i+1})\hat{\theta}_{i+1,f_{i+1}(x_{i+1},a),b} | F_i^p\right] \quad (5)$$

by the sample average, based only on the $p$-th sample path,

$$\delta\left[\max_a r_{i+1}(a, F_{i+1}^p) + \sum_b \psi_{i+1,b}(F_{i+1}^p)\hat{\theta}_{i+1,f_{i+1}(x_{i+1},a),b}\right]. \quad (6)$$

In step (ii) LSMN performs a two-norm regression on these CFA estimates to determine the vector of CFA weights $\hat{\theta}_{i,x_{i+1}}$.

### 3.2. LSML

We define a VFA for each stage $i$ and state in this stage as a linear combination of $B_i$ basis functions of the exogenous part of this state. Let $\phi_{i,b}$ be the $b$-th VFA basis function at stage $i$. The weight associated with this basis function for the endogenous part of the state $x_i$ is $\beta_{i,x_i,b}$. The VFA at stage $i$ and state $(x_i, F_i)$ is $\sum_b \phi_{i,b}(F_i)\beta_{i,x_i,b}$. Obtaining this VFA entails estimating the vector of weights $\beta_{i,x_i} := (\beta_{i,x_i,1}, \ldots, \beta_{i,x_i,B_i})$.

LSML computes a VFA by approximating the value function SDP (2), which requires evaluating expectations of next stage VFAs, that is, the term $\mathbb{E}[\sum_b \phi_{i+1,b}(F_{i+1})\beta_{i+1,x_{i+1},b}|F_i]$ for each stage $i$ and term structure $F_i$. LSML is appealing if it is possible to choose basis functions, in the context of a given term structure model, to avoid approximating these expectations. Our LSML presentation is based on choosing each basis function $\phi_{i+1,b}(F_{i+1})$ such that the expectation $\mathbb{E}[\phi_{i+1,b}(F_{i+1})|F_i]$ is a function of the term structure $F_i$ that can be computed in essentially closed form, that is,

$$\mathbb{E}[\phi_{i+1,b}(F_{i+1})|F_i] = \overline{\phi}_{i,i+1,b}(F_i) \quad (7)$$

for some known function $\overline{\phi}_{i,i+1,b}(F_i)$. The closed-form condition (7) generalizes the martingale condition (C3) in Glasserman and Yu (2004). We comment on the practicality of (7) after presenting LSML.

**Algorithm 2:** LSML.

**Inputs:** Number of regression sample paths $P$ and set of basis function vectors $\{\phi_i, i_{-(0)}\}$ that satisfy (7).
**Initialization:** Generate the set of $P$ term structure regression sample paths $\{F_i^p, (i,p)\}$; $\hat{\beta}_{N,x_N} := 0$ for each $x_N$.
**For** each $i = N-1$ to 1 **do**:
**For** each $x_i$ **do**:

(i) **For** each $p$ **do**: Compute the VFA estimate

$$v_i(x_i, F_i^p) := \max_a r_i(a, F_i^p) + \delta\sum_b \overline{\phi}_{i,i+1,b}(F_i^p)\hat{\beta}_{i+1,f_i(x_i,a),b}.$$

(ii) Perform a two-norm regression on the VFA estimates in set $\{v_i(x_i, F_i^p), p\}$ to determine the vector of VFA weights $\hat{\beta}_{i,x_i}$.

Algorithm 2 summarizes the LSML steps by modifying the description of this method for American options by Glasserman and Yu (2004) and replacing their martingale condition with our closed-form condition (7). Let $\phi_i := (\phi_{i,1}, \ldots, \phi_{i,B_i})$. The inputs to LSML are the number of sample paths and a set of vectors of VFA basis functions that satisfy (7). LSML starts by generating a set of $P$ term structure Monte Carlo regression sample paths and initializing the stage $N$ weight vector $\hat{\beta}_{N,x_N}$ to zero for each $x_N$. At each stage $i$, starting from stage $N-1$ and moving backward to stage 1, and for each endogenous state $x_i$ it executes steps (i) and (ii). In step (i) LSML computes estimates $v_i(x_i, F_i^p)$ of the stage $i$ VFA for each sample path $p$ by replacing the stage $i+1$ value function $V_{i+1}(f_i(x_i, a), F_{i+1})$ on the right hand side of (2) by the known stage $i+1$ VFA $\sum_b \phi_{i+1,b}(F_{i+1})\beta_{i+1,f_i(x_i,a),b}$ and exploiting (7). In step (ii) LSML performs a two-norm regression on these estimates to determine the vector of VFA weights $\hat{\beta}_{i,x_i}$.

The closed-form condition (7) is satisfied by certain functions of the exogenous state component when using term structure models that are widespread in commodity, energy, fixed income, and operations and hydro-power reservoir management applications both in practice and in the academic literature (Clewlow & Strickland, 2000; Cortazar & Schwartz, 1994; Eydeland & Wolyniec, 2003; Graves et al., 1986; Heath & Jackson, 1994; Ho & Lee, 1986; Maragos, 2002; Nandalal & Bogardi, 2007; Veronesi, 2010; Zhao et al., 2011; Zhao et al., 2013). In these models, which we use in our numerical study discussed in §5, the term structure evolution is governed by a multidimensional diffusion model. In this continuous time setting, we denote by $F(t, T_j)$ the value of the element of the

term structure at time $t \in [T_0, T_j]$ with maturity at time $T_j$, $j \in \mathcal{I}$. Hence, if $t = T_i$, $i \in \mathcal{I}$, and $j > i$, then $F(t, T_j) \equiv F_{i,j}$. Given a fixed number $K$ of stochastic factors, the evolution of $F(t, T_j)$ is governed by the stochastic differential equations

$$\frac{dF(t,T_j)}{F(t,T_j)} = \sum_{k=1}^{K} \sigma_{j,k}(t) dW_k(t), \quad \forall j \in \mathcal{I} \setminus \{0\}, t \in (T_0, T_j], \tag{8}$$

$$dW_k(t) dW_{k'}(t) = 0, \quad \forall k, k' \in \{1, \ldots, K\}, k \neq k', \tag{9}$$

where $\sigma_{j,k}(t)$ is the time $t$ loading coefficient on the Brownian motion $W_k$ for the term structure element $F(t, T_j)$. This model captures seasonality in the changes in the term structure elements because the loading factors are time dependent. Seasonality in the term structure levels is embedded in the initial (time $T_0$) term structure.

Under models (8) and (9), it is possible to compute (sometimes approximate) conditional expectations of certain classes of functions of future term structure elements as essentially closed-form functions of current term structure elements, that is, $\overline{\phi}_{i,i+1,b}(F_i)$ in (7). We provide two examples of such classes of functions (see Haug, 2006 for a catalog):

1. All polynomials of term structure elements. For example, when $i' > i$, we can use the property $\mathbb{E}[F_{i',j}|F_{i,j}] = F_{i,j}$ to compute expectations of functions that are linear in the term structure elements, and the property $\mathbb{E}[F^2_{i',j}|F_{i,j}] = F^2_{i,j} \exp(\sum_{k \in \mathcal{K}} \int_{T_i}^{T_{i'}} \sigma^2_{j,k}(t) dt)$ to compute expectations of quadratic functions of such elements (these properties are easy to verify).
2. Prices of call and put options on the term structure elements: $\mathbb{E}[\max(0, F_{i',j} - \mathsf{K})|F_{i,j}]$ and $\mathbb{E}[\max(0, \mathsf{K} - F_{i',j})|F_{i,j}]$, where $i' > i$ and $\mathsf{K} \in \mathbb{R}$ (see Section 1.1.3 of Haug, 2006 for explicit formulas for these prices).

These two types of basis functions are commonly used to represent a CFA (Boogert & De Jong, 2008; Cortazar et al., 2008; Longstaff & Schwartz, 2001).

### 3.3. Contrast between LSMN and LSML

LSMN and LSML approximate the continuation function SDP (4) and the value function SDP (2), respectively. Basis functions are one layer of approximation in both of these methods. In addition, LSMN embeds one more approximation layer because it obtains CFA estimates in step (i) of Algorithm 1 by replacing the expectation in (5) with the single sample path, and hence potentially high variance, sample average (6). In contrast, LSML has no additional approximation layer when condition (7) holds. Therefore, when using the same basis functions and number of regression sample paths, we intuitively expect the deviations of the LSML VFA from the exact value function to be smaller than the deviations of the LSMN CFA from the exact continuation function. We provide some theoretical support for this intuition in Online Supplement A. In other words, although we anticipate that for a sufficiently large number of regression sample paths LSMN/L yield essentially equivalent C/VFAs in terms of the accuracy and precision of their corresponding bounds (see Section 4), we expect that LSMN will need more such sample paths than LSML to achieve comparable bounds. Because Algorithms 1 and 2 require a similar number of operations, we predict that attaining such bounding performance will be faster with LSML than LSMN.

## 4. Bounding the real option value

In this section we discuss how to use the C/VFAs computed by LSMN/L to estimate lower and dual bounds on a real option value. This material is in part based on Glasserman (2004, Chapter 8), Bertsekas (2007, Section 6.1.1), Brown et al. (2010), and Powell (2011, Chapter 6).

To estimate a lower bound on a real option value one generates a set of $H$ term structure evaluation Monte Carlo sample paths $\{F^h_i, (i, h)\}$, starting from the term structure $F_0$ at time $T_0$, and simulates the policy induced by the LSML/N V/CFA as now discussed. When using the LSML VFA, we compute the action associated with this policy at stage $i$ and state $(x_i, F_i)$ by solving the optimization model

$$\max_a r_i(a, F_i) + \delta \sum_b \overline{\phi}_{i,i+1,b}(F_i) \hat{\beta}_{i+1, f_i(x_i,a), b}, \tag{10}$$

which is analogous to the optimization executed by LSML in its step (i). When using the LSMN CFA, such an action is obtained by solving the optimization model

$$\max_a r_i(a, F_i) + \sum_b \psi_{i,b}(F_i) \hat{\theta}_{i, f_i(x_i,a), b}, \tag{11}$$

which is related to the optimization performed by LSML in its step (i). A lower bound on the real option value $V_0(x_0, F_0)$ is estimated by averaging the sums of the time $T_0$ discounted rewards gained from executing the actions obtained, according to (10) or (11), along each sample path. We expect roughly equal computational effort for estimating a lower bound using the LSML/N V/CFAs for a given number of evaluation sample paths and employing the same basis functions.

The quality of an estimated lower bound can be assessed by estimating a dual bound using the information relaxation and duality framework (see Brown et al., 2010, and references therein). This approach relies on *feasible* dual penalties $p_i(f_i(x_i, a), F_{i+1}, F_i)$ that penalize knowledge of the future information $F_{i+1}$ when performing action $a$ in state $(x_i, F_i)$ at stage $i$: The feasibility requirement is $\mathbb{E}[p_i(f_i(x_i, a), F_{i+1}, F_i)|F_i] \leq 0$ (Brown et al., 2010). Consider the same set of $H$ term structure sample paths used to estimate a lower bound. Feasible dual penalties are used to formulate the following deterministic dynamic program for the $h$-th term structure sample path with stage $i$ value function $U^h_i(\cdot)$:

$$U^h_i(x_i) = \max_a r_i(a, F^h_i) - p_i(f_i(x_i, a), F^h_{i+1}, F^h_i) + \delta U^h_{i+1}(f_i(x_i, a)), \tag{12}$$

for each pair $(i, x_i)$ with boundary conditions $U^h_N(x_N) := 0$ for each $x_N$. A dual bound estimate on the sought real option value is the sample average $\sum_h U^h_0(x_0)/H$.

As shown by Brown et al. (2010), ideal, but unknown, dual penalties that would lead to a tight dual bound are

$$\delta\{V_{i+1}(f_i(x_i, a), F_{i+1}) - \mathbb{E}[V_{i+1}(f_i(x_i, a), F_{i+1})|F_i]\}. \tag{13}$$

Replacing the value functions in (13) with their corresponding VFAs (Brown et al., 2010), using (7), and ignoring the discount factor for simplicity yields the feasible dual penalties

$$\sum_b \phi_{i+1,b}(F_{i+1}) \hat{\beta}_{i+1, f_i(x_i,a), b} - \sum_b \overline{\phi}_{i,i+1,b}(F_i) \hat{\beta}_{i+1, f_i(x_i,a), b}. \tag{14}$$

When using the LSMN CFA a feasible dual penalty analogous to (14) is

$$\max_a r_{i+1}(a, F_{i+1}) + \sum_b \psi_{i+1,b}(F_{i+1}) \hat{\theta}_{i+1, f_{i+1}(x_{i+1},a), b} - \mathbb{E}\left[\max_a r_{i+1}(a, F_{i+1}) + \sum_b \psi_{i+1,b}(F_{i+1}) \hat{\theta}_{i+1, f_{i+1}(x_{i+1},a), b} | F_i\right]. \tag{15}$$

Obtaining the first term in (15) requires solving an optimization model. In general the second term in (15) cannot be computed exactly because of the presence of the maximization inside the expectation. It is standard (Andersen and Broadie, 2004, Glasserman, 2004, Section 8.7, Haugh & Kogan, 2007) to replace this expectation by a sample average approximation, which requires performing an additional inner Monte Carlo simulation and corresponding

optimizations. In contrast, evaluating (14) is free of this simulation and these optimizations. Thus, given the same number of evaluation sample paths and basis functions, we expect that estimating a dual bound based on (15) be more computationally demanding than doing so based on (14).

## 5. Numerical study

In this section we compare both the quality of the lower and dual bounds estimated using the LSMN/L C/VFAs and the associated computational effort on crude oil swing option and natural gas storage option instances. In Section 5.1 we discuss a specific energy futures term structure model and its calibration. We describe our instances in Section 5.2. In Section 5.3 we present the basis functions that we use to specify our C/VFAs. In Section 5.4 we discuss our results.

### 5.1. Forward curve model and its calibration

In our application, the term structure is an energy forward curve. We used models (8) and (9) to represent its dynamics. We choose each function $\sigma_{m,k}(\cdot)$ in this model to be right continuous and piecewise constant on each interval $[T_i, T_{i+1})$ (Blanco et al., 2002; Secomandi et al., 2015). That is, we set $\sigma_{j,k}(t)$ equal to the constant $\sigma_{j,k,i}$ when $t$ belongs to the interval $[T_i, T_{i+1})$. Under this specification, we can equivalently rewrite models (8) and (9) as

$$F(t', T_j) = F(t, T_j) \exp\left[-\frac{1}{2}(t'-t)\sum_{k=1}^{K}\sigma_{j,k,i}^2 + \sqrt{t'-t}\sum_{k=1}^{K}\sigma_{j,k,i}Z_k\right], \tag{16}$$

for each $i \in \mathcal{I}$, $j \in \{i+1, \ldots, N-1\}$, $t \in [T_i, T_{i+1})$, and $t' \in (T_i, T_{i+1}]$ with $t' > t$, where $(Z_k, k = 1, \ldots, K)$ is a vector of $K$ independent standard normal random variables. We generate forward curve sample paths by Monte Carlo simulation based on (16).

We use ten years of NYMEX crude oil and natural gas futures prices, observed from 1997 to 2006, to estimate sample variance-covariance matrices of the daily log futures price returns for each month for both commodities. We perform a principal component analysis of these matrices to estimate the loading coefficients $\sigma_{j,k,i}$ (see Blanco et al., 2002; Secomandi et al., 2015, Section 5 for more details). We choose the number of factors $K$ equal to 3 and 7 for crude oil and natural gas, respectively, as they are the smallest numbers of factors explaining more than 99 percent of the total observed variance in each of our monthly data sets.

### 5.2. Instances

We set the number of stages $I$ equal to twenty four. Each time interval $[T_i, T_{i+1}]$ represents one month. We create four *price* instances for both crude oil and natural gas by including in the time $T_0$ forward curve, $F_0$, the first twenty four elements of the respective forward curve for these energy sources observed on the first trading date of January, April, July, and October 2006, respectively. We take these months as representative of Winter, Spring, Summer, and Fall. Following Lai et al. (2010), we use risk free interest rates equal to 4.87 percent, 4.74 percent, 5.05 percent, and 5.01 percent for the Winter, Spring, Summer, and Fall price instances, respectively.

We generate our swing option instances by adding to each crude oil price-instance the number of swing rights $n$, which we vary between 1 and 10 in increments of 1, and setting each base load capacity $q_i$ equal to 1, each swing capacity $Q_i$ equal to 0.2, and each price $\mathsf{K}_i$ equal to the price at time $T_0$ of the futures with maturity at time $T_i$. We thus obtain forty swing option instances.

Our storage option instances are based on our natural gas price-instances. Their operational parameters are specified following Lai et al. (2010). In particular, we add to each such price-instance a normalized maximum storage space $\bar{x}$ equal to 1, an initial inventory $x_0$ equal to 0, and high, moderate, and low injection and withdrawal capacity pairs as defined in Lai et al. (2010). This process results in twelve natural gas storage instances.

### 5.3. Basis functions

We use CFAs and VFAs based on the same basis functions. For our swing option instances, we use as basis functions polynomials of the forward prices, which are standard in the LSM literature (Boogert & De Jong, 2008; Cortazar et al., 2008; Longstaff & Schwartz, 2001), and call and put options. The use of these option prices as basis functions is based on the observation that the reward function and the *exact* value function when the number of swing rights equals the number of exercise dates ($n = N$) can be modeled using pairs of these prices. The polynomial basis functions are 1, $F_{i,j}$, and $F_{i,j}^2$ for $j \in \{i, \ldots, N-1\}$, and $F_{i,j}F_{i,j'}$ for $j, j' \in \{i, \ldots, \min\{i+4, N-1\}\}$, $j' > j$. The call option and put option price basis functions, respectively, are $\mathbb{E}[\max(0, F_{j,j} - F_{0,j})|F_{i,j}]$ and $\mathbb{E}[\max(0, F_{0,j} - F_{j,j})|F_{i,j}]$ for $j \in \{i, \ldots, N-1\}$. Thus, as discussed in Section 3.2, our chosen basis functions satisfy the closed-form condition (7). For our storage option instances, we use only the polynomials just described as basis functions.

### 5.4. Results

We estimate our lower and dual bounds using $H = 100{,}000$ evaluation sample paths. We present all our estimated bounds as ratios of the dual bounds estimated using the LSML VFA for $P$ equals to 1000, because, as discussed below, they are essentially tight on all our instances; that is, they are essentially equal to the best estimated lower bounds on all our instances. We label these dual bound estimates as reference dual bound estimates. Our chosen number of evaluation sample paths ensures that the standard errors of all our estimated bounds are less than 0.5 percent of their respective reference dual bound estimates. When estimating dual bounds using the penalties (15), which rely on the LSMN CFA, we approximate expectations by sample averages based on 100 inner samples.

Fig. 1 displays the LSML/N-based dual bound estimates as percentages of the reference dual bound estimates on the January and April swing option instances with three exercise rights ($n = 3$). The results for the July and October swing option instances are similar and can be found in Online Supplement B. Our findings for the instances with more exercise rights are analogous and are not reported here for brevity. Fig. 2 presents the LSML/N-based dual bound estimates as percentages of the reference dual bound estimates on the January and April storage option instances with a high capacity restriction. The results for the January and April instances with medium and low capacity restrictions are comparable and are presented in Online Supplement B. Our findings for the July and October instances agree with the ones discussed here and are omitted for conciseness. Figs. 1 and 2 include the standard errors of our dual bound estimates expressed as ratios of the reference dual bound estimates, but they are not visible for most values of $P$ because they are close to zero.

The accuracy and precision of the LSML-based dual bound estimates are insensitive to the number of samples used to obtain a VFA, that is, these bounds converge when $P$ equals 1000. This insensitivity does not extend to the LSMN-based dual bound estimates, which however become comparable to the LSML-based ones for sufficiently large $P$; 10,000 on our instances. This

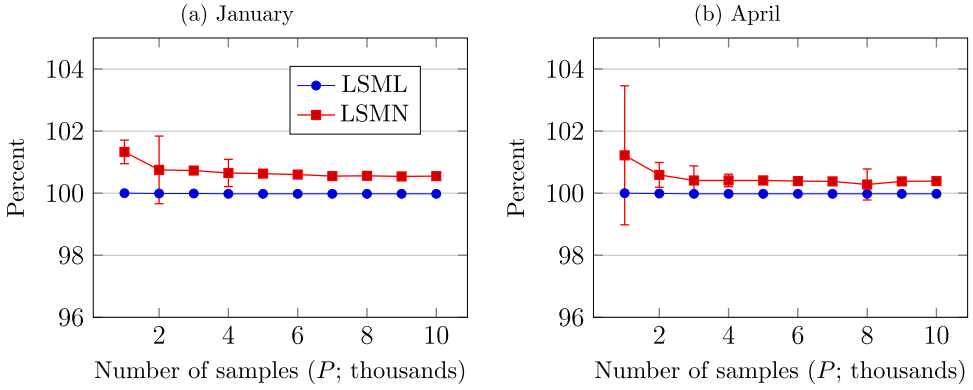


**Fig. 1.** Convergence of the LSML/N-based dual bound estimates as percentages of the reference dual bound estimates on the January and April swing option instances with three exercise rights ($n = 3$).

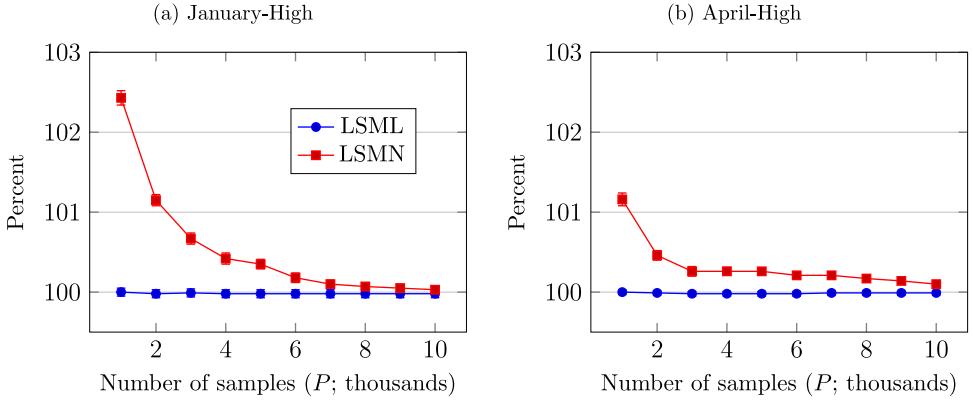


**Fig. 2.** Convergence of the LSML/N-based dual bound estimates as percentages of the reference dual bound estimates on the January and April storage option instances with a high capacity restriction.

**Table 1**
Average CPU seconds needed for computing V/CFAs and estimating lower and dual bounds on a subset of the swing option instances and on the storage option instances (the LSML VFA and LSMN CFA are estimated with $P$ equal to 1000 and 10,000, respectively).

| Swing option | | | | | | |
|---|---|---|---|---|---|---|
| $n$ | V/CFA | | Lower bound | | Dual bound | |
| | LSML | LSMN | LSML VFA | LSMN CFA | LSML VFA | LSMN CFA |
| 1 | 0.5 | 6.4 | 15.1 | 13.5 | 14.1 | 1,699 |
| 10 | 0.6 | 7.1 | 16.4 | 14.5 | 16.8 | 2,424 |
| Storage option | | | | | | |
| Capacity | LSML | LSMN | LSML VFA | LSMN CFA | LSML VFA | LSMN CFA |
| High | 0.5 | 5.5 | 3.5 | 2.3 | 13.6 | 20,149 |
| Moderate | 0.7 | 8.2 | 4.2 | 3.3 | 22.4 | 56,550 |
| Low | 0.9 | 10.0 | 4.7 | 4.1 | 27.6 | 87,828 |

observed convergence difference is consistent with our discussion in Section 3.3.

The LSML/N-based lower bound estimates exhibit a convergence pattern similar to the one discussed for the dual bound estimates. We thus omit a detailed discussion of this behavior here for brevity (Online Supplement B contains figures displaying these bounds). We only state that with sufficiently many regression sample paths the LSMN-based lower bound estimates become essentially equal to the analogous LSML-based estimates, which is expected given what we pointed out in Section 3.3. Specifically, the respective percentage ratios of these lower bound estimates and their corresponding reference dual bound estimates are at least 99.5 percent and 98.5 percent across our swing and storage option instances. Moreover, the standard errors of these lower bound estimates are no larger than 0.39 percent of their respective reference dual bound estimates. Thus, the LSML/N-based lower bound estimates are essentially equal to the exact swing and storage option values on our instances, and our dual reference bound estimates are essentially tight.

Table 1 reports the average CPU times required to run LSML and LSMN with a number of regression samples equal to 1000 and 10,000, respectively (LSMN needs more samples to converge, as discussed earlier), and estimate the lower and dual bounds based on their V/CFAs for a subset of the swing option instances and for the storage option instances. LSML and LSMN require from half to one CPU second and six to ten CPU seconds, respectively, to estimate VFAs and CFAs that lead to accurate and precise bounds. Thus, as expected given the discussion in Section 3.3, achieving this level of bounding performance with LSML is more computationally burdensome than doing so with LSMN. However, this

computational difference is small. Consistent with the discussion in Section 4, estimating the LSML/N-based lower bounds takes comparable effort. Specifically, between four and fifteen CPU seconds for LSML and two and fifteen CPU seconds for LSMN. The total CPU effort exerted to compute the LSMN/L C/VFAs and estimate their corresponding lower bounds is thus roughly the same on our instances. In stark contrast, the LSML-based dual bounds can be estimated two to three orders of magnitude faster compared to the LSMN-based dual bounds: Between fourteen and twenty-eight CPU seconds versus between half and twenty four CPU *hours*. Although we expect the estimation of a dual bound to be faster when using a VFA rather than a CFA, as pointed out in Section 4, this observed difference is both substantial and remarkable.

## 6. Conclusions

We compare LSMN and LSML for the valuation and management of crude oil swing and natural gas storage options using a common energy term structure evolution model. Applied to realistic instances of these real options, the C/VFA obtained by LSMN/L yield similarly accurate and precise lower and dual bound estimates. The computational efforts exerted to compute the LSMN/L C/VFAs and estimate their corresponding lower bounds are comparable. Instead, estimating dual bounds based on the LSML VFAs instead of the LSMN CFAs takes seconds rather than minutes or hours. This difference occurs because with our chosen term structure models using the LSML VFAs allows bypassing the nested simulations and the optimizations that are needed when employing the LSMN CFAs. Our research thus brings to light LSML, a largely unexplored LSM variant, as a more appealing method than LSMN to obtain the inputs required for estimating dual bounds on the values of energy swing and storage options, two common real options, when using typical futures term structure models. Beyond these real options, our findings have potential relevance for the computation of dual bounds on the values of other early or multiple exercise options, as well as of optimal policies for capacity investment and inventory/production management models with supply/demand forecast updates.

## Acknowledgments

The authors thank the review team for their comments that helped them to improve the quality of the initial version of this paper. This research is based upon work supported by the National Science Foundation under Grant CMMI 1129163. The second author was also supported by ONR Grant N00014-12-10032. The third author is a Faculty Affiliate of the Scott Institute for Energy Innovation at Carnegie Mellon University.

## Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.ejor.2016.06.020.

## References

Adkins, R., & Paxson, D. (2011). Reciprocal energy-switching options. *Journal of Energy Markets, 4*(1), 91–120.
Andersen, L., & Broadie, M. (2004). Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science, 50*(9), 1222–1234.
Arvesen, Ø., Medbø, V., Fleten, S. E., Tomasgard, A., & Westgaard, S. (2013). Linepack storage valuation under price uncertainty. *Energy, 52*(1), 155–164.
Barbieri, A., & Garman, M. B. (1996). Putting a price on swings. *Energy and Power Risk Management, 1*(6), 17–19.
Bäuerle, N., & Riess, V. (2016). Gas storage valuation with regime switching. *Energy Systems*. (Forthcoming).
Bertsekas, P. B. (2007). *Dynamic programming and optimal control*: vol. 2 (3rd ed.). Belmont, MA, USA: Athena Scientific.
Blanco, C., Soronow, D., & Stefiszyn, P. (2002). Multi-factor models for forward curve analysis: An introduction to principal component analysis. *Commodities Now*, 76–78.
Boogert, A., & De Jong, C. (2008). Gas storage valuation using a Monte Carlo method. *The Journal of Derivatives, 15*(3), 81–98.
Boogert, A., & De Jong, C. (2011/12). Gas storage valuation using a multifactor price process. *The Journal of Energy Markets, 4*(4), 29–52.
Boomsma, T. K., Meade, N., & Fleten, S. E. (2012). Renewable energy investments under different support schemes: A real options approach. *European Journal of Operational Research, 220*(1), 225–237.
Brandão, L., Dyer, J. S., & Hahn, W. J. (2005). Using binomial decision trees to solve real-option valuation problems. *Decision Analysis, 2*(2), 69–88.
Broadie, M., & Cao, M. (2008). Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance, 8*(8), 845–861.
Brown, D. B., Smith, J. E., & Sun, P. (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research, 58*(4), 1–17.
Carmona, R., & Ludkovski, M. (2010). Valuation of energy storage: An optimal switching approach. *Quantitative Finance, 10*(4), 359–374.
Carriere, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics, 19*(1), 19–30.
Chandramouli, S. S., & Haugh, M. (2012). A unified approach to multiple stopping and duality. *Operations Research Letters, 40*(4), 258–264.
Clewlow, L., & Strickland, C. (2000). *Energy derivatives: Pricing and risk management*. London, UK: Lacima.
Cortazar, G., Gravet, M., & Urzua, J. (2008). The valuation of multidimensional American real options using the LSM simulation method. *Computers & Operations Research, 35*(1), 113–129.
Cortazar, G., & Schwartz, E. S. (1994). The valuation of commodity contingent claims. *The Journal of Derivatives, 1*(4), 27–39.
Denault, M., Simonato, J. G., & Stentoft, L. (2013). A simulation-and-regression approach for stochastic dynamic programs with endogenous state variables. *Computers & Operations Research, 40*(11), 2760–2769.
Desai, V. V., Farias, V. F., & Moallemi, C. C. (2012). Pathwise optimization for optimal stopping problems. *Management Science, 58*(12), 2292–2308.
Detemple, J. (2006). *American-style derivatives: Valuation and computation*. London, UK: Taylor & Francis.
Devalkar, S. K., Anupindi, R., & Sinha, A. (2011). Integrated optimization of procurement, processing, and trade of commodities. *Operations Research, 59*(6), 1369–1381.
Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton, NJ, USA: Princeton University Press.
Enders, P., Scheller-Wolf, A., & Secomandi, N. (2010). Interaction between technology and extraction scaling real options in natural gas production. *IIE Transactions, 49*(9), 643–655.
EnergyQuants (2016). Swing planner and storage planner. http://www.energyquants.com/software/.
Eydeland, A., & Wolyniec, K. (2003). *Energy and power risk management*. Hoboken, NJ, USA: John Wiley and Sons Inc.
Felix, B. J., & Weber, C. (2012). Gas storage valuation applying numerically constructed recombining trees. *European Journal of Operational Research, 216*(1), 178–187.
Fontes, D. (2008). Fixed versus flexible production systems: A real options analysis. *European Journal of Operational Research, 188*(1), 169–184.
Gamba, A. (2003). *Real options valuation: A Monte Carlo approach*. University of Calgary. Working paper
Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. New York, NY, USA: Springer.
Glasserman, P., & Yu, B. (2004). Simulation for American options: Regression now or regression later?. In H. Niederreiter (Ed.), *Monte Carlo and quasi-Monte Carlo methods 2002* (pp. 213–226). Berlin, Germany: Springer-Verlag.
Goel, A., & Gutierrez, G. J. (2011). Multiechelon procurement and distribution policies for traded commodities. *Management Science, 47*(12), 2228–2244.
Graves, S., Meal, H., Dasu, S., & Qiu, Y. (1986). Two-stage production planning in a dynamic environment. In S. Axsäter, C. Schneeweiss, & E. Silver (Eds.), *Multi-stage production planning and inventory control*, Lecture Notes in Economics and Mathematical Systems (pp. 9–43). Berlin, Germany: Spinger-Verlag.
Guthrie, G. (2009). *Real options in theory and practice*. NY, USA: Oxford University Press.
Gyurkó, L. G., Hambly, B., & Witte, J. H. (2015). Monte Carlo methods via a dual approach for some discrete time stochastic control problems. *Mathematical Methods of Operations Research, 81*(1), 109–135.
Hahn, W. J., & Dyer, J. S. (2008). Discrete time modeling of mean-reverting stochastic processes for real option valuation. *European Journal of Operational Research, 184*(2), 534–548.
Hahn, W. J., & Dyer, J. S. (2011). A discrete time approach for modeling two-factor mean-reverting stochastic processes. *Decision Analysis, 8*(3), 220–232.
Haug, E. G. (2006). *The complete guide to option pricing formulas* (2nd ed.). NY, USA: McGraw-Hill.
Haugh, M. B., & Kogan, L. (2007). Duality theory and approximate dynamic programming for pricing American options and portfolio optimization. In J. R. Birge, & V. Linetsky (Eds.), *Financial engineering*, Handbooks in Operations Research and Management Science*: vol. 15* (pp. 925–948). Amsterdam, The Netherlands: Elsevier.

Heath, D., & Jackson, P. L. (1994). Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions, 26*(3), 17–30.
Ho, T. S. Y., & Lee, S. B. (1986). Term structure movements and pricing interest rate contingent claims. *Journal of Finance, 41*(5), 1011–1029.
Hull, J. (2014). *Options, futures and other derivatives* (9th). Englewood Cliffs, NJ, USA: Pearson Prentice Hall.
Iida, T., & Zipkin, P. (2006). Approximate solutions of a dynamic forecast-inventory model. *Manufacturing & Service Operations Management, 8*(4), 407–425.
Jaillet, P., Ronn, E. I., & Tompaidis, S. (2004). Valuation of commodity-based swing options. *Management Science, 50*(7), 909–921.
Khansa, L., & Liginlal, D. (2009). Valuing the flexibility of investing in security process innovations. *European Journal of Operational Research, 192*(1), 216–235.
Kouvelis, P., Chambers, C., & Wong, H. (2006). Supply chain management research and production and operations management: Review, trends, and opportunities. *Production and Operations Management, 15*(3), 449–469.
KYOS (2013). Guidelines for valuation of real options in energy markets. KYOS Energy Consulting white paper.
Lai, G., Margot, F., & Secomandi, N. (2010). An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research, 58*(3), 564–582.
Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies, 14*(1), 113–147.
Maragos, S. (2002). Valuation of the operational flexibility of natural gas storage reservoirs. In E. Ronn (Ed.), *Real options and energy management using options methodology to enhance capital budgeting decisions* (pp. 431–456). London, UK: Risk Publications.
Matlab (2015). Financial instruments toolbox: User's guide.
Mazières, D., & Boogert, A. (2013). A radial basis function approach to gas storage valuation. *The Journal of Energy Markets, 6*(2), 19–50.
Nadarajah, S., Margot, F., & Secomandi, N. (2015). Relaxations of approximate linear programs for the real option management of commodity storage. *Management Science, 61*(12), 3054–3076.
Nandalal, K. D. W., & Bogardi, J. J. (2007). *Dynamic programming based operation of reservoirs: applicability and limits*. Cambridge, UK: Cambridge University Press.
Porteus, E. (2002). *Foundations of stochastic inventory theory*. CA, USA: Stanford Business Books.
Powell, W. B. (2011). *Approximate dynamic programming: Solving the curses of dimensionality* (2nd ed.). Hoboken, NJ, USA: John Wiley & Sons.
Secomandi, N. (2010). Optimal commodity trading with a capacitated storage asset. *Management Science, 56*(3), 449–467.
Secomandi, N., Lai, G., Margot, F., Scheller-Wolf, A., & Seppi, D. (2015). Merchant commodity storage and term structure model error. *Manufacturing and Service Operations Management, 17*(3), 302–320.
Secomandi, N., & Seppi, D. J. (2014). Real options and merchant operations of energy and other commodities. *Foundations and Trends in Technology, Information and Operations Management, 6*(3–4), 161–331.
Shreve, S. (2004). *Stochastic calculus for finance II: Continuous-time models*. New York, New York, USA: Springer.
Smith, J. (2005). Alternative approaches for solving real-options problems. *Decision Analysis, 2*(2), 89–102.
Smith, J. E., & McCardle, K. F. (1998). Valuing oil properties: Integrating option pricing and decision analysis approaches. *Operations Research, 46*(2), 198–217.
Smith, J. E., & McCardle, K. F. (1999). Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research, 47*(1), 1–15.
Thompson, M. (2012). *Natural gas storage valuation, optimization, market and credit risk management*. Queens University. Working paper.
Thompson, M. (2013). Optimal economic dispatch and risk management of thermal power plants in deregulated markets. *Operations Research, 61*(4), 791–809.
Thompson, M., Davison, M., & Rasmussen, H. (2004). Valuation and optimal operation of electric power plants in competitive markets. *Operations Research, 52*(4), 546–562.
Triantis, A. J., & Hodder, J. E. (1990). Valuing flexibility as a complex option. *The Journal of Finance, 45*(2), 549–565.
Trigeorgis, L. (1996). *Real options: Managerial flexibility and strategy in resource allocation*. Cambridge, MA, USA: The MIT Press.
Tsitsiklis, J. N., & Van Roy, B. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks, 12*(4), 694–703.
Veronesi, P. (2010). *Fixed income securities: Valuation, risk, and risk management*. Hoboken, NJ, USA: John Wiley and Sons Inc.
Wang, T., Atasu, A., & Kurtuluş, M. (2012). A multiordering newsvendor model with dynamic forecast evolution. *Manufacturing & Service Operations Management, 14*(3), 472–484.
Wang, T., & Dyer, J. S. (2010). Valuing multifactor real options using an implied binomial tree. *Decision Analysis, 7*(2), 185–195.
Wu, O. Q., Wang, D. D., & Qin, Z. (2012). Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management, 14*(3), 455–471.
Zhao, T., Cao, X., & Yang, D. (2011). Effect of streamflow forecast uncertainty on real-time reservoir operation. *Advances in Water Resources, 34*(4), 495–504.
Zhao, T., Zhao, J., Yang, D., & Wang, H. (2013). Generalized martingale model of the uncertainty evolution of streamflow forecasts. *Advances in Water Resources, 57*, 41–51.
Zipkin, P. (2000). *Foundations of inventory management*. NY, USA: McGraw-Hill.