

LUND UNIVERSITY
FACULTY OF ENGINEERING

MASTER'S THESIS

Pricing swing options in the
electricity market

Author:
Caroline Olofsson

Supervisor:
Magnus Wiktorsson

*A thesis submitted in fulfilment of the requirements
for the degree of Master of Science*

in the

Mathematical Statistics
Centre for Mathematical Sciences

April 20, 2015

Abstract

The thesis deals with how to price swing options in the electricity market by using a least squares Monte Carlo method. This is a simulation method which uses a backwards moving algorithm where the optimal decision is calculated at every time step. Regression is used for the optimal decision and in this thesis both a polynomial regression and a cubic smoothing spline are used. They are both shown to be rather good estimators for the regression.

Two variation of contracts are priced. For the first only one exercise right that can be used when exercising and for the second one several exercise rights can be used when exercising. Volume restrictions are also used. The algorithm implemented in this thesis give similar results to the ones of previous authors and when we can not compare with other authors it give us results fairly close to our expectations.

The thesis also examines the optimal exercise strategy for a swing option and the boundaries for when to use an exercise right are calculated.

Acknowledgments

This master's thesis was written at the department for Mathematical Statistics at Lund University during the autumn semester of 2014/2015.

I would like to thank my supervisor Magnus Wiktorsson at Lund University for the valuable feedback and contributions to this thesis and for presenting this interesting topic to me.

I also want to thank my family and friends for their support in the process of writing my thesis.

Contents

1	Introduction	4
1.1	Electricity Markets	4
1.2	Swing Options	4
1.3	The aim of the thesis	6
2	Theory	7
2.1	Spot price process	7
2.2	Former research of spot price models	7
2.3	Former research of pricing swing options	9
3	Least Squares Monte Carlo method	11
3.1	The LSM algorithm	11
3.2	Pseudo code LSM	12
3.3	Example	12
3.4	Extended LSM	15
3.5	Pseudo code XLSM	15
3.6	Extended LSM example	17
3.7	Variation of swing option contracts	20
3.7.1	Contract 1	20
3.7.2	Contract 2	20
3.8	Matlabs built-in function	21
3.9	Early exercise strategy	21
4	Computational results	23
4.1	Comparing with existing algorithms	23
4.2	Contract 1	23
4.3	Contract 2	27
4.4	Early exercise strategy	29
5	Conclusion and Outlook	33
5.1	Conclusion	33
5.2	Outlook	34
	Bibliography	34

Chapter 1

Introduction

1.1 Electricity Markets

In the 90s the electricity market started to get deregulated in some countries and now many more have followed. In a regulated market the spot price for electricity and prices for different contracts are set by the regulators. In a deregulated market the prices is set by demand and supply. Determining the price process of electricity is difficult since electricity cannot be stored and the demand is uncertain and inelastic. Since the electricity prices are volatile it is good for a buyer of electricity to hedge against the risk of the price going up. Therefore it is common to use a forward contract with a predetermined price and quantity. The seller of the contract also gets a decrease in risk since they are certain to sell an amount for a fixed price over an amount of time. This makes it easier for them to invest in new facilities. Swing options can be seen as a risk management tool since it lets the buyer have flexibility in the contract by buying less or more quantity and the buyer can choose between different opportunities to use it. Buyers usually do not know when they might need less or more volume and since electricity can not be stored, a swing option is a good hedging tool.

1.2 Swing Options

A swing option is an exotic derivative with multiple exercise rights, which means that the holder of the option can exercise more than once during the lifetime of the option. Swing options are mainly found in the energy sector for purchasing oil, gas and electricity. In this thesis we are only focusing on the electricity market. Swing options are traded over the counter (OTC) and for electricity there are electricity exchanges (e.g. Nord Pool, APX and NYMEX). A swing option can be settled either with physical delivery or financially. Like an American option it is possible to exercise at any point in time before the maturity, although it is common that these are

discrete points in time (like a Bermudan option), which is also assumed in this thesis. Usually can only one exercise right at most be exercised at each exercise opportunity.

It is common that the swing option is written in conjunction with a set of forward contracts where the option holder has contracted to buy a certain amount of volume at a certain price over a certain amount of time, a so called base load agreement. Due to variability in consumption the holder might want to buy more or less volume at certain time points. The holder of the swing option can then swing up (buy more volume) or down (sell volume) at the strike price K . K is set either at the initialization of the contract or at a future date before the duration of the contract. The strike price can be the same during the whole contract or different for every exercise opportunity (Jaillet, Ronn and Tompaidis (2004)).

Contracts specify over a time interval when the delivery of electricity is to take place, not at a specific time point. For example a contract that is written at time 0 and is specified over $[T_1, T_2]$ with $0 \leq T_1 \leq T_2$. Then the exercise opportunities are $T_1 \leq \tau_1 < \tau_2 < \dots < \tau_n \leq T_2$ where n is the number of exercise opportunities and we have for the number of exercise rights N that $N \leq n$.

When using an exercise right, there is usually a local volume restriction for how much the buyer is allowed to swing up or down at each exercise opportunity. There is usually also a global volume restriction for how much the buyer is allowed to swing up or down during the lifetime of the swing option.

Assuming one unit can be bought with each exercise right and there are no global constraints, we can get the upper and lower bound for the option (Jaillet, Ronn and Tompaidis (2004)). The upper bound is given by N identical Bermudan options. The price of the swing option is lower since it is restricted to only exercising the rights on different opportunities whereas the Bermudan options can all be exercised on the same opportunity. The lower bound is given by a set of European options where the exercise opportunities are calculated as the optimal N exercise opportunities among the set of all possible exercise opportunities. Since these the exercise points are pre-determined the price of a swing option with it's flexible exercise opportunities is higher. For the extremes $N = 1$ and $N = n$, if the swing option only has one exercise right, it has the same price as a Bermudan option. If the swing option has as many exercise rights as exercise opportunities, then the price is the same as a string of European options with each and every exercise opportunity as expiration date.

In this thesis two variation of swing option contracts are implemented. They are based on the Least squares Monte Carlo (LSM) algorithm that Dörr (2003) described which is an extension of the LSM algorithm that Longstaff and Schwartz (2002) used to price American and Bermudan options with. For the first contract only one right can be exercised at each opportunity.

There is a local volume restriction but no global volume restriction. For the second contract there is no restriction on the number of exercise rights. Instead the global volume restriction is discretized into volume blocks but work in the same way as the exercise rights did in the first contract. The difference is also that more than one volume block can be exercised at each opportunity (up to the local volume restriction, which is also discretized).

In this thesis no penalty function is assumed but it can be used. With a penalty function it would be possible to exercise beyond the volume restrictions, although it would result in a penalty. We also assume that it is not specified how many of the exercise rights are reserved for upswings or downswings. The algorithm can for each exercise right choose to swing up or down depending on which will optimize the value. This is also the case for the second contract when we use global volume instead of exercise rights.

1.3 The aim of the thesis

The aim of this thesis is to use a realistic spot price model for electricity spot prices to determine the price of a swing option. Since swing options are mostly traded over the counter it is more difficult to find the specifics of swing contracts and their prices. Therefore in this thesis we only compare with other articles on swing options. In order to do that, spot price models of other authors are used to compare our algorithm with the ones of the other authors. It is desirable to know when to use the exercise rights and therefore the spot price boundaries for when to exercise are also investigated. The programming language used for this thesis is MATLAB.

The method chosen to estimate the price is the extended Least Squares Monte Carlo method, which is a simulation method. Other methods which can be used to price swing options are for example dynamic programming approaches or finite differences to calculate the price.

Chapter 2

Theory

2.1 Spot price process

It is concluded that the spot price of electricity exhibit these characteristics among others (Dörr (2003)):

Seasonal pattern: Since the use of electricity is not the same throughout the year, the spot price process has a seasonal pattern. For countries with very cold winters, the electricity is at its highest in the winter. For countries with warm summers, electricity is also high in summer due to the use of air-conditioning. There are also seasonal variations over the week with high use during weekdays and low use during the weekends. This is due to that many industries, which use a lot of electricity, are closed on weekends. Also during the day there is a change in the use of electricity, it is lower during the night when people sleep and many industries are closed.

Mean-reversion: Electricity prices tend to revert to a mean value whenever it starts to drift. Prices tend to increase when there is a the demand increases and when the supply decrease. This will be adjusted as the supply will be increased and the price will go down. The same if the demand is low and the supply is high the price will go down and the supply will decrease and thus rendering the price to increase again.

Price spikes: Positive price spikes occur when the current demand approaches the maximum supply. This is usually due to a failure in the production or distribution of electricity. The price will increase rapidly and when the supply has increased the price will fall back to normal level, since this will occur within a short amount of time the price process will get the shape of a spike. Negative price spikes do also occur, but are not as common.

2.2 Former research of spot price models

Since the deregulation of the energy market a variety of research on how to model the spot price process and how to price derivatives has been done.

Lucia and Schwartz proposed in 2000 a model for the spot price process which is a stationary mean-reverting process (an Ornstein-Uhlenbeck process) with seasonal pattern. Hambly, Howison and Kluge proposed in 2005 a model that is an extension of Lucia and Schwartz model with the addition of spike prices process. Deng also proposed in 2000 a mean-reverting jump diffusion models but without a seasonality pattern. Instead he took the volatility into account in his models. Cartea and Figueroa (2005) also used a one-factor mean-reversion jump diffusion model which is adjusted to incorporate seasonality effects but it differs from Hambly, Howison and Kluge model in the sense that Hambly, Howison and Kluge used an independent jump process.

The model presented by Hambly, Howison and Kluge (2009) for the spot price process S use the characteristics presented in 2.1. It contains a deterministic periodic function $f(t)$ which should represent the seasonality, an Ornstein-Uhlenbeck (OU) process X and a mean-reverting process Y which exhibits jumps which represents price spikes.

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t, \end{aligned} \tag{2.2.1}$$

where α and β are independently the speed of mean-reversion, σ is the volatility, W_t is a standard Brownian motion, N_t is a Poisson process and the independent identically distributed process J_t represents the jump size.

The spot price process used in this thesis is mainly the one described in Hambly, Howison and Kluge (2009), but some others are applied in order to try the accuracy of the algorithm, namely Meinshausen and Hambly (2004) and Dörr (2003).

Meinshausen and Hambly use an AR(1) model for the logarithm of the energy price where the price S_t on day t is simulated with:

$$\log S_t = (1 - k)(\log S_{t-1} - \mu) + \mu + \sigma e_t, \tag{2.2.2}$$

where e_t is a standard normally distributed random variable with unit variance and k , μ and σ are constants.

Dörr uses among others a spot price model he calls the one-factor model:

$$S_t = F(t)e^{Y_t} \tag{2.2.3}$$

$$dY_t = -\alpha Y_t dt + \sigma(t)dW_t \tag{2.2.4}$$

where α is the mean-reversion factor, $F(t)$ is a deterministic mean reversion level, $\sigma(t)$ is the deterministic volatility and W_t denotes a standard Brownian motion. For the case where σ is constant it is discretized to:

$$Y_{t+1} = Y_t e^{-\alpha h} + \sigma \sqrt{\frac{(1 - e^{-2\alpha h})}{2\alpha}} e_{t+1}, \quad (2.2.5)$$

where e_{t+1} again is a standard normal variable.

The model of Hambly, Howison and Kluge is discretized to:

$$\begin{aligned} X_{t+1} &= e^{-\alpha h} X_t + \sigma e_{t+1} \\ Y_{t+1} &= e^{-\beta h} Y_t + e^{-\beta h(1-u)} J_{t+1} \\ J_{t+1} &= -\log(1 - \tilde{u}) (\mathbb{1}_{\tilde{u} < p} (\frac{1}{n_1} + \frac{1}{n_2}) - \frac{1}{n_2}) \end{aligned} \quad (2.2.6)$$

where u , \tilde{u} and \bar{u} are independent uniformly distributed random variables, e_{t+1} is again a standard normal variable, h is an increment of time and n_1 and n_2 are the mean size of the jumps inverted.

An example of simulated spot price processes with Hambly, Howison and Kluges model is seen in figure 2.1. The parameters used are the ones that they use in their article, i.e. $K = 1$, $\alpha = 7$, $\beta = 200$ and $\sigma = 1.4$. The below figure is done with 180 time points occurring daily. Starting points X_0 and Y_0 are zero and no seasonality is assumed.

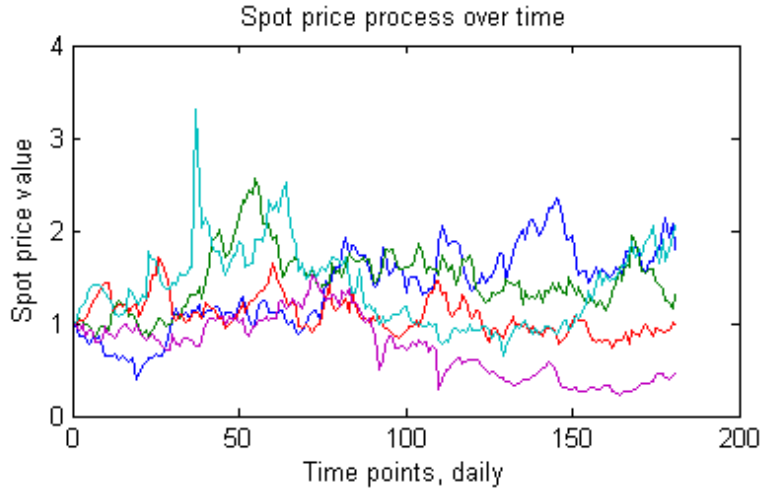


Figure 2.1: Spot price processes over time using the model of Hambly, Howison and Kluge, without seasonality.

2.3 Former research of pricing swing options

Various methods has been used for pricing swing options. Carriere (1996) presented a simulation algorithm which uses regression to approximate conditional expectations that can be used to price American options. This is similar to the work of Longstaff and Schwartz (2002) who used a least square

Monte Carlo based approach to pricing American options. Dörr (2003) extended this approach to swing options. Cartea and Williams (2008) also uses an extension to the Longstaff and Schwartz (2002) algorithm to price swing options.

Keppo (2004) modeled electricity swing options using regular electricity derivatives. He showed that a swing option can be replicated by using regular electricity forwards and call options.

Jaillet, Ronn and Tompaidis (2004) used a dynamic programming approach for the pricing of swing options, more precisely they used a binomial/trinomial forest which is an extended version with several layers of the traditional binomial/trinomial tree used to price options with only one exercise right. As the spot price model they used a one-factor model with mean-reverting process with seasonality.

Hambly, Howison and Kluge (2009) used a grid method to price swing options. The grid method uses approximations to the conditional density of the spot process. The grid method is a modification of the method used by Jaillet, Ronn and Tompaidis but Hambly, Howison and Kluge use a grid instead of a trinomial tree. The downside with this method was that the speed of the algorithm was not satisfactory.

Chapter 3

Least Squares Monte Carlo method

The least squares Monte Carlo method (LSM) uses the optimal stopping rule to decide when to exercise the rights of the option. The disadvantage with this is that we can practically never reach the optimal stopping rule, which means that the price is just a lower bound of the true price (Cartea and Williams (2008)). The advantage is that the method is not model dependent which means that a multi-factor model is as easily priced as a simpler model.

In order to understand the contracts presented in this thesis we first look at the method presented by Longstaff and Schwartz (2002) and then the extension of the method that Dörr (2003) presented.

3.1 The LSM algorithm

The method developed by Longstaff and Schwartz (2002) for American and Bermudan options uses least square regression to decide whether to exercise the option or keep it. This decision is made for each simulated spot price path and at each time step as the algorithm moves backwards in time. In order to make the decision we need to determine the value in continuing with the option versus the cash flow received when exercising the option at the current time point. We can express the continuation value as the expectation of the future payoffs conditional on the information up to the current time point. We are now looking at an option that can only swing in one direction.

Since we use a backwards moving algorithm we start from the last exercise opportunity. If the option has not been exercised already, it is always optimal to exercise if it is in the money at the last exercise opportunity. We move backwards and at each time point evaluate whether the value of exercising is greater than the value of continuation, which we get by using

least squares regression. We regress the discounted future cash flows realized from continuing onto a finite set of basis functions of our values for the spot price. The regression is done by using the values from all of the paths. The set of the basis functions for the regression is in this thesis polynomial regression but it could also be Legendre or Chebyshev for example. The regression can also be done using a smoothing spline, which we will look at later. The continuation value for a path w with values X at time t_{K-1} is (Dörr (2003)):

$$F(X, t_{K-1}) = \sum_{j=0}^{\infty} a_j(t_{K-1}) B_j(X) \quad (3.1.1)$$

where X is the price, a_j are coefficients and B_j is the set of basis functions. For example if we have a second order polynomial regression then: $F(X, t_{K-1}) = a_0(t_{K-1}) + a_1(t_{K-1})X + a_2(t_{K-1})X^2$.

A pseudo code for the algorithm is given next, followed by a practical example of how the algorithm works. The option used is an American put option (downswing) with payoff if $K > S_t$ for each time t .

3.2 Pseudo code LSM

For this pseudo code a polynomial regression is used with MATLABs built-in functions "polyfit" and "polyval". The cash flow matrix C is where all the cash flows are saved and M is the number of trajectories, C is the cash flow matrix where all the cash flows from exercising are saved.

3.3 Example

In order to explain the principle we here give a small example with is set up as the example of Longstaff and Schwartz (2002). For a deeper understanding the reader is referred to their example.

In this example we use 10 simulated trajectories with starting point 1.0 at $t=0$ and maturity at t_4 . We will here only show 5 of the trajectories, as more trajectories will not add to the understanding of the principle. The spot price process used is 2.2.1 but without the jump process. The strike price is $K=1$ and we assume a risk-free rate of 5%. A polynomial regression is used of order 2. The spot prices are:

Algorithm 1 Longstaff and Schwartz algorithm

```

1: set  $C_i^n = \max(K - S_i^n, 0)$  for  $i=1, \dots, M$ 
2: for  $t = (n - 1) \rightarrow 1$  do
3:   for  $i = 1 \rightarrow M$ , where  $S_i^t < K$  do
4:      $X_i^t \leftarrow S_i^t$ 
5:      $Y_i^t \leftarrow$  future values of C discounted back to t
6:   end for
7:    $pol^t = \text{polyfit}(X^t, Y^t, 3)$ 
8:   for  $i = 1 \rightarrow M$ , where  $S_i^t < K$  do
9:      $continuation_i^t = \text{polyval}(pol, X_i^t)$ 
10:     $exercise_i^t = K - S_i^t$ 
11:    if  $exercise_i^t > continuation_i^t$  then
12:       $C_i^t = exercise_i^t$ 
13:       $C_i^m = 0$ , where  $m > t$ 
14:    end if
15:  end for
16: end for
17:  $V_i \leftarrow$  values of C discounted to  $t_0$  over all  $t$  for  $i=1, \dots, M$ 
18: OptionValue =  $\frac{1}{M} \sum_{i=1}^M V_i$ 

```

Spot price paths

Path	t=0	t=1	t=2	t=3	t=4
1	1.0000	0.9388	0.8681	0.8322	0.9264
2	1.0000	1.0057	1.0080	1.0644	1.0860
3	1.0000	0.9151	0.9544	0.9420	0.9568
4	1.0000	0.9219	1.0006	1.0677	1.1976
5	1.0000	0.9995	1.1189	1.0558	0.9945
⋮	⋮	⋮	⋮	⋮	⋮

We start at t_4 and exercise the option for the trajectories that are in the money:

Cash flow matrix at time t=4

Path	t=1	t=2	t=3	t=4
1	—	—	—	0.0736
2	—	—	—	0
3	—	—	—	0.0432
4	—	—	—	0
5	—	—	—	0.0055
⋮	⋮	⋮	⋮	⋮

Then we move on to the t_3 . Let X denote the spot price at t_3 and Y denote the cash flows received after t_3 discounted back to t_3 , for the trajectories that are in the money (here 1 and 3). We then have:

Regression basis at time t=3		
Path	Y	X
1	0.0736×0.9512	0.8322
2	—	—
3	0.0432×0.9512	0.9420
4	—	—
5	—	—
\vdots	\vdots	\vdots

Let us use the basis functions $B_1 = a_1X^0$, $B_2 = a_2X^1$ and $B_3 = a_3X^2$, i.e. basic polynomial regression of order 2. The conditional expectation function is then $\mathbb{E}[Y|X] = 4.06681 - 8.70601X + 4.71905X^2$. The continuation values we get when inserting the X-vector into the function are:

Early exercise decision at t=3		
Path	Continuation	Early exercise
1	0.0899	0.1678
2	—	—
3	0.0533	0.0580
4	—	—
5	—	—
\vdots	\vdots	\vdots

This means that it is optimal to exercise for both the first and third trajectory. The cash flow matrix now looks like:

Cash flow matrix at time t=3				
Path	t=1	t=2	t=3	t=4
1	—	—	0.1678	0
2	—	—	0	0
3	—	—	0.0580	0
4	—	—	0	0
5	—	—	0	0.0055
\vdots	\vdots	\vdots	\vdots	\vdots

Note that the cash flows for trajectory 1 and 3 will move from t_4 to t_3 . We only have one exercise right so there can only be one cash flow received for each trajectory.

This goes on in the same manner until t_1 . The cash flow matrix at point t_1 now looks like this:

Path	t=1	t=2	t=3	t=4
1	0	0	0.1678	0
2	0	0	0	0
3	0.0849	0	0	0
4	0.0781	0	0	0
5	0.0005	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots

We have now filled the cash flow matrix and can calculate the option price by discounting all cash flows back to t_0 and then taking the average over all paths. The price in this case is 0.0611. We will now go into the details of the extended LSM that Dörr (2003) describes. Both up- and downswing will now be used which means that all the trajectories will be in the money.

3.4 Extended LSM

When using the extended LSM (XLSM) we need to add one more dimension, i.e. number of exercise rights left to use, to our already existing dimensions for paths and time. We can let this extra dimension to be represented by levels, one for each number of exercise rights left to use. The algorithm work in a similar way to the one by Longstaff and Schwartz (2002). The difference is that when comparing whether to exercise or not, we need to compare keeping the option as it is (staying at a certain level) or exercising one right and continuing with one less exercise right (one level lower). We also need to consider that when we decide to exercise a right for a certain path, the subsequent cash flows for that path need to be substituted to the subsequent cash flow of the option with one exercise right less than the original one. This becomes obvious when considering that if we would not do that then payoffs could be received at more time points than we have exercise rights. The substitution need to be done before the level under is evaluated since the level under might exercise one right at the current time point and the subsequent cash flow we substitute to will be one too few.

3.5 Pseudo code XLSM

For this pseudo code we use the same functions and notations as before, but now we also have exercise rights, where N is the number of exercise rights for the option.

Algorithm 2 Extended LSM

```
1: for  $J = 1 \rightarrow N$  do
2:   for  $t = n \rightarrow n - (J - 1)$  do
3:     set  $C_i^{tJ} = \max(K - S_i^t)$  for  $i=1, \dots, M$ 
4:   end for
5: end for
6: for  $t = (n - 1) \rightarrow 1$  do
7:   for  $i = 1 \rightarrow M$  do
8:      $X_i^t \leftarrow S_i^t$ 
9:     for  $J = 1 \rightarrow \min(n - t, N)$  do
10:       $Y_i^{tJ} \leftarrow C_i^{mJ}$  discounted back to  $t$ , where  $m > t$ 
11:    end for
12:     $exercise_i^t = K - S_i^t$ 
13:  end for
14:  for  $J = 1 \rightarrow \min(n - t, N)$  do
15:     $pol^{tJ} = \text{polyfit}(X^t, Y^{tJ}, 3)$ 
16:    for  $i = 1 \rightarrow M$  do
17:       $continuation_i^{tJ} = \text{polyval}(pol^{tJ}, X_i)$ 
18:    end for
19:  end for
20:  if  $\min(n - t, N) > 1$  then
21:    for  $J = \min(n - t, N) \rightarrow 2$  do
22:      for  $i = 1 \rightarrow M$  do
23:        if  $exercise_i^t + continuation_i^{t^{J-1}} > continuation_i^{t^J}$  then
24:           $C_i^{tJ} = exercise_i^t$ 
25:           $C_i^{mJ} = C_i^{m^{J-1}}$ , where  $m > t$ 
26:        end if
27:      end for
28:    end for
29:  end if
30:  if  $exercise_i^t > continuation_i^{t^1}$  then
31:     $C_i^{t^1} = exercise_i^t$ 
32:     $C_i^{m^1} = 0$ , where  $m > t$ 
33:  end if
34: end for
35: for  $J = 1 \rightarrow N$  do
36:    $V_i^J \leftarrow$  values of  $C$  discounted to  $t_0$  over all  $t$  for  $i=1, \dots, M$ 
37:    $OptionValue^J = \frac{1}{M} \sum_{i=1}^M V_i^J$ 
38: end for
```

3.6 Extended LSM example

We use the same trajectories for the price as before. The difference is that we now have 3 exercise rights instead of one, although we can still only exercise one right at each opportunity. They all have the same strike price $K=1$. In this example we now also assume that we can both swing up (buy at the price K) and swing down (sell at the price K), which means that the payoff is now $|K - S_t|$. This example follows the same structure as the one that Dörr used. This example is however shorter and we will still only show 5 trajectories. For a more elaborate example the reader is referred to Dörr (2003).

We start again from the last time point. The difference is that for level three we can in the initial step exercise not only at T but at $T-1$ and at $T-2$. This is because when we reach $T-2$ with all three exercise rights left it is optimal to use them all at the last three time points. The same goes for level two, where we can exercise the remaining two rights at time $T-1$ and T . The cash flow matrix after the initial step then is:

Path	t=1	t=2	t=3	t=4
1	–	0.1319	0.1678	0.0736
2	–	0.0080	0.0644	0.0860
3	–	0.0456	0.0580	0.0432
4	–	0.0006	0.0677	0.1976
5	–	0.1189	0.0558	0.0055
⋮	⋮	⋮	⋮	⋮

Path	t=1	t=2	t=3	t=4
1	–	–	0.1678	0.0736
2	–	–	0.0644	0.0860
3	–	–	0.0580	0.0432
4	–	–	0.0677	0.1976
5	–	–	0.0558	0.0055
⋮	⋮	⋮	⋮	⋮

1 exercise left				
Path	t=1	t=2	t=3	t=4
1	—	—	—	0.0736
2	—	—	—	0.0860
3	—	—	—	0.0432
4	—	—	—	0.1976
5	—	—	—	0.0055
⋮	⋮	⋮	⋮	⋮

This means that at time t_3 , the cash flow matrix is already done for level two and three and they will remain the same. We only need to update level one, which can be calculated almost the same as in our first example above, but now all trajectories are in the money. The cash flow matrix for level one after using LSM is then:

1 exercise rights left, cash flow at t=3				
Path	t=1	t=2	t=3	t=4
1	—	—	0.1678	0
2	—	—	0	0.0860
3	—	—	0	0.0432
4	—	—	0.0677	0
5	—	—	0	0.0055
⋮	⋮	⋮	⋮	⋮

At time t_2 we still don't need to update level three since we already have the cash flows for that time point. We want to evaluate level one and two so we need the spot price vector X and the vector Y_1 for the cash flows at level one discounted back to t_2 and Y_2 for the cash flows at level two discounted back to t_2 .

Regression			
Path	Y_2	Y_1	X
1	$0.1678 \times 0.9512 + 0.0736 \times 0.9512^2$	0.1678×0.9512	0.8681
2	$0.0644 \times 0.9512 + 0.0860 \times 0.9512^2$	0.0860×0.9512^2	1.0080
3	$0.0580 \times 0.9512 + 0.0432 \times 0.9512^2$	0.0432×0.9512^2	0.9544
4	$0.0677 \times 0.9512 + 0.1976 \times 0.9512^2$	0.0677×0.9512	1.0006
5	$0.0558 \times 0.9512 + 0.0055 \times 0.9512^2$	0.0055×0.9512^2	1.1189
⋮	⋮	⋮	⋮

Each Y is regressed separately onto X and gives two different vectors of continuation values. The values are shown in the table below. Also shown is the values we would receive for level two when choosing to exercise one

right and continue with only one right left. For level one we need to choose between exercising our only exercise right or continuing with it.

Early exercise decision at $t=2$

Path	Continue 2	Exercise 1 & Continue 1	Continue 1	Exercise 1 & stop
1	0.2788	0.2997	0.1678	0.1319
2	0.0856	0.0376	0.0297	0.0080
3	0.1326	0.1085	0.0629	0.0456
4	0.0901	0.0334	0.0328	0.0006
5	0.0942	0.1571	0.0382	0.1189
\vdots	\vdots	\vdots	\vdots	\vdots

For the trajectories at level two where it is optimal to exercise we take the cash flows for the same trajectories from level one. When we decide to exercise for level 1 for a certain trajectory we substitute the previous cash flow for the new one, as in our first example. As mentioned earlier we need to exercise the highest level first.

Note that the cash flow when exercising trajectories for level two is not the values under "Exercise 1 & Continue 1". If we choose to exercise 1 and continue with only one right left the cash flow will be from exercising one right, i.e. the same cash flow that we would get for those trajectories in level 1 if we would choose to exercise.

So for t_2 the cash flows for 1 and 2 exercise rights left are now:

2 exercises left

Path	t=1	t=2	t=3	t=4
1	—	0.1319	0.1678	0
2	—	0	0.0644	0.0860
3	—	0	0.0580	0.0432
4	—	0	0.0677	0.1976
5	—	0.1189	0	0.0055
\vdots	\vdots	\vdots	\vdots	\vdots

1 exercise rights left

Path	t=1	t=2	t=3	t=4
1	—	0	0.1678	0
2	—	0	0	0.0860
3	—	0	0	0.0432
4	—	0	0.0677	0
5	—	0.1189	0	0
\vdots	\vdots	\vdots	\vdots	\vdots

The same is then done for t_1 , but then also for level 3. For level 3 we need to compare continuing with three exercise rights or using one and continuing with two exercise rights. When we have all the cash flows it is time to calculate the price. When calculating the price for the option with 3 exercise right, we use the third level cash flow table. If we want the price of an option with 2 rights we do not need to rerun the algorithm, we only need to look at the cash flow table for level 2. This is because exercise rights are not affected by what happens in the exercise levels above. This means that if we have calculated the price for a certain number of exercise rights, we can easily obtain the price for an option with less exercise rights. The price of the swing option is 0.2678, 0.1975 and 0.1080 for three, two and one exercise rights.

3.7 Variation of swing option contracts

The two different variation of swing option contracts used in this thesis is described below.

3.7.1 Contract 1

The first contract used to price swing options in this thesis is the extended LSM method described above. There is a local volume restriction but no global one. The local volume is assumed to be maximized always, which is known as "bang bang". It means that when faced with the opportunity to execute all local volume it is also optimal to do so. In this thesis, for contract 1 it is assumed that this will always happen and therefore the algorithm is written so that it will always use all local volume. This is intuitive, since for those exercise opportunities where the most payoff can be made, we want to use as much volume as possible.

3.7.2 Contract 2

In the second contract we use volume blocks for our exercise levels instead of exercise rights. The volume blocks are our global volume restriction discretized into a number of exercise levels. The global volume is the accumulated deviation from the contracted volume for the whole time period, where at each time period a certain deviation is allowed, which is the local volume. Since we want to be able to exercise more than one volume block at each opportunity (up to the local volume restriction) we can now exercise more than one level at a time. The local volume is now not assumed to always be maximized, although we are expecting it to happen. In this contract the local volume restriction is also discretized (with the same factor as the global volume restriction) and the algorithm will choose how many volume blocks

will be exercised at each opportunity. Dividing the global volume into very small blocks will cost a lot of time so we are dividing into bigger volume blocks in this thesis. It is still a realistic contract since when purchasing volume it is often sold in bigger volume units. The volume restrictions can be expressed as:

$$V_{global} \geq \sum_{t=1}^T |V_{contracted,day_t} - V_{delivered,day_t}| \quad (3.7.1)$$

$$V_{local} \geq |V_{contracted,day_t} - V_{delivered,day_t}| \quad (3.7.2)$$

An example is a global volume restriction of 1000 units over 20 days and a local volume restriction of 300 units per day. The discretization factor is 100 units which means that we have 10 exercise levels and for every day we can chose to exercise 0 ,1, 2 or 3 levels.

3.8 Matlabs built-in function

Matlab has a built in function called "**hswingbyls**" which calculates the price of a swing option. The function uses a minimum and maximum local volume but no restriction on global volume. This algorithm also uses the bang bang principle for the optimal early exercise strategy. The regression can be done using interpolating spline or by using polynomials of the third order. This function does the same as our contract 1.

3.9 Early exercise strategy

The regression technique used in this thesis is a simple polynomial regression. A desirable aspect which can then easily be obtained when running the algorithm is to be able to get the threshold of prices for when it is optimal to exercise, the early exercise strategy. This is done by saving the polynomials received during the regression and then solving the equation we get when we want the s where the continuation value of staying at the current exercise level J is equal to the value of exercising L number of levels and continuing on the level $J-L$:

$$f(t, S, J) = f(t, S, J - L) + L\Phi(S), \quad (3.9.1)$$

where $f(t, S, J)$ is the same as in 3.1.1 and has the shape $a_0 + a_1s + a_2s^2 + a_3s^3 + \dots$ and $\Phi(S)$ depend on if it is the boundary for upswings or downswing we want to find, i.e. for upswings we have $\Phi(S) = S - K$ and for downswing $\Phi(S) = K - S$. Depending on the degree of the polynomial we get different amount of roots. We only want to use the non-imaginary roots. If more than one non-imaginary root exist we choose the one closest to K . When

we have the s for each point in time and for each L for a certain number of exercise rights, we can create a figure that show how the boundary changes with time and for different L for that amount of exercise rights.

Chapter 4

Computational results

4.1 Comparing with existing algorithms

We first start with comparing our algorithm with the ones that Dörr (2003) and Meinshausen and Hambly (2004) used. In their reports the only time we can compare to their algorithm is when they only used upswings.

Dörr uses 10 time steps, where $\Delta_t = 1$, $K = 20$, $\alpha = 0.5$ and 0.05 , the mean reversion level $F=20.7387$, $\sigma=0.392$ and 6 exercise right (only upswings) are used. Different spot prices are used to depict how the option value changes with the spot price. Two different α are used. Dörr does not write out the values but looking at a graph taken from Dörr (2003) in figure 4.2 it seems that the values we get with our algorithm in figure 4.1 are pretty much the same.

Meinshausen and Hambly set the values of the constants to be: $\sigma = 0.5$, $k=0.9$ and $\mu=0$ with $K=0$ and $S_1 = 1$. They use $T = 1000$ and up to $n=100$ exercise opportunities. $B_1 = a_1X^0$ and $B_2 = a_2X^1$ are the basis functions used. The results show that our algorithm gives the same results as theirs.

4.2 Contract 1

From now on we will use the spot price model of Hambly, Howison and Kluge (2009) with $S_0 = 1$, $K = 1$, $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, 10 time points (daily) and no seasonality, if not stated otherwise.

For our first contract we have a local volume restriction but not a global volume restriction. We can compare this contract to the built-in function in MATLAB. Our algorithm gives the same result as the built-in algorithm in MATLAB if both up- and downswings are allowed, even if the regression function is not always the same. If not both up- and downswings are allowed the results will differ about 0,6 % for $M=100000$, which is an acceptable result. The underlying trajectories, number of swing rights and exercise opportunities are of course set the same.

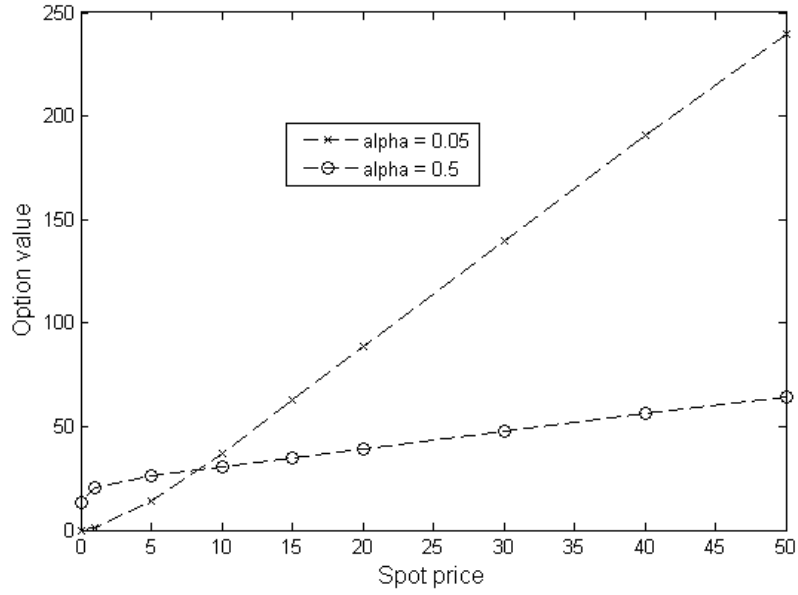


Figure 4.1: The values of the option price with our algorithm, as a function of the spot price. Two different α are used.

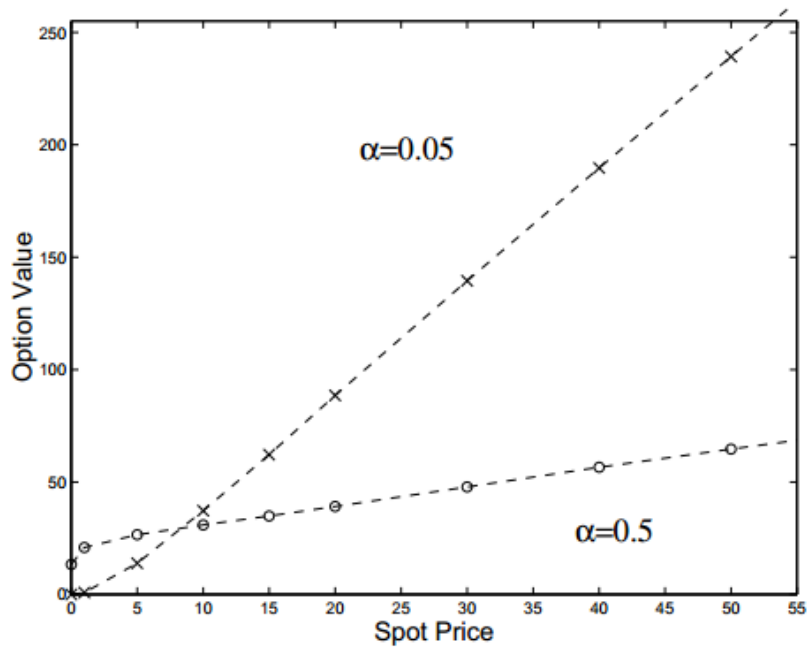


Figure 4.2: The values of the option price with Dörrs algorithm, as a function of the spot price. Two different α are used.

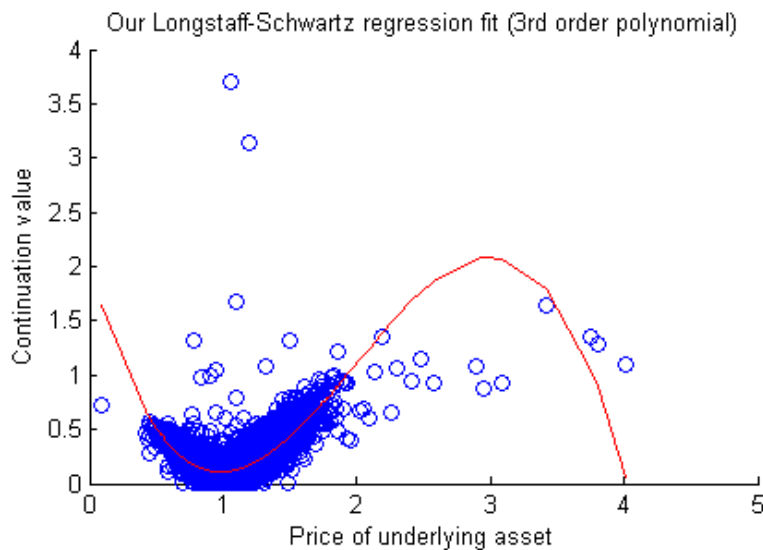


Figure 4.3: Regression between the underlying price and the continuation value with the discounted future cash flows as blue scatter, at the last exercise date before maturity using a third order polynomial

An example of how the regression looks like is shown in figure 4.3. 10000 trajectories are used and the number of exercise rights are 5. The volume when swinging up or down is of the unit one. A third order polynomial is used for the regression.

The figure shows blue scatter with the spot prices on the x-axis and the future cash flows discounted back to time $T-1$ on the y-axis. The red line is the conditional expectation function with continuation values as a function of the spot price.

The last exercise date before maturity is used because this should be when the trajectories are the most spread, and hence the regression is better. The closer we get to the initial time t_0 the spot prices of the trajectories will be around the value 1 (S_0) and closer to each other, which means that the regression will not be as good. This could be prevented by assuming a starting point that is not t_0 but rather an arbitrary point in time. We will use this later.

The function seems to fit the points rather well, which would mean that the regression is probably accurate. Intuitively one could think that a polynomial of order two would be a good fit, by looking at the shape of the curve the scatter plots create. This is not the case though. With a second order polynomial the curving around the central points of interest, i.e. in this case around 1, is too flat. A third order polynomial is used instead, since it will give a better fit around the central points. However, the function will then have the problem that it will not be a good fit around

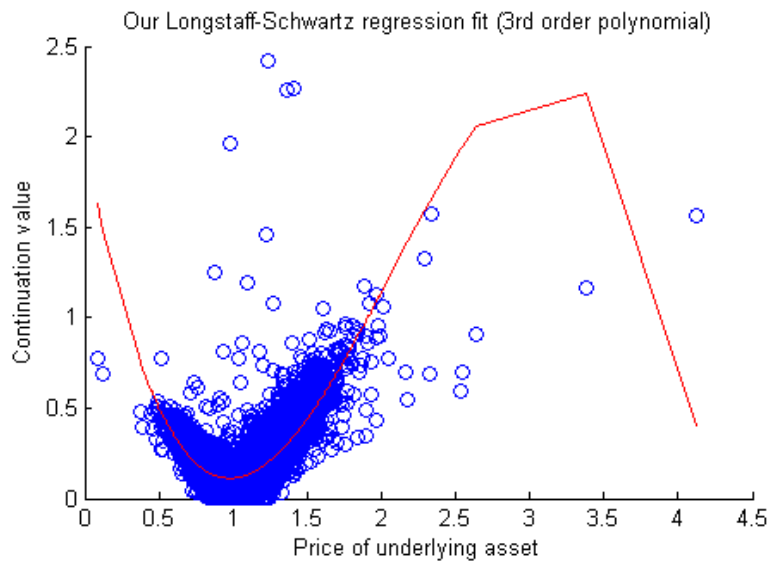


Figure 4.4: Regression between the underlying price and the continuation value with the discounted future cash flows as blue scatter, at the last exercise date before maturity using a third order polynomial

the extreme values of the spot price, which can be seen in figure 4.3 and 4.4. This is not a concern though since at the extreme values of the spot price, the option will usually be exercised and then it does not matter if the continuation value is not accurate. We will now use a spline instead and look at the results.

As we can see in figure 4.5 the regression looks much better with the use of a spline. The settings are the same as in the previous two figures, except that now we use a spline for the regression. The regression with a spline is usually better at the central points of interest. Although sometimes the spline will wobble around these points. This can be prevented to some extent by changing the smoothing parameter, which in our case is set to 0.7. For a more smooth curve we can decrease this parameter, but then the curve will be further away from the scatter points.

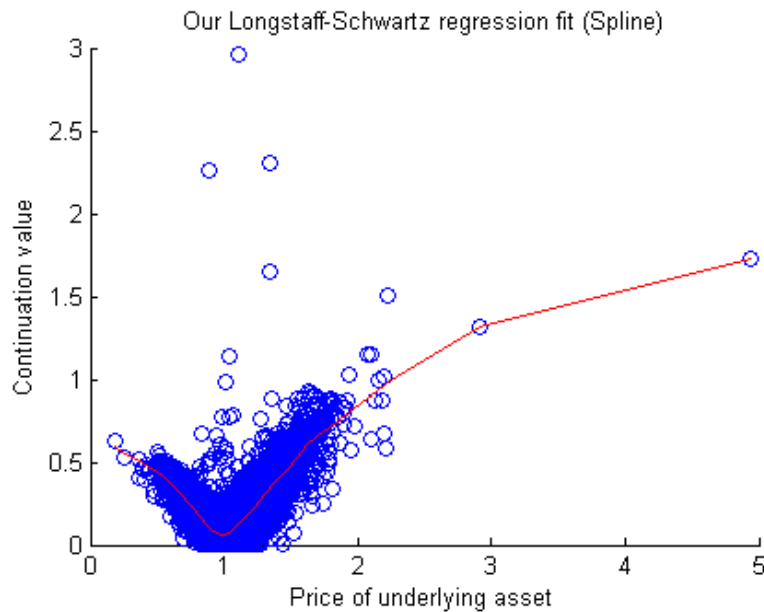


Figure 4.5: Regression between the underlying price and the continuation value with the discounted future cash flows as blue scatter, at the last exercise date before maturity using a spline

The MATLAB function "csaps" is used by both our algorithm and the built-in function in MATLAB, which means that the figure looks the same for both algorithms. The results will in contrast differ a bit. If we use a spline we get an approximate 0.1% deviation when both up- and downswings are allowed and approximately 0.3 % if only one of up- or downswings are allowed, which are both acceptable results.

Another thing to test is to see if the algorithm uses up all of the exercise rights. Since both up and downswings are possible at all exercise opportunities and therefore for each opportunity there is always a positive payoff to be made. So it is always optimal to use up all exercise rights. This can easily be calculated and the result shows that the algorithm indeed uses all the exercise rights.

4.3 Contract 2

In the second contract we are also interested in the price of the swing option but the difference from the previous section is that we do not have other work to compare with since we have designed this contract ourselves. But we can adjust the parameters in contract 2 to make it equal to contract 1 and compare the results, i.e. by letting only one level be exercised at a time in contract 2 and having as many exercise levels in contract 2 as exercise

rights in contract 1. The results show that the values are the same down to the tenth significant digit.

Like with contract 1 we are interested to see if all global volume (compared to exercise rights in the previous section) and all local volume is used, which it should. All global volume is not always used, but almost. This means that there is a deviation from what we expect. This is for when the local volume restriction for up and down are the same. When they are not the deviation is bigger. When only up- or downswings are allowed, the global volume is not always used up, which is reasonable since a profit is made by swinging in only one direction and so the global volume used will depend on the spot price process.

It is optimal to use all local volume when exercising, but our algorithm does not always do that. One example can be seen below in figure 4.6, where the red dots represent an exercise of two levels at the same time and the blue dots represent an exercise of one level. The blue dots are moved 0.1 time units to the left in order for them to be more visual. The global volume has four levels, two can be exercised at most at each time point (both up- and downswings). As we can see the local volume is maximized most of the time, but sometimes just one level is exercised (in this case four times).




Figure 4.6: Local volume used, red dots represent an exercise of two levels at the same time and the blue dots represent an exercise of one level. Total amount of exercise levels are four. The blue dots are moved 0.1 time units to the left.

Figure 4.7 show an example with the same settings as above but now the global volume has five exercise level. As we can see, the exercise of only one level is more frequent which is to be expected since when we now have

five exercise levels, which is an odd number.




Figure 4.7: Local volume used, red dots represent an exercise of two levels at the same time and the blue dots represent an exercise of one level. Total amount of exercise levels are 5. The blue dots are moved 0.1 time units to the left.

In the figures we can see at which spot prices the exercises are made. This is an indicator of the early exercise boundary which we will show the results for next.

4.4 Early exercise strategy

We are interested to find the optimal exercise values for the spot price at each time point. By using the technique in section 3.9 we can find the empirical values for the boundaries. The results are depicted in figures. In all the figures used in this section a swing option with seven exercise levels and local restriction of three levels exercised at the same time (both up and down) is used. The settings of the spot price process is the same as before.

Figure 4.8 shows the lines drawn between the prices where it is optimal to swing up if the spot price is above the line or swing down if the spot price is below the line. At t_{10} which is outside the figure we already know that if we have exercise levels left then we should use them. What the figure tells us is that at t_9 we should also use all the levels we can with the local restriction. At t_8 we should always use one exercise right.

The idea with the early exercise boundaries are to be able to determine when to exercise the option. In the figure, two trajectories are also included.

If we use the boundaries in the figure, it would mean that for trajectory 1, 3 levels should be exercised at t_7 , 1 at t_8 and 3 at t_9 . For trajectory 2 we should exercise 3 levels at t_6 , 3 levels at t_7 and one level at t_8 (since we will only have one left then).

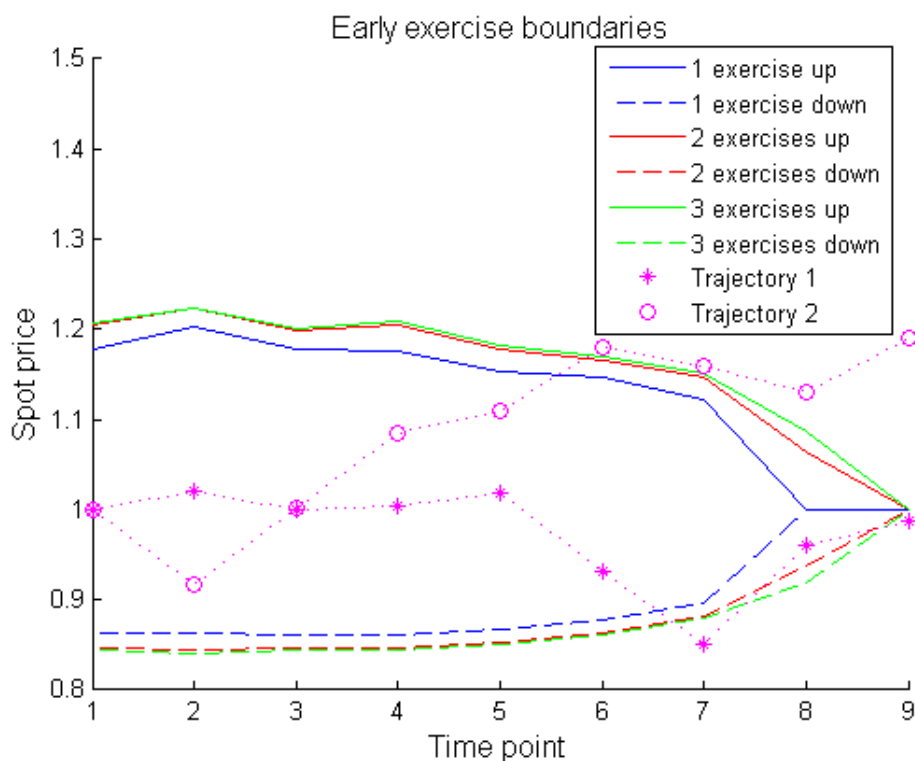


Figure 4.8: Early exercise boundaries for the spot price in the form of lines through time, global volume is 7 exercise levels and local volume is 3 exercise levels (both up and down). Trajectory 1 and 2 are examples of spot price processes.

For this spot price model (Hambly, Howison and Kluge) the lines are not smooth. This is due to the existence of price spikes. The lines below are smoother than the ones above. This is because the values of the spot price process are sometimes extreme, i.e. far away from K , for prices above K . This will give a less accurate estimate for the above price lines. The prices below K are bounded by zero, so the spot prices are there not so extreme. The price lines above are further away from K than the ones below, which is because the spot price process is moving upward.

The line for using 3 exercises should be further from K than the one with using 2 exercises and so on but this is not always the case. The lines lie relatively close together which is reasonable considering that when it is optimal to exercise one or two levels it is usually optimal to exercise 3 levels

(if possible).

In order to get a smoother lines we can use the trick we mentioned earlier. We assume that another point than t_0 is the starting point for the trajectories, to make them more spread out. Since they are independent and time homogeneous it does not matter at which point in time we start our algorithm. This is done by creating a longer spot price process but only using the last time points, e.g. 10 out of 30 time points. When using the trick we get figures (e.g. figure 4.9) that look like before and after several iterations we conclude that there is not much difference. This is probably due to the fact that the spot price processes are already well dispersed without this trick, when including jumps. When jumps are not included the figure will become too volatile at the first time steps to be of any use. This is because the values of the underlying price distribution is not as dispersed as for the case including jumps. In order to fix this we can use the trick.

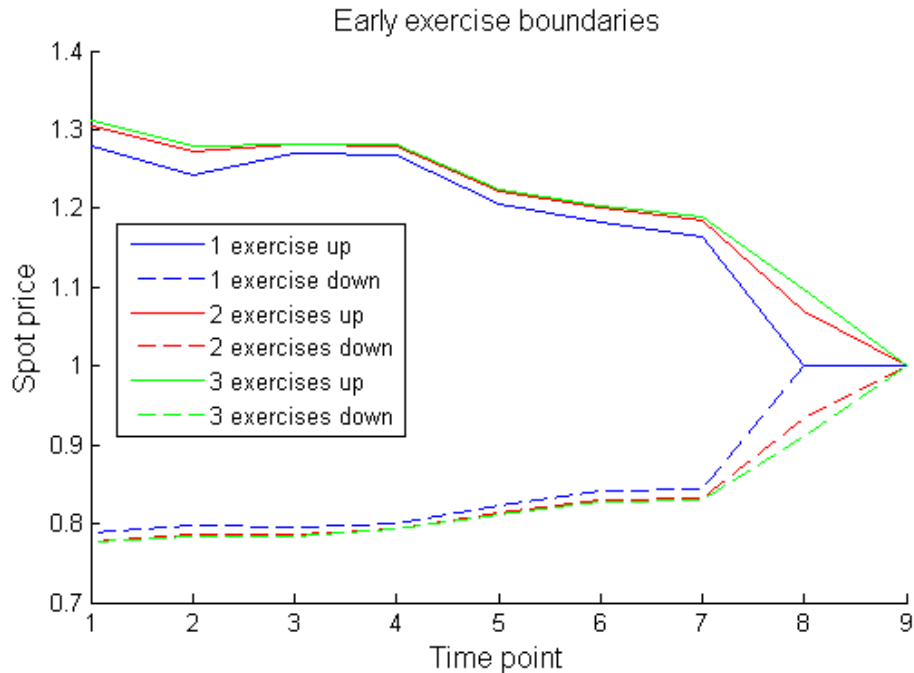


Figure 4.9: Early exercise boundaries for the spot price in the form of lines through time, global volume is 7 exercise levels and local volume is 3 exercise levels (both up and down). Here the trick to use later time points is used.

Figure 4.10 shows the early exercise boundaries with the model of Hamby, Howison and Kluge but without price spikes. Now we have used our trick. As we can see the lines are now much smoother, which they should be.

We have now looked at two different regression types and the impact

of volume restrictions when pricing swing options. We have also seen how the optimal exercise strategy can be found. It is now time to make some conclusions of our findings.




Figure 4.10: Early exercise boundaries for the spot price in the form of lines through time, global volume is 7 exercise levels and local volume is 3 exercise levels (both up and down). Here the trick to use later time points is used and there are no jumps in the spot price process.

Chapter 5

Conclusion and Outlook

5.1 Conclusion

In this thesis we have examined how to price a swing option with two different contracts and when the optimal time to exercise is.

Our algorithm seem to work as well as Dörr (2003) and Meinshausen and Hambly (2004). It also seem to work as well as the built in MATLAB function for our first contract. The spline seem to be a better fit for the regression, but can sometimes be unstable at the area where we have greatest interest. The polynomial regression will usually give a worse regression but will only be unstable around the extreme points.

The algorithm does sometimes not give us the expected results, as with the local and global volume used, but it is usually close. This might be due to some error in the algorithm. The early exercise boundaries seem also not to be accurate since we would expect them to be smooth but they are not always smooth. It can even so be an indicator of the optimal early exercise strategy.

We tried to used the trick of using spot price trajectories with a later starting point than t_0 to get the trajectories more spread out. This gave a better regression and increased the accuracy of the algorithm.

To maximize the profit of the swing option, one could use the early exercise boundaries to find the optimal stopping strategy. The strategy is easily applicable, however it can not be used exactly when we have a contract with physical delivery. It might be profitable to use the swing option to buy extra electricity, but it does not always mean that it is a good idea to do so if the buyer do not have any use for the extra electricity. Since electricity can not be stored the buyer would have to resell it to someone else. This would induce a profit but the profit might be less than the costs for buying and selling the electricity. Therefore it might not always be a good idea to use the swing option even though it theoretically it optimal to do so.

5.2 Outlook

A further expansion of the contracts used in this thesis could be to use both number of exercise rights with local and global volume at the same time. It could also be to divide the global volume restrictions to a certain amount for upswings and for downswings. It would be a much more complicated task than here where we only look at the deviation from the baseload volume.

There are great variation of swing option contracts and since they are often bought over the counter they can be adjusted to fit the specific buyer and seller. Therefore there are many different variations of swing options to price and there are also numerous ways to do so.

References

- Aleksandrov, N. and Hambly, B. M. (2010). *A dual approach to multiple exercise option problems under constraints*. *Mathematical Methods of Operations Research* 71(3), p. 503-533.
- Carriere, J. (1996). *Valuation of Early-Exercise Price of Options Using Simulations and Nonparametric Regression*. *Insurance: Mathematics and Economics*, 19, p. 19-30.
- Cartea, A. and Figueroa M. (2005). *Pricing in Electricity Markets: a Mean Reverting Jump Diffusion Model with Seasonality*. *Applied Mathematical Finance*, 12(4), p.313-335.
- Cartea, A. and Williams, T. (2008). *UK Gas Markets: the Market Price of Risk and Applications to Multiple Interruptible Supply Contracts*. *Energy Economics*, 30(3).
- Deng, S. (2000). *Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes*. Working paper, PWP-073, University of California Energy Institute.
- Dörr, U. (2003). *Valuation of Swing Options and Examination of Exercise Strategies by Monte Carlo Techniques*. Master thesis, University of Oxford.
- Jaillet P., Ronn, E. I. and Tompaidis S. (2004). *Valuation of Commodity-Based Swing Options*. *Management Science* 50, p. 909-921.
- Keppo, J. (2004). *Pricing of Electricity Swing Options*. *The Journal of Derivatives*, Spring, Vol. 11, No. 3: pp. 26-43.
- Hambly, B., Howison, S. and Kluge, T. (2009). *Modelling spikes and pricing swing options in electricity markets*. *Quantitative Finance*, 9:8, p. 937-949, DOI: 10.1080/14697680802596856.
- Longstaff, F. and Schwartz, E. (2002). *Valuing American Options by Simulation: A Simple Least-Squares Approach*. *Review of Financial Studies* 14, p. 113- 147.

Lucia , J. J. and Schwartz, E. S. (2002). *Electricity prices and power derivatives: Evidence from the Nordic Power Exchange*. Review Deriv. Research, 5, p. 5-50.

Meinshausen, N. and Hambly, B. (2004). *Monte Carlo Methods for the Valuation of Multiple Exercise Options*. Mathematical Finance 14, p. 557–583.

Appendix A

MATLAB function

`SwingPrice = hswingbyls(Paths, Times, RateSpec, Settle, Maturity, Strike, ExerciseDates, NumSwings, DCQ, minDCQ, maxDCQ, useSpline, P, LSplot)`

where the inputs are:

- Paths are the simulated trajectories.
- Times are the time points for the exercise opportunities.
- RateSpec is used to represent the interest-rate term structure.
- Settle is the settlement date, i.e. the date before the first exercise opportunity.
- Maturity is the date of maturity.
- Strike is the strike price K .
- ExerciseDates is a vector of all possible exercise dates between Settle and Maturity. If left empty, it will be possible to exercise at all time points.
- NumSwings are the number of exercise rights.
- DCQ is the daily contract quantity, which for our comparison is simplest set to 0. Note that it does not have to be set to 0, if it is set to $DCQ > 0$, we would just have to adjust the local volume restriction to that. The base load value of the contract is not included in SwingPrice.
- minDCQ is the level that the local quantity can not go below.
- maxDCQ is the level that the local quantity can not go above.
- useSpline is set to true if a spline is preferred over polynomials.
- P is the smoothing parameter.
- LSplot is set to true if we want to create a plot of the regression between the underlying price and the continuation value at the exercise date before maturity.