



HAL
open science

Numerical methods for the pricing of Swing options: a stochastic control approach

Christophe Barrera-Esteve, Florent Bergeret, Charles H Dossal, Emmanuel Gobet, Asma Meziou, Rémi Munos, Damien Reboul-Salze

► **To cite this version:**

Christophe Barrera-Esteve, Florent Bergeret, Charles H Dossal, Emmanuel Gobet, Asma Meziou, et al.. Numerical methods for the pricing of Swing options: a stochastic control approach. *Methodology and Computing in Applied Probability*, 2006, *Methodology and Computing in Applied Probability*, 8 (4), pp.517-540. 10.1007/s11009-006-0427-8 . inria-00117175

HAL Id: inria-00117175

<https://inria.hal.science/inria-00117175v1>

Submitted on 1 Dec 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Numerical methods
for the pricing of Swing options:
a stochastic control approach

Christophe BARRERA-ESTEVE * Florent BERGERET †
Charles DOSSAL ‡ Emmanuel GOBET §¶ Asma MEZIOU ||
Rémi MUNOS ** Damien REBOUL-SALZE ††

11th February 2006

Abstract

In the natural gas market, many derivative contracts have a large degree of flexibility. These are known as *Swing* or *Take-Or-Pay* options. They allow their owner to purchase gas daily, at a fixed price and according to a volume of their choice. Daily, monthly and/or annual constraints on the purchased volume are usually incorporated. Thus, the valuation of such contracts is related to a stochastic control problem, which we solve in this paper using new numerical methods. Firstly, we extend the Longstaff-Schwarz methodology (originally used for Bermuda options) to our case. Secondly, we propose two efficient parameterizations of the gas consumption, one is based on neural networks and the other on finite elements. It allows us to derive a local optimal

*Gaz de France - Direction de la Recherche - Projet Valoris - 361, avenue du Président Wilson - BP 33
- 93211 Saint Denis La Plaine Cedex - FRANCE. Email: christophe.barrera-estevé@gazdefrance.com

†Gaz de France. Email: florent.bergeret@gazdefrance.com

‡Ecole Polytechnique - Centre de Mathématiques Appliquées - 91128 Palaiseau Cedex - FRANCE.

Email: dossal@cmapx.polytechnique.fr

§Corresponding author

¶ENSIMAG - INP Grenoble - Laboratoire de Modélisation et Calcul - BP 53 - 38041 Grenoble Cedex
9 - FRANCE. Email: emmanuel.gobet@imag.fr. Previously at Ecole Polytechnique.

||Ecole Polytechnique. Email: meziou@cmapx.polytechnique.fr

**Ecole Polytechnique. Email: remi.munos@polytechnique.fr

††Previously at Gaz de France.

consumption law using a stochastic gradient ascent. Numerical experiments illustrate the efficiency of these approaches. Furthermore, we show that the optimal purchase is of bang-bang type.

Key words: Swing options, Monte Carlo simulations, bang-bang control, parametric consumption, stochastic gradient.

AMS 2000 Subject Classification numbers: 90C31, 91B02, 93Exx.

Introduction

Since natural gas is a physical asset, the associated energy market is different indeed from the non-physical equity markets. In fact, the model needs to incorporate specific features, such as seasonality effects: we refer for these aspects to the paper by Jaillet, Ronn and Tompaidis (2004) and references therein. Another difference, on which we focus in this work, comes from the nature of financial contracts written on gas. These latter are aimed at accommodating the uncertain volume of energy purchased by the end user and then, include many flexibilities of delivery. Namely, the owner of such a contract is usually allowed to purchase some amount of gas at some dates. The purchase may be daily, or at a fixed number of dates chosen by the owner. The volume of gas is subject to daily as well as periodic (monthly or annual) constraints. Moreover, the purchase price may be fixed or predetermined according to some factors, such as the past values of the crude oil. These options are known as *Swing* or *Take-or-Pay* options (see the description in Thompson (1995)). They are also traded on other energy markets (e.g. electricity). However, as we will consider in the sequel the case of *storage contracts* that are specific to the gas industry, we will mainly focus on this asset in our discussions. The aim of the paper is to propose new numerical approaches for the pricing of Swing options and to compare them with existing methods.

The valuation of such contracts is known to be a challenging task: their price depends on the price of the underlying asset and the cumulative consumption. Actually, it may be written as a solution of a dynamic programming equation (see Thompson (1995); Lari-Lavassani, Simchi and Ware (2001); Clewlow, Strickland and Kaminski (2001); Raikar and

Ilić (2001); Jaillet, Ronn and Tompaidis (2004)) with a two-dimensional state variable (see Paragraph 2.1 for a better description). At each date, given a value of the underlying asset and the cumulative consumption, the price is computed by maximizing over the possible consumptions the immediate payoff plus the expected cash flows. In practice, a methodology extending the usual binomial tree method can be performed: it is the so-called *forest tree* approach (see the references above). This can be successfully implemented in the case of a one-factor model on the forward prices. However, this discretizes a continuous dynamics in space and thus incorporates errors in relation to this discretization procedure. Furthermore, for multi-factor models, the computational complexity blows up and the method is not feasible in practice. For these reasons, it may be more efficient to combine the backward dynamic programming equation with a forward simulation of the underlying asset, analogously to what can be done for Bermuda options (where only one exercise date has to be determined in an optimal way). So here, we propose an adaptation of the Longstaff-Schwarz methodology (Longstaff and Schwartz (2001)) to our case (where optimal volumes at each date have to be computed). This adaptation is discussed in Section 2.

One inherent feature of the dynamic programming equation is that we have an idea of the price today, only when all the intermediate optimal consumptions have been computed, in any state, between the current date and the maturity. As the purchase frequency increases (it is daily in our numerical examples), this iteration takes more and more time and the exact resolution of this nonlinear problem turns out to be too costly. So, instead of computing the optimal consumption, we may prefer to identify a fairly good consumption rule which would provide a price close to the optimal one. This is the so-called *policy search algorithm*, which transforms a stochastic control problem into a parametric optimization one (see Kushner and Yang (1991)). This approach is developed in Section 3. The success of this method depends strongly on our ability to find a nice parameterization of the consumption. Under some assumptions, the Pontryagin principle enables us to prove that the optimal consumption is of *bang-bang* type: this new result is justified in Appendix A. Then, based on this observation, we derive two relevant parameterizations and test them numerically.

1 Statement of the problem and notations

Here we give the notation which will be used throughout the paper.

We consider contracts with maturity T . Since the spot price of the gas does not correspond to a tradable instrument, the valuation and the hedging of contracts on gas are performed through the use of forward contracts with different maturities $t \in [0, T]$ (see the approach given in Jaiilet, Ronn and Tompaidis (2004)): their prices are denoted by $(F(s, t))_{0 \leq s \leq t}$. The interest rates are assumed to be constant, equal to r and for the sake of simplicity, we consider a one-factor yield, as in Jaiilet, Ronn and Tompaidis (2004):

$$\frac{dF(s, t)}{F(s, t)} = \sigma e^{-\alpha(t-s)} dW_s, \quad 0 \leq s \leq t. \quad (1.1)$$

In the last expression, the initial condition is defined by the forward curve $(F(0, t))_{0 \leq t \leq T}$




Figure 1: Gas Forward Curve

(corresponding to the seasonality, as shown on Figure 1). It would be easy to incorporate extra factors or a convenience yield term, without major modifications to what follows. Within our framework, the spot price at time t is given by $F_t = F(t, t)$. In Equation (1.1), the dynamics is given directly under the usual risk-neutral probability measure \mathbf{Q} . W defines a linear Brownian motion and so, the probability \mathbf{Q} transforms the forward contract processes (or equivalently the discounted option prices) into martingales. We recall that the true dynamics of the forward contracts is not under \mathbf{Q} but under the historical probability; however, arbitrage arguments show that the fair value of a financial contract has to be evaluated as the \mathbf{Q} -expectation of its discounted cash-flow (see Karatzas

(1988) for American options, or Jaillet, Ronn and Tompaidis (2004) for Swing options). Since we focus on pricing issues, we consider the dynamics only under \mathbf{Q} .

Now, let us describe a general Swing contract. We assume that energy is purchased at some fixed dates $0 = t_0 < \dots < t_i < \dots < t_N = T$. The volume at time t_i is denoted by q_{t_i} and subject to the constraint

$$q_{\min} \leq q_{t_i} \leq q_{\max}. \quad (1.2)$$

The cumulative volume purchased up to time t_i is given by $Q_{t_i} = \sum_{j=0}^{i-1} q_{t_j}$. It must satisfy the terminal constraint

$$Q_{\min} \leq Q_T \leq Q_{\max}. \quad (1.3)$$

To simplify, the purchase price is assumed to be constant, equal to K . However if necessary, it can be extended to a deterministic time function. Hence, the buyer of the contract gets $q_{t_i}(F_{t_i} - K)$ at time t_i . Furthermore, when the global purchase constraints are violated, he may apply some penalties at maturity. These can be proportional to the spot price F_T at maturity and to the over or under-consumption: $-A F_T [Q_T - Q_{\min}]_- - B F_T [Q_T - Q_{\max}]_+$, where A and B are positive constants. To be concise, we denote this penalty by $\mathcal{P}_T(F_T, Q_T)$ in the sequel. For a given consumption strategy $q = (q_{t_i})_{0 \leq i \leq N}$, the fair price is thus given by

$$\mathcal{J}(q) = \mathbf{E} \left(\sum_{i=0}^{N-1} e^{-rt_i} q_{t_i} (F_{t_i} - K) + e^{-rT} \mathcal{P}_T(F_T, Q_T) \right) \quad (1.4)$$

(the expectation is performed under the risk-neutral probability measure \mathbf{Q}).

Case 1 *In the further numerical experiments about Swing contracts, we shall consider the same following set of parameters: $\sigma = 70\%$, $\alpha = 4$, $r = 0\%$, $q_{\min} = 0$, $q_{\max} = 6$, $Q_{\min} = 1300$, $Q_{\max} = 1900$, $K = 20$, $T = 1$ year with daily exercise ($N = 365$). We will use the forward curve described in Figure 1. To discuss the impact of seasonality in the numerical procedure, we will also consider a flat forward curve $F(0, t_i) = 20, \forall 0 \leq i \leq N$.*

The preceding description has to be slightly modified for *storage contracts*. In this case, an empty tank is rent to the owner of the contract and allows him to either buy some gas on the market and inject it in the tank, or withdraw it and sell it on the market. At each

date, the profit is given by

$$\psi(F_{t_i}, q_{t_i}) = \begin{cases} -q_{t_i}(F_{t_i} + c_I) & \text{if } q_{t_i} \geq 0 \text{ (Injection),} \\ -q_{t_i}(F_{t_i} - c_W) & \text{if } q_{t_i} \leq 0 \text{ (Withdrawal),} \\ 0 & \text{if } q_{t_i} = 0 \text{ (Same level for the storage),} \end{cases}$$

where $c_I > 0$ (resp. $c_W > 0$) denotes some strictly positive injection (resp. withdrawal) cost. Moreover, the tank capacity has some volume limitations, which imposes constraints like $0 \leq f_1(t) \leq Q_t \leq f_2(t)$ for some deterministic functions f_1 and f_2 . For the next optimization problem, we can formulate these physical constraints in terms of a smooth penalization : $-C \sum_{i=0}^N e^{-rt_i} [(Q_{t_i} - f_1(t_i))_-]^2 - C \sum_{i=0}^N e^{-rt_i} [(Q_{t_i} - f_2(t_i))_+]^2$, with a large constant C . If we consider this penalization as a financial penalty, the fair price of this *storage contract* is then given by

$$\mathcal{J}(q) = \mathbf{E} \left(\sum_{i=0}^N \Psi_{t_i}(q_{t_i}, F_{t_i}, Q_{t_i}) \right), \quad (1.5)$$

where $\Psi_t(q, F, Q) = e^{-rt} (\psi(F, q) - C[(Q - f_1(t))_-]^2 - C[(Q - f_2(t))_+]^2)$ if $t < t_N$ and $\Psi_{t_N}(q, F, Q) = e^{-rt_N} (-C[(Q - f_1(t_N))_-]^2 - C[(Q - f_2(t_N))_+]^2)$.

Case 2 *The storage contract we will consider in the next numerical experiments is such that the tank is empty at the maturity of the contract ($Q_{\min} = Q_{\max} = 0$) and the tank capacity equals 20 units say (this determines the function f_2). The parameters will be set as follows: $\sigma = 70\%$, $\alpha = 4$, $r = 0\%$, $q_{\min} = -0.2$, $q_{\max} = 0.4$, $f_1 \equiv 0$, $f_2(t_i) = \min(q_{\max}i, 20, q_{\min}(i - N))$, $c_I = 0.6$, $c_W = 0.2$, $K = 20$, $T = 1$ year with daily exercise ($N = 365$). As for the Swing contract, we will use the forward curve on Figure 1, but also the flat forward curve $F(0, t_i) = 20$.*

These descriptions of Swing options and storage contracts can be unified within the following formulation of the price

$$\mathcal{J}(q) = \mathbf{E} \left(\sum_{i=0}^{N-1} \Psi_{t_i}(q_{t_i}, F_{t_i}, Q_{t_i}) + \mathcal{P}_T(F_T, Q_T) \right). \quad (1.6)$$

In the above, q is a given consumption strategy, $\Psi(\cdot)$ denotes some appropriate instantaneous profit and $\mathcal{P}_T(\cdot)$ some terminal penalty.

To price these instruments, we make the crucial assumption that the consumption decisions

are taken according to the sole evolution of the underlying asset, and not the real needs in gas. So, the price is given by

$$\sup_q \mathcal{J}(q).$$

It is a stochastic optimal control problem. In the next section, we state the associated dynamic programming equation (see Bertsekas and Shreve (1978), Bertsekas (1995)) from which we derive some numerical methods.

2 Numerical methods based on dynamic programming equation

First, for both theoretical and numerical purposes, it is better to rewrite $(F_t)_{0 \leq t \leq T}$ using an extra state variable $X = (X_t)_{0 \leq t \leq T}$ defined by $X_t = \sigma e^{-\alpha t} \int_0^t e^{\alpha u} dW_u$. Note that X is an Ornstein-Uhlenbeck process, which solves $dX_t = -\alpha X_t dt + \sigma dW_t$. Yet, this property will not be really used in the sequel. Moreover, an immediate application of Itô's formula yields $F(s, t) = F(0, t) \exp[e^{-\alpha(t-s)} X_s - \frac{\sigma^2}{4\alpha}(e^{-2\alpha(t-s)} - e^{-2\alpha t})]$. We particularly have

$$F_{t_i} = F(0, t_i) \exp\left(X_{t_i} - \frac{\sigma^2}{4\alpha}(1 - e^{-2\alpha t_i})\right). \quad (2.7)$$

2.1 Writing the Dynamic Programming Equation

As already mentioned, the valuation of Swing contracts can be embedded into a stochastic control framework. Indeed, the dynamics of the system is described by a discrete-time Markov chain $(F_{t_i}, Q_{t_i})_{0 \leq i \leq N}$ that is controlled by the local consumption $(q_{t_i})_{0 \leq i \leq N-1}$. Given a time t_i and a state (F, Q) , the price is given by

$$P(t_i, F, Q) = \sup_{(q_{t_j})_{i \leq j \leq N-1}} \mathbf{E}\left(\sum_{j=i}^{N-1} \Psi_{t_j}(q_{t_j}, F_{t_j}, Q_{t_j}) + \mathcal{P}_T(F_T, Q_T) \mid F_{t_i} = F, Q_{t_i} = Q\right).$$

In the last expression, the sup is taken over all admissible consumptions. A standard result states that, under mild assumptions, there exists an optimal Markovian consumption, given by the dynamic programming equation (also known as Bellman's equation).

Proposition 2.1 *Assume that for some positive constants p and C , the following inequalities hold for any $t_i \in [0, T]$, $F > 0$, $q \in [q_{\min}, q_{\max}]$ and $(Q, Q') \in [-N(q_{\min})_-, N(q_{\max})_+]$:*

$$\begin{aligned} |\Psi_{t_i}(q, F, Q)| + |\mathcal{P}_T(F, Q)| &\leq C(1 + F^p), \\ |\Psi_{t_i}(q, F, Q') - \Psi_{t_i}(q, F, Q)| &\leq C(1 + F^p)|Q' - Q|. \end{aligned}$$

Then, there exists an optimal Markov consumption $q^(t, F, Q)$, which is given by the argmax in the following dynamic programming equation :*

$$\begin{aligned} P(t_i, F, Q) &= \max_{q \in [q_{\min}, q_{\max}]} [\Psi_{t_i}(q, F, Q) + \mathbf{E}(P(t_{i+1}, F_{t_{i+1}}, Q + q) | F_{t_i} = F)], \quad (t_i < T), \\ P(T, F, Q) &= \mathcal{P}_T(F, Q). \end{aligned} \tag{2.8}$$

Note that this result holds for both Swing and storage contracts.

PROOF. This is an immediate application of Proposition 8.5. in Bertsekas and Shreve (1978). First, the conditions on the growth of instantaneous and terminal profits together with the classical estimates $\mathbf{E}(\sup_{0 \leq t \leq T} F_t^p) < \infty$ enable us to check assumptions (F^+) and (F^-) in Proposition 8.5. Then, it remains to prove that the supremum in (2.8) is attained. This immediately follows from the continuity of $P(t_{i+1}, F, Q)$ w.r.t. Q (which is a easy consequence of our Lipschitz continuity hypothesis). Finally, since the dynamics of the forward contracts does not depend on $(q_{t_i})_i$, the usual conditional expectation w.r.t. (F_{t_i}, Q_{t_i}) can be simplified using the relation $Q_{t_{i+1}} = Q_{t_i} + q_{t_i}$. It provides the dynamic programming equation as expressed in (2.8).

□

Paragraphs 2.2 and 2.3 respectively implement the $\max_{q \in [q_{\min}, q_{\max}]}$ operation and the computation of $\mathbf{E}(\cdot | F_{t_i} = F)$ for solving approximately the dynamic programming equation (2.8).

2.2 Towards an implementation of the dynamic programming equation

As usual, solving the dynamic programming equation (2.8) is far from being trivial because of the conditional expectation and the optimization w.r.t. q at each time. For a numerical resolution, this formulation must be simplified. To do so, we first discretize the set of admissible values of q_{t_i} : $q_{t_i} \in \{q_{\min}, q_{\min} + \Delta q, \dots, q_{\max} - \Delta q, q_{\max}\}$ where Δq denotes a




Figure 2: Restriction of the set of admissible purchased volumes when the global constraints can not be violated.

positive consumption step. Consequently, the purchased cumulative volume $Q_{t_i} = \sum_{j=0}^{i-1} q_{t_j}$ belongs to a set of possible values $\{Q_{t_i}^1, Q_{t_i}^2, \dots, Q_{t_i}^{M_{t_i}}\}$, where M_{t_i} denotes the number of possible cumulative volumes that are purchased up to time t_i . If it is not allowed to break the global purchase constraints Q_{\min} and Q_{\max} (i.e. the applied penalties are infinite), the set $\mathcal{A} = \{(t_i, Q_{t_i}) : 0 \leq i \leq N\}$ of admissible cumulative consumptions can be efficiently limited. This will be extremely useful to speed up the resolution of (2.8). For example, let us consider a contract with the characteristics $Q_{\max} = 1$, $Q_{\min} = 0$, $\Delta q = q_{\max} = 1$ and $q_{\min} = -2$, over the period $\{t_0 = 0, t_1, \dots, t_4 = T\}$. The set \mathcal{A} is represented on Figure 2. For a given couple (t_i, Q_{t_i}) , denote by $\text{Adm}(t_i, Q_{t_i})$ the set of discrete admissible values of q_{t_i} . If the constraints can be broken, then $\text{Adm}(t_i, Q_{t_i}) = \{q_{\min}, q_{\min} + \Delta q, \dots, q_{\max} - \Delta q, q_{\max}\}$. Otherwise it can be much smaller (see Figure 2). Thus, for any $(t_i, Q_{t_i}) \in \mathcal{A}$, (2.8) becomes

$$P(t_i, F, Q_{t_i}) = \max_{q \in \text{Adm}(t_i, Q_{t_i})} [\Psi_{t_i}(q, F, Q_{t_i}) + \mathbf{E}(P(t_{i+1}, F_{t_{i+1}}, Q_{t_i} + q) | F_{t_i} = F)]. \quad (2.9)$$

In light of (2.9), we should compute as many functions $P(t_i, \cdot, Q_{t_i})$ as elements in the set \mathcal{A} . However, this latter can be quite large (see Example 1), and thus much time consuming. This is an inherent drawback of this approach and as we will later see, we do not meet these issues with the approaches based on parameterized consumptions.

Example 1 Consider a storage contract where, the initial level of gas is 17 units, the maximum possible volume is 24 units and the final volume can take any value between 2

units and 17 units. The set \mathcal{A} represented on Figure 3, gives an idea about the complexity. The daily injection capacity is one unit per exercise whereas the withdrawal capacity is two units. We consider in this example 20 possible exercise dates. Here, q_{t_i} can take any integer value between -2 and 1 . We can yet speed up the resolution if we further assume that the consumption is of bang-bang type (see Theorem 3.1 below). In this case, we can simply take $\Delta_q = q_{\max} - q_{\min}$.




Figure 3: States where the price function $P(t_i, Q_{t_i}, \cdot)$ should be computed

To complete the resolution of the dynamic programming equation (2.9), we develop in the sequel two different strategies.

1. Using a trinomial tree to describe the evolution of $(F_{t_i})_{0 \leq i \leq N}$. It leads to the so-called forest of trees algorithm (see Lari-Lavassani, Simchi and Ware (2001); Jaillet, Ronn and Tompaidis (2004)).
2. Using Monte Carlo simulations of $(F_{t_i})_{0 \leq i \leq N}$, which are much more adapted to multi-factor models. They enable to compute conditional expectations involved in (2.9) (with another representation) using regression techniques (analogous to the Longstaff-Schwartz approach (2001) for Bermuda options). This approach has not been used in the existing literature on Swing options.

2.3 Forest of trees

2.3.1 Description of the algorithm

In light of (2.7), we need to construct a trinomial tree, in order to describe how $(X_{t_i})_{0 \leq i \leq N}$ (and therefore $(F_{t_i})_{0 \leq i \leq N}$) evolves. We briefly remind the construction of such a tree (for details, see Lari-Lavassani, Simchi and Ware (2001)). Each node can either go up, or go down or even stay in the same physical state. Three kinds of branches are then available to fit the mean-reverting behavior of X . See Figure 4, where case 2 and 3 correspond respectively to the top and the bottom of the trinomial tree. The spacing between nodes




Figure 4: On the left, three possible branches. On the right, trinomial tree.

is equal to $\sigma\sqrt{3\Delta t}$, where Δt denotes the discretization time step. Table 1 gives the probabilities on each alternative branching. These latter are computed in order to match the first two moments, and their sum is naturally equal to one. As probabilities must not be negative, we deduce that necessarily $j_{\max} = \left\lceil \frac{0.184}{\alpha\Delta t} \right\rceil$, and this bounds the tree's width since $j \in [-j_{\max}, j_{\max}]$. Denote by $X_{t_i}^j$ the value of X_{t_i} on node j in the trinomial tree, and

	Case 1	Case 2 ($j = j_{\max}$)	Case 3 ($j = -j_{\max}$)
p_{up}	$\frac{1}{6} - \frac{1}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$	$\frac{7}{6} - \frac{3}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$	$\frac{1}{6} + \frac{1}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$
p_{middle}	$\frac{2}{3} - [\alpha j\Delta t]^2$	$-\frac{1}{3} + 2\alpha j\Delta t - [\alpha j\Delta t]^2$	$-\frac{1}{3} - 2\alpha j\Delta t - [\alpha j\Delta t]^2$
p_{down}	$\frac{1}{6} + \frac{1}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$	$\frac{1}{6} - \frac{1}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$	$\frac{7}{6} + \frac{3}{2}\alpha j\Delta t + \frac{1}{2}[\alpha j\Delta t]^2$

Table 1: Probabilities depending on the vertical position j in the tree.

by L_{t_i} the number of nodes at time t_i . To each point $(t_i, Q_{t_i}^{j_{t_i}})_{j_{t_i}=1, \dots, M_{t_i}}$ on the graphic, corresponds a range $(X_{t_i}^j)_{j=1, \dots, L_{t_i}}$ of nodes in the tree (at time t_i). Formula (2.7) gives the corresponding spot prices of gas $(F_{t_i}^j)_{j=1, \dots, L_{t_i}}$. Consider now the algorithm for the forest of trees. By Equation (2.9), the price $P(t_0, F_{t_0}^1, Q_{t_0}^1)$ of the option is computed backwards in the forest, in the following way.

- **Initialization.** Compute the terminal values $P(t_N, F_{t_N}^j, Q_{t_N}^k)$ for $j = 1, \dots, L_{t_N}$ and $k = 1, \dots, M_{t_N}$.
- **Iteration.** Assume now that the values $P(t_{i+1}, F_{t_{i+1}}^j, Q_{t_{i+1}}^k)$ for $j = 1, \dots, L_{t_{i+1}}$ and $k = 1, \dots, M_{t_{i+1}}$ are known. Hence, for each possible cumulative volume $(t_i, Q_{t_i}^k)$ with a spot price equal to $F_{t_i}^j$, we compute the value $\Psi_{t_i}(q, F_{t_i}^j, Q_{t_i}^k) + \mathbf{E}(P(t_{i+1}, F_{t_{i+1}}^k + q) | F_{t_i}^j)$, for any q in $\text{Adm}(t_i, Q_{t_i}^k)$, using the probabilities in the tree (see Table 1). The maximum over q gives $P(t_i, F_{t_i}^j, Q_{t_i}^k)$.

Then, iterating till time t_0 provides the price of the contract.

2.3.2 Numerical results

We give numerical results for the two cases described in Section 1. We also plot the optimal gas consumption with respect to a natural gas sample path. We observe a bang-bang behavior for the optimal consumption, as it is theoretically justified in Section 3.1.

Example 1: Swing contract. The values of parameters are set to those of Case 1. In the table below, both forward curves are considered. To have comparable results with different consumption frequencies, we adjust q_{\min} and q_{\max} . For instance, with a weekly consumption, we set $q_{\min} = 0$ and $q_{\max} = 6 \times 7$.

Consumption period	Flat forward curve	Real forward curve
7 days	2700	4439
3 days	2691	4587
2 days	2704	4607
1 day	2717	4611




Figure 5: Swing contract. Optimal gas decision for a natural gas sample path. On the left, with a flat forward curve. On the right, with the real forward curve of Figure 1.

Consumption period	Flat forward curve	Real forward curve
7 days	66.69	224.9
4 days	67.36	239
2 days	67.75	242.2
1 day	67.92	242.3

Example 2: storage contract. The values of parameters are set to those of Case 2.

For both cases, the price converge as the consumption frequency increases. The convergence is slower with the real forward curve because of the high seasonality. In the sequel, the above values are taken as references for the other numerical methods. We also note that the realized optimal consumption is often of bang-bang type (this is theoretically confirmed by Theorem 3.1).

2.4 Monte Carlo method using Longstaff-Schwartz methodology (2001)

2.4.1 Description of the algorithm

We need to simulate N_{MC} independent realizations of $F = (F_{t_i})_{0 \leq i \leq N}$. To do so and in view of (2.7), it suffices to simulate $(X_{t_i})_{0 \leq i \leq N}$, which can be done by writing $X_{t_i} = \sigma e^{-at_i} \sum_{j=0}^{i-1} g_j$, where $g_j = \int_{t_j}^{t_{j+1}} e^{\alpha u} dW_u$. Then, since $(g_j)_j$ are independent Gaussian random variables with a zero mean and a variance equal to $\int_{t_j}^{t_{j+1}} e^{2\alpha u} du = \frac{e^{2\alpha t_{j+1}} - e^{2\alpha t_j}}{2\alpha}$,




Figure 6: Storage contract. Optimal gas decision for a natural gas sample path. On the left, with a flat forward curve. On the right, with the real forward curve of Figure 1.

the simulation of $(X_{t_i})_{0 \leq i \leq N}$ is standard. Note that an extension to multi-factor models is straightforward with this approach (even if we do not present related results here).

We now explain how the algorithm is devised. Once the optimal consumption $q_{t_i}^{*,m}$ for the m -th path determined, the price is computed via a Monte Carlo approximation of the form

$$\frac{1}{N_{MC}} \sum_{m=1}^{N_{MC}} \left[\sum_{i=0}^{N-1} \Psi_{t_i}(q_{t_i}^{*,m}, F_{t_i}^m, Q_{t_i}^{*,m}) + \mathcal{P}_T(F_T^m, Q_T^{*,m}) \right].$$

For this, we first have to find at any given date t_i , the optimal strategy for all simulated paths $(F^m)_{1 \leq m \leq N_{MC}}$ and all admissible cumulative consumptions Q (i.e. $(t_i, Q) \in \mathcal{A}$). We therefore need to measure the quantity

$$R(t_i, F, Q + q) = \mathbf{E} \left(\sum_{j=i+1}^{N-1} \Psi_{t_j}(q_{t_j}^*, F_{t_j}, Q_{t_j}^*) + \mathcal{P}_T(F_T, Q_T^*) \mid F_{t_i} = F, Q_{t_{i+1}} = Q + q \right),$$

where $(q_{t_j}^*)_{j>i}$ denotes the optimal control at time t_j and $(Q_{t_j}^*)_{j>i}$ the corresponding cumulated consumption. This is the best expected return from date t_{i+1} till the end of the contract. Assume, for a while, that $R(t_i, F, Q + q)$ is numerically approximated by $\hat{R}(t_i, F, Q + q)$, using N_{MC} Monte Carlo simulations. Then the optimal strategy, at time t_i , for the m -th path and the admissible cumulative consumption Q is determined by:

$$q_{t_i}^*(F_{t_i}^m, Q) = \arg \max_{q \in \text{Adm}(t_i, Q)} [\Psi_{t_i}(q, F_{t_i}^m, Q) + \hat{R}(t_i, F_{t_i}^m, Q + q)].$$

Hence, by estimating the conditional expectation $R(t_i, F, Q + q)$ for each date and each possible cumulative consumption Q , the optimal exercise strategy is completely specified

along each path. Following the Longstaff-Schwartz methodology, we fix (t_i, Q) and approximate the random variable $R(t_i, F_{t_i}, Q)$ as the \mathbb{L}_2 -projection of $\sum_{j=i+1}^{N-1} \Psi_{t_j}(q_{t_j}^*, F_{t_j}, Q_{t_j}^*) + \mathcal{P}_T(F_T, Q_T^*)$, on the space spanned by a set of selected basis functions $[\Phi_l(F_{t_i})]_{0 \leq l \leq L}$. It leads to $\hat{R}(t_i, F, Q) = \sum_{l=0}^L \hat{\alpha}_l(t_i, Q) \Phi_l(F)$, where the coefficients $[\hat{\alpha}_l(t_i, Q)]_{0 \leq l \leq L}$ solve the least-squares problem

$$\arg \min_{(\alpha_l)_l} \sum_{m=1}^{N_{MC}} \left(\sum_{j=i+1}^{N-1} \Psi_{t_j}(q_{t_j}^{*,m}, F_{t_j}^m, Q_{t_j}^{*,m}) + \mathcal{P}_T(F_T^m, Q_T^{*,m}) - \sum_{l=0}^L \alpha_l \Phi_l(F_{t_i}^m) \right)^2$$

with an initial consumption $Q_{t_{i+1}} = Q + q$. Various methods are available to solve numerically this problem, depending on the volatility and the exercise frequency. For a general reference on least-squares problems, see Golub and Van Loan (1996). Iterating the procedure, we obtain backwards an approximation $q_{t_i}^*(F_{t_i}^m, Q)$ along all simulated paths and all possible cumulative consumptions. The price of the contract easily follows.

2.4.2 Numerical results

The pricing has been performed with 5000 simulations (plus antithetic), with the first four canonical polynomials ($\Phi(x) = x^i$, $0 \leq i \leq 3$) and with various consumption frequencies. We have not reported any result for periods smaller than 4 days, since the resolution of the least-squares problems usually generates instabilities for small initial times t_i .

Example 1: Swing contract. For the following results, the statistical errors due to

Consumption period	Flat forward curve	Real forward curve
15 days	2450	4410
10 days	2510	4450
5 days	2560	4480
4 days	2640	4510
Forest of trees (daily purchase)	2717	4611

Monte Carlo simulations are about 100 units at the 95% level. Hence, this approach gives good results.

Example 2: storage contract. The values of parameters are set to those of Case 2. We directly plot the results w.r.t. the frequency.

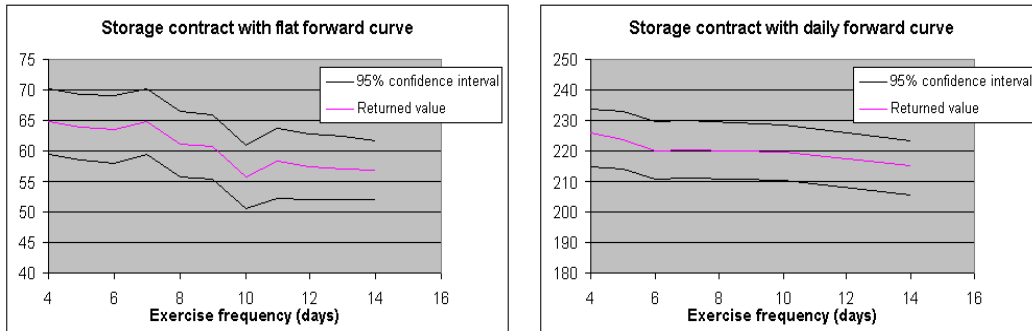


Figure 7: Flat forward curve on the left. Real forward curve on the right.

Concerning the functions basis, we have observed that the numerical values are quite insensitive to the number of polynomials from the moment that the degree 3 is achieved. The statistical error due to simulations limits any possible improvement (when we take more basis functions).

Moreover, in our numerical experiments (similarly to the forest method), it appears that the consumption is of bang-bang type. To prepare the connection with the further consumption parameterization approach, we plot on Figure 8 the optimal injection threshold for the storage contract (see Paragraph 3.3.1). Here, we take the flat curve rescaled at level 1 for convenience. The two first axes stand respectively for the time and the volume of stored gas. The surface shows the dependence of the gas price w.r.t the volume and the optimal time for injection. Note that the price for which we have to inject decreases as the




Figure 8:

volume of the stored gas increases. When the tank is empty, we inject at *any* cost. On the contrary, when the tank is almost full, the price must be much lower to inject some more gas. These observations are intuitively satisfying. Irregularities of the optimal frontier are due to statistical errors in the Monte Carlo algorithm. However, the overall shape seems quite smooth. Thus, modeling this optimal frontier with a smooth parameterization may lead to efficient procedures. This is developed in the sequel.

3 Parametric approximation

In this section, we develop numerical procedures based on parameterized consumptions. The purpose is to avoid using the dynamic programming equation (2.8), whose resolution can be much time consuming. The approach below leads to accurate estimations of $\sup_q \mathcal{J}(q)$, provided that the parameterization is relevant enough. To this end, we first determine some intrinsic properties of the optimal consumption, which will help in setting the parameterization.

3.1 Theoretical results: bang-bang consumption

Hereafter, we state a remarkable theoretical result concerning the necessary optimality conditions satisfied by the optimal consumption, in the Swing contract (with a penalty a little smoother than $\mathcal{P}_T(F, Q) = -AF [Q - Q_{\min}]_- - BF [Q - Q_{\max}]_+$). The proof is relegated to Appendix A. It states that the optimal consumption is of bang-bang type under some assumptions. This has been numerically observed with the procedures presented before.

Theorem 3.1 *Consider Problem¹(1.6) with $\Psi_{t_i}(q, F) = e^{-rt_i}q(F - K)$ and $\mathcal{P}_T(F, Q) = -e^{-rT}FP(Q)$, P being a continuously differentiable function.*

If the following condition holds

$$\mathbf{P} \left[e^{-rt_i} (F_{t_i} - K) + \mathbf{E} \left(e^{-rT} F_T P'(Q_T^*) | F_{t_i}, Q_{t_i}^* \right) = 0 \right] = 0, \quad (3.10)$$

¹For the storage contract, we have not been able to derive analogous rigorous results because of the lack of smoothness of $\Psi_{t_i}(q, F)$ w.r.t. q . However, formal computations show that the optimal consumption essentially takes three values $\{q_{\min}, 0, q_{\max}\}$.

the optimal consumption at time t_i is necessarily of bang-bang type and given by

$$q^*(t_i, F_{t_i}, Q_{t_i}^*) = q_{\max} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K) + \mathbf{E}(e^{-rT}F_T P'(Q_T^*)|F_{t_i}, Q_{t_i}^*) > 0} \\ + q_{\min} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K) + \mathbf{E}(e^{-rT}F_T P'(Q_T^*)|F_{t_i}, Q_{t_i}^*) < 0}.$$

Condition (3.10) appears to be hard to check since it involves the unknown optimal consumption. However, it can be expected (even if we have not been able to prove it) that for $t_i < T$, the law of the random variable $e^{-rt_i}(F_{t_i} - K) + \mathbf{E}(e^{-rT}F_T P'(Q_T^*)|F_{t_i}, Q_{t_i}^*)$ has a density w.r.t. the Lebesgue measure. Then, 0 will be a null point and Condition (3.10) will be satisfied.

Note that Theorem 3.1 is strongly related to the stochastic maximum principle (Pontryagin principle) for optimal control problems. However, this latter strictly holds in the continuous case (both in time and space). Since the cumulative consumption has a continuous state space, it has been possible to derive optimality conditions in our context.

Moreover, for less smooth penalties of the form $P(Q) = -A [Q - Q_{\min}]_- - B [Q - Q_{\max}]_+$, one can easily see that there are certain regions where the derivative of the terminal penalty does not depend any more on the whole consumption. The proposition below explicitly characterizes the consumption threshold in these special cases.

Proposition 3.2 *Consider the Swing contract described by (1.4).*

1. *If $Q_{t_i} \in [Q_{\min} - q_{\min} (N - i), Q_{\max} - q_{\max} (N - i)]$, the terminal constraints will be satisfied whatever happens ($Q_T \in [Q_{\min}, Q_{\max}]$). The consumption threshold is therefore equal to the strike K and the optimal consumption takes the following form (see Zone 1 in Figure 12)*

$$q_{t_i} = q_{\max} \mathbf{1}_{F_{t_i} > K} + q_{\min} \mathbf{1}_{F_{t_i} \leq K}.$$

2. *If $Q_{t_i} + q_{\min} (N - i) \geq Q_{\max}$, we are certain to be in over-consumption at the time horizon and $P'(Q_T) = -A$. Then using the relation $F(t_i, T) = \mathbf{E}(F_T|F_{t_i})$, the optimal consumption is in the form of (cf. Zone 2 in Figure 12)*

$$q_t = q_{\max} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K) - A e^{-rT}F(t_i, T) > 0} + q_{\min} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K) - A e^{-rT}F(t_i, T) \leq 0}.$$

3. If $Q_{t_i} + q_{\max} (N - i) \leq Q_{\min}$, we are certain to be in under-consumption at the time horizon and $P'(Q_T) = B$. The optimal consumption is given by (cf. Zone 3 in Figure 12)

$$q_t = q_{\max} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K)+B} e^{-rT} F(t_i,T) > 0 + q_{\min} \mathbf{1}_{e^{-rt_i}(F_{t_i}-K)+B} e^{-rT} F(t_i,T) \leq 0.$$

In order to improve optimization procedures, these specific cases have been taken into account in the parameterization of the gas consumption.

3.2 Using a neural network

3.2.1 Description and motivations

We first propose a nonlinear parameterization of the consumption function $q_t = q(t, F_t, Q_t)$, using a neural network (see Haykin (1994) for general references on neural networks). We have selected for our study a one-hidden-layer neural network, whose architecture is described in Figure 9. The input layer $(x_i)_{1 \leq i \leq 3}$ is connected to the state and time variables




Figure 9: Architecture of the neural network.

: $x_1 = t$, $x_2 = F_t$ and $x_3 = Q_t$. There is one hidden layer with L neurons, and the output layer returns the level of optimal consumption at time t . Actually, the number of hidden units to use is far from clear. It is related to the complexity of the underlying function that the network is trying to model. A good starting point is to choose a number of units equal to half the sum of the number of input and output units. We have conducted a certain number of experiments with different configurations and finally selected a network with 4 hidden neurons, in preference to a larger one with a negligible improvement. The network

is defined by a matrix and a vector of weights (the parameters) : the input weights $\{\theta_{kl}^{\text{in}}\}$ and the output weights $\{\theta_l^{\text{out}}\}$. The output of the network is given by

$$y = \sum_{l=1}^4 \theta_l^{\text{out}} s \left(\sum_{k=0}^3 \theta_{kl}^{\text{in}} x_k \right),$$

where $(\theta_{0l}^{\text{in}})_l$ are the bias weights (we set $x_0 = 1$) and s is the Sigmoid function: $s(x) = \frac{1}{1+e^{-x}}$. The local consumption is then defined by

$$q = q_{\min} + (q_{\max} - q_{\min}) s(y), \quad (3.11)$$

which still belongs to the range $[q_{\min}, q_{\max}]$.

The motivation of such a parameterization is the following. First, neural networks are very sophisticated techniques capable of modeling extremely complex functions (in particular, nonlinear ones with large numbers of variables), and at the same time, the level of the user knowledge needed to successfully apply neural networks is much lower than would be the case using, for example, some more traditional nonlinear statistical methods. In addition, the expected payoff $\mathcal{J}(q)$ defined in (1.6) becomes $\mathcal{J}(\theta)$, with $\theta = (\theta_{kl}^{\text{in}}, \theta_l^{\text{out}} ; 0 \leq k \leq 3 ; 1 \leq l \leq 4)$, and in the sequel the optimization w.r.t. θ will be performed using an estimation of the gradient $\nabla \mathcal{J}(\theta)$. This justifies why we here consider a smooth parameterization. However, thanks to Theorem 3.1, we may expect the optimal consumption to be of bang-bang type. This behavior can be seen as a limit case where the output weights tend to infinity. It means that y represents in a nonlinear way the bang-bang threshold : for large positive values of y , $q \approx q_{\max}$ and for large negatives values of y , $q \approx q_{\min}$.

3.2.2 The gradient algorithm

From an initial random configuration of weights, the stochastic gradient algorithm makes, in expectation, a move towards points of higher values in the search space, in order to locate a local maximum (see Benveniste, Metivier and Priouret (1990), Kushner and Yin (1997) for general references on stochastic approximation algorithms). We thus define a recursive sequence of parameters (θ_n) by setting $\theta_{n+1} = \theta_n + \gamma_n \widehat{\nabla_{\theta} \mathcal{J}(\theta_n)}$, where $\widehat{\nabla_{\theta} \mathcal{J}(\theta_n)}$ is a gradient estimator of the mapping $\theta \mapsto \mathcal{J}(\theta)$ and (γ_n) is a suitable sequence of

nonnegative step sizes. We have chosen to fix γ_n to a constant throughout our numerical experiments, but the difficult part is to decide how large the steps should be. Large steps may converge more quickly, but may also overstep the solution or go off in the wrong direction. In contrast, very small steps may go in the correct direction, but they also require a large number of iterations. The correct setting for the step size is application-dependent, and typically chosen by experiments. The algorithm stops in a high point, which is a local maximum. Throughout the numerical experiments we have performed, the values to which the algorithm has converged were all very close to the global optimum (presumably given by the method of Section 2.3). As for $\nabla_{\theta}\mathcal{J}(\theta)$, it has been computed by a classic Monte-Carlo approach, using a path-wise method (see Kushner and Yang (1991) and recently Gobet and Munos (2005)). It is thus expressed as an expectation :

$$\begin{aligned} \nabla_{\theta}\mathcal{J}(\theta) = & \mathbf{E} \left(\sum_{i=0}^{N-1} \nabla_{\theta}q(\theta, t_i, F_{t_i}, Q_{t_i}) \partial_q \Psi_{t_i}(q_{t_i}, F_{t_i}, Q_{t_i}) \right. \\ & \left. + \partial_{\theta} Q_{t_i} \partial_Q \Psi_{t_i}(q_{t_i}, F_{t_i}, Q_{t_i}) \right) + \mathbf{E} (e^{-rT} \partial_{\theta} Q_T \partial_Q \mathcal{P}_T(F_T, Q_T)). \end{aligned} \quad (3.12)$$

The quantity $\nabla_{\theta}q(\theta, t_i, F, Q)$ results in differentiating the parameterization of the consumption defined in (3.11), which is straightforward. For the Monte Carlo evaluations of $\nabla_{\theta}\mathcal{J}(\theta)$, we need to simulate $(\partial_{\theta}Q_{t_i})_{0 \leq i \leq N}$. It is obtained using the equation $Q_{t_{i+1}} = Q_{t_i} + q_{t_i}$, from which we deduce

$$\partial_{\theta}Q_{t_{i+1}} = \partial_{\theta}Q_{t_i} + \partial_{\theta}q(\theta, t_i, F_{t_i}, Q_{t_i}) + \partial_Q q(\theta, t_i, F_{t_i}, Q_{t_i}) \partial_{\theta}Q_{t_i},$$

with an initial condition $\partial_{\theta}Q_0 = 0$.

3.2.3 Numerical results

Example 1: Swing contract. As can be seen in section 3.1, we expect the optimal consumption to be of bang-bang type. Thus in order to take into account this (remarkable) result, a threshold function can be additionally applied to the output unit :

$$q_{\text{threshold}} = q_{\text{max}} \mathbf{1}_{q_{\text{network}} > \frac{q_{\text{max}} + q_{\text{min}}}{2}} + q_{\text{min}} \mathbf{1}_{q_{\text{network}} \leq \frac{q_{\text{max}} + q_{\text{min}}}{2}}.$$

In the above expression, q_{network} and $q_{\text{threshold}}$ stand for the output neuron and the threshold value, respectively. Here, the parameters are the ones given in Case 1. We also consider




Figure 10: Price of the Swing contract w.r.t. the number of iterations. On the left, with a flat forward curve. On the right, with the forward curve of Figure 1.

	Neural network before thresholding	Neural network after thresholding	Forest of trees
Flat forward curve	2659 ± 28	2681 ± 28	2717
Real forward curve	4555 ± 31	4584 ± 31	4611

Table 2: Final price of the Swing contract, before and after thresholding the consumption (the values after \pm equal the half of the 95% confidence interval width).

$A = B = 1$. The step size for the stochastic-gradient algorithm is constant and equal to 3. Denote by

- N_{MC} the required number of Monte Carlo simulations to calculate an estimator of the gradient, for a given set of parameters,
- N_{IT} the number of iterations on the parameters in the optimization stage.

The algorithm can therefore be split into two stages:

- 1) a first optimization stage, where we set for example $N_{MC} = 100$ and a large number of iterations $N_{IT} = 3000$, so that the algorithm presumably converges.
- 2) a second pricing stage with the last set of parameters and a much higher N_{MC} (it can be set equal to 100000 for example), in order to reduce the confidence interval

on the price.

Results are reported in Table 2. Note that for both forward curves, the results are far by 1% from the optimal price computed by the tree method, and thresholding at the exit of the neural network seems to be relevant. Furthermore, we can see that these two tests provide powerful insights concerning the type of convergence that may be observed in the optimization algorithms (independently of the selected parameterization). For the first forward curve, we can see that the algorithm requires only 1000 iterations to converge, and thus it is useless to go until 3000 iterations. On the contrary, for the second forward curve, we have a slower optimization, which is even performed in two steps (as if a local optimum were reached at an intermediate stage). The quality of results and the type of convergence are thus strongly related to the selected forward curve.

Example 2: storage contract. The values of parameters are set to those of Case 2. The step size for the stochastic-gradient algorithm is constant and equal to 0.1. The terminal penalties are always satisfied in this case. In fact, as already mentioned, the constraints in the storage case can be formulated as follows

$$f_1(t) \leq Q_t \leq f_2(t), \forall t \in [0, T].$$

Denote by $q_{params}(t)$ the output of the neural network at time t .

- $Q_t \geq f_1(t)$ for every $t \in [0, T]$. Thus as soon as $|Q_t - f_1(t)| \leq |q_{min}| dt$, we replace the output of the network by :

$$q_t^1 = \begin{cases} q_{params}(t), & \text{if } q_{params}(t) + Q_t dt \geq f_1(t) \\ \frac{f_1(t) - Q_t}{dt}, & \text{otherwise.} \end{cases}$$

- $Q_t \leq f_2(t)$ for every $t \in [0, T]$. Thus if $q_t^1 dt + Q_t > f_2(t)$, we replace q_t^1 by $\frac{f_2(t) - Q_t}{dt}$.

Finally, the retained control at time t takes the form of

$$q_t = \min \left[q_t^1, \frac{f_2(t) - Q_t}{dt} \right] = \min \left[\max \left(q_{params}(t), \frac{f_1(t) - Q_t}{dt} \right), \frac{f_2(t) - Q_t}{dt} \right].$$

Note that the above consumption is not of class C_1 any more. However, the stochastic gradient algorithm is still applicable, since there is only a finite set of points where the function is not differentiable.

	Neural network	Forest of trees
Flat forward curve ($N_{IT} = 3000$)	64.7 ± 0.4	67,9
Real forward curve ($N_{IT} = 6000$)	218.7 ± 0.6	242,3

Table 3: Final price of the storage contract (the values after \pm equal the half of the 95% confidence interval width).

The results reported in Table 3 are computed with $N_{IT} = 3000$ (resp. $N_{IT} = 6000$) iterations for the flat forward curve (resp. for the forward curve from Figure 1), still using $N_{MC} = 100000$ paths for the final value. Note that the neural network gives a less




Figure 11: Price of the storage contract w.r.t. the number of iterations. On the left, with a flat forward curve. On the right, with the forward curve of Figure 1.

satisfying price in the case of the real forward curve. The strong seasonality component of the curve is undoubtedly an important explaining factor. Adding more hidden neurons certainly improves the quality of the price, but also implies longer computing times.

Let us finally stress that no thresholding at the output of the network has been applied here. Unlike the Swing contract, we have not been able in this case to find a relevant thresholding rule at the output of the network.

3.3 Direct parameterization of the purchase threshold

3.3.1 For the Swing contract

Unlike the previous approach with neural networks, here the parameterization of $q(t_i, F, Q)$ does not need to be smooth with respect to parameters. We rather focus on the bang-bang behavior which can be expected from Theorem 3.1. Thus, we here assume that

- $q^*(t_i, F_{t_i}, Q_{t_i})$ takes only two values q_{\min} and q_{\max} ,
- $q^*(t_i, F_{t_i}, Q_{t_i})$ is an increasing function w.r.t F_{t_i} .

The second property is heuristic : we expect that a higher price of gas will lead to a higher consumption at a fixed price K . Previous algorithms based on the dynamic programming equation (see Section 2) give results that confirm this heuristic. From these two properties we deduce that if q is optimal, there exists a threshold S^* depending on both t_i and Q_{t_i} such that

$$q_{t_i} = q_{\max} \mathbf{1}_{\{F_{t_i} > S^*(t_i, Q_{t_i})\}} + q_{\min} \mathbf{1}_{\{F_{t_i} \leq S^*(t_i, Q_{t_i})\}}.$$

Our approach is to model this optimal threshold price as a function of t and Q_t . We can divide the set of couples (t_i, Q_{t_i}) into four zones.

- Zone 1 corresponds to a consumption such that the consumer is sure to respect the global constraints. The optimal threshold is equal to the strike K .
- Zone 2 (resp.3) corresponds to an over (resp. under) consumption. In these two cases the optimal threshold is an explicit function of t_i . If the global penalties are infinite, these two zones are useless.
- Zone 4 is the only one where the threshold has to be found.

Even if F_{t_i} fluctuates a lot, we expect that the threshold S^* changes a little with respect to t_i and Q_{t_i} . For instance, in the special case where F_{t_i} is deterministic with $r = 0$, S^* is constant. Consequently, $S^*(t_i, Q_{t_i})$ may be approximated by a smooth function of (t_i, Q_{t_i}) . Figure 8 about the injection optimal decision confirms that the threshold surface may be




Figure 12: Parameterization of the threshold as a function of both (t, Q_t) and $(t - \frac{Q_t - q_{\min}}{q_{\max} - q_{\min}}, Q_t - q_{\min} \times t)$.

approximated by a smooth one for these contracts. Actually, it is more efficient to express S^* as a function of

$$\left(i - \frac{Q_{t_i} - q_{\min} \times i}{q_{\max} - q_{\min}}, Q_{t_i} - q_{\min} \times i \right),$$

the support of the function S^* being now a rectangle. The parameters θ of the function S^* are values of the threshold at a fixed position spread, on a uniform grid with a chosen size (5×5 for example). The value of S^* outside the grid is equal to the strike if the couple (t_i, Q_{t_i}) is in Zone 1, and to a bilinear interpolation of the four nearest neighbors, otherwise.

3.3.2 Optimization of parameters

With this parameterization, the function $\mathcal{J}(\theta)$ to maximize is not smooth, and a specific optimization procedure has to be carried out.

Since parameters θ are threshold values, they can be initialized near the value K . The idea of the optimization is the following. With a given set of simulations, we can evaluate $\hat{\mathcal{J}}(\theta)$ for any value of θ . We are looking for the optimal parameter by applying coordinate-wise iterative perturbations of the parameters. The size of the perturbation equals ε and parameters are iteratively chosen according to a statistical test based on the empirical mean and variance of the improvement. To get better and better parameters we increase the size of the set of simulations in several stages. We start with $l = 0$ and $N_{MC} = 10000$. We define $\theta_{j_0}^+ = (\theta_1, \dots, \theta_{j_0} + \varepsilon, \dots, \theta_d)$ and $\theta_{j_0}^- = (\theta_1, \dots, \theta_{j_0} - \varepsilon, \dots, \theta_d)$. The algorithm can be described as follows.

- Stage l
 - Simulate a set of $2^l N_{MC}$ independent paths $(F_{t_i})_i$.
 - Denote by $\hat{\sigma}(X)$ the empirical standard deviation of X computed with this set of paths.
 - While $\hat{\mathcal{J}}(\theta)$ is improved
 - * Compute the empirical mean of $\hat{\mathcal{J}}(\theta)$ using this set of paths.
 - * For each parameter coordinate θ_{j_0}
 - if $\hat{\mathcal{J}}(\theta_{j_0}^+) - \hat{\mathcal{J}}(\theta) > 1.65 \times \hat{\sigma}(\mathcal{J}(\theta_{j_0}^+) - \mathcal{J}(\theta))$ or
 - if $\hat{\mathcal{J}}(\theta_{j_0}^-) - \hat{\mathcal{J}}(\theta) > 1.65 \times \hat{\sigma}(\mathcal{J}(\theta_{j_0}^-) - \mathcal{J}(\theta))$
- change the value of θ_{j_0} and continue the optimization with the best of these three sets of parameters.
- $l \rightarrow l + 1$.

Iterations are stopped when $2^l N_{MC}$ is too large for computations or when there is no more improvement.

3.3.3 Applications to storage contract

A similar strategy can be used for the storage contract, since we expect that the optimal consumption $q(t, Q_t, F_t)$ takes only three values. In this case we build two surfaces associated to the two thresholds, using the same method. We choose a size of grid and change the values of nodes, one by one, at each optimization step: nodes of the first surface, then nodes of the second one. Since the previous change of variables is useless, we parameterize the two surfaces with t_i and Q_{t_i} .

3.3.4 Results

Example 1: Swing contract. For different forward curves, the solution and the optimal parameters computed with this algorithm do not depend on the initial set of parameters. Furthermore, no local maximum phenomenon has been noticed. The quality of results mainly depends on the forward curves. With a 4×4 grid we can reach, in the best cases,

98% of the performance obtained by a dynamic programming approach. The following results are computed with $\varepsilon = 0.1$, $l = 5$ steps (corresponding to a maximum set of 320 000 different simulations) and 95% confidence intervals.

Forward curve	Size of grid	$\hat{\mathcal{J}}$	Forest of trees
real	4×4	4565 ± 25	4611
real	10×10	4578 ± 25	4611
flat	4×4	2696 ± 23	2718
flat	10×10	2692 ± 23	2718

Example 2: storage contract. These results mainly depend on the choice of ε and the size of grids. ε needs to be small enough to have reliable results but a too small value of ε may also generate heavy computations. In our case grids were too coarse to use a very small value of ε . Adaptive grids, which become more precise as t grows, may be more efficient than uniform ones. In the storage contract, a grid that ensures continuity of the threshold surface at edges of Zone 1, gives good results as well. It appears that initialization and order of parameters θ_j in the optimization procedure are not crucial.

Forward curve	Size of grid	$\hat{\mathcal{J}}$	Forest of trees
real	4×4	225.3 ± 1.2	242.5
real	10×10	230 ± 1.2	242.5

4 Conclusion

In the paper, we have presented several numerical methods to price Swing options. The well-known forest of trees method allows to deal efficiently with one factor models. We have extended the Longstaff-Schwartz method (initially used for Bermuda options) to our framework. This authorizes multi-factor models. These two procedures take advantage of the dynamic programming equation and thus, lead to accurate valuations of the Swing contract. Numerical experiments confirm this feature. We develop alternative algorithms based on relevant parameterizations of the consumption. These new methods can also handle multi-factor models. Moreover, unlike the previous ones, they have the advantage to give intermediate prices throughout the optimization stage.

To guide our choice of parameterization, we have proven that the optimal purchase is

of bang-bang type. The first parametric method is based on a neural network, which can asymptotically reproduce the bang-bang behavior. The objective function is smooth w.r.t. parameters and gradient stochastic methods can be used for the optimization. The results are quite satisfying. The second parametric method directly models the bang-bang threshold as a function of the time and the cumulative consumption. This seems to be particularly adapted to the problem and thus leads to better results.

A Proof of Theorem 3.1

Let $q(\cdot)$ be an optimal control and $\bar{q}(\cdot)$ an admissible control in $[q_{\min}, q_{\max}]$. For $\epsilon \in [0, 1]$, we define a third admissible control

$$q_{t_j}^\epsilon = \begin{cases} q_{t_i} + \epsilon(\bar{q}_{t_i} - q_{t_i}) & j = i, \\ q_{t_j} & j \neq i. \end{cases}$$

The objective function corresponding to q^ϵ is given by

$$\mathcal{J}(q^\epsilon) = \mathbf{E} \left(\sum_{j=0}^{N-1} e^{-rt_j} q_{t_j}^\epsilon (F_{t_j} - K) + e^{-rT} F_T P(Q_T^\epsilon) \right).$$

Since q is optimal, we get $\mathcal{J}(q^\epsilon) \leq \mathcal{J}(q)$, for every $\epsilon \in [0, 1]$, and therefore $\frac{\partial \mathcal{J}(q^\epsilon)}{\partial \epsilon} \Big|_{\epsilon=0} \leq 0$.

We can easily compute this derivative and obtain

$$\begin{aligned} \frac{\partial \mathcal{J}(q^\epsilon)}{\partial \epsilon} \Big|_{\epsilon=0} &= \lim_{\epsilon \downarrow 0} \frac{\mathcal{J}(q^\epsilon) - \mathcal{J}(q)}{\epsilon} \\ &= \mathbf{E} \left\{ (\bar{q}_{t_i} - q_{t_i}) \left[e^{-rt_i} (F_{t_i} - K) + e^{-rT} F_T P'(Q_T) \right] \right\} \\ &= \mathbf{E} \left\{ (\bar{q}_{t_i} - q_{t_i}) \left[e^{-rt_i} (F_{t_i} - K) + \mathbf{E} (e^{-rT} F_T P'(Q_T) | F_{t_i}, Q_{t_i}) \right] \right\} \leq 0. \end{aligned}$$

The above inequality holds true for all admissible control $(\bar{q}_{t_j})_j$, thus in particular for

$\bar{q}_{t_j} = q_{t_j} + (q - q_{t_i}) \mathbf{1}_{\lambda_{t_i} > 0} \mathbf{1}_{i=j}$, where

$$H_{t_i} = e^{-rt_i} (F_{t_i} - K) + \mathbf{E} (e^{-rT} F_T P'(Q_T) | F_{t_i}, Q_{t_i})$$

and $\lambda_{t_i} = (q - q_{t_i}) H_{t_i}$ ($q \in [q_{\min}, q_{\max}]$ being fixed). This gives $\mathbf{E} \left(\mathbf{1}_{\lambda_{t_i} > 0} \lambda_{t_i} \right) \leq 0$. In other words, λ_{t_i} is negative outside some negligible sets. Thus we have proven the following optimality condition

$$(q - q_{t_i}) \left[e^{-rt_i} (F_{t_i} - K) + \mathbf{E} (e^{-rT} F_T P'(Q_T) | F_{t_i}, Q_{t_i}) \right] \leq 0, \forall q \in [q_{\min}, q_{\max}],$$

and the desired result follows at once.

References

- [1] A. Benveniste, M. Metivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, New York, 1990.
- [2] D.P. Bertsekas. *Dynamic programming and optimal control. Vol. 1 and 2*. Belmont, MA: Athena Scientific, 1995.
- [3] D.P. Bertsekas and S.E. Shreve. *Stochastic optimal control*, volume 139 of *Mathematics in Science and Engineering*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1978. The discrete time case.
- [4] L. Clewlow, C. Strickland, and V. Kaminski. Valuation of Swing contracts. Technical report, Energy Power Risk Management, Risk Waters Group, July 2001.
- [5] G. Golub and C.F. Van Loan. *Matrix computations*. 3rd ed. Baltimore, MD: The Johns Hopkins Univ. Press. xxvii, 694 p. , 1996.
- [6] E. Gobet and R. Munos. Sensitivity analysis using Itô-Malliavin calculus and martingales. Application to stochastic control problem. *SIAM Journal on Control and Optimization*, 43(5):1676–1713, 2005.
- [7] S. Haykin. *Neural Networks: A Comprehensive Foundation*. McMillan, New York, 1994.
- [8] P. Jaillet, E.I. Ronn, and S. Tompaidis. Valuation of commodity-based Swing options. *Management Science*, 50:909–921, 2004.
- [9] I. Karatzas. On the pricing of American options. *Applied Mathematics and Optimization*, 17(1):37–60, 1988.
- [10] H.J. Kushner and J. Yang. A Monte Carlo method for sensitivity analysis and parametric optimization of nonlinear stochastic systems. *SIAM J. Control Optim.*, 29(5):1216–1249, 1991.

- [11] H.J. Kushner and G. Yin. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, Berlin and New York, 1997.
- [12] A. Lari-Lavassani, M. Simchi, and A. Ware. A discrete valuation of Swing options. *Canadian Applied Mathematics Quarterly*, 9(1):35–74, 2001.
- [13] F. Longstaff and E.S. Schwartz. Valuing american options by simulation: A simple least squares approach. *The Review of Financial Studies*, 14:113–147, 2001.
- [14] S. Raïkar and M. Ilić. Interruptible physical transmission contracts for congestion management. *Power engineering society summer meeting, 2001, IEEE*, 3:1639–1646, 2001.
- [15] A.C. Thompson. Valuation of path-dependent contingent claims with multiple exercise decisions over time: the case of Take-or-Pay. *Journal of Financial and Quantitative Analysis*, 30:271–293, 1995.