NORWEGIAN UNIVERSITY OF SCIENCE AND TECHNOLOGY

PROJECT THESIS

Real option valuation of expansion and abandonment options in offshore petroleum production

Authors:
Øystein Dahl HEM
Alexander SVENDSEN

Academic Supervisor:
Stein-Erik FLETEN

Abstract

We value two real options related to offshore petroleum production by considering expansion of an offshore oil field by tying in a satellite field, and the option of early decommissioning when the operator is able to sell the production unit. It is assumed that the decision to invest in the main field is already taken. Even if the satellite field is not profitable to develop at current oil prices, the option to tie in such satellites can have a significant value if the oil price increases. This implies that investing in excess production capacity as the production unit is built in order to obtain such real options can be profitable. We also investigate whether the option to sell the FPSO and decommission the field before its planned abandonment date is valuable. This could be a plausible action if the oil price fall below marginal cost, or if production volumes are lower than expected. Our results imply that this option does not have much value for reasonable cost assumptions. Two sources of uncertainty are considered, oil price risk and production uncertainty. The oil price is shown to be the most important factor affecting the value. By analyzing long-term price data, we find that modeling the oil price as a geometric Brownian motion is a better fit than using a mean reverting model. A mean-reverting production level is used to model production uncertainty to take into account the stochastic nature of well production levels. The option valuation is based on the Least Squares Monte Carlo algorithm, which is efficient when modeling multiple stochastic processes.

December 14, 2009

1 Introduction

We study the flexibility related to investment timing in offshore oil exploration and production. Offshore oil production can require large investments in infrastructure, offshore and onshore facilities and well drilling costs. These costs are to a large degree sunk once the investment has been made. Since 2000 oil prices have been increasingly volatile, creating uncertainty about whether marginal projects can deliver a sufficient return on the investment. This should make the problem of optimal investment timing interesting for both practitioners assessing investment opportunities and officials and researchers forecasting the future level of investment in petroleum production. The main decision in a petroleum production project is when and if the field should be built and the main investment to be made. Depending on the field and the technology used to produce it, the operator might also have available choices after the field has started to produce. Often there are smaller reservoirs surrounding the main field. Generally too small to warrant an independent production unit, these reservoirs can be produced by the production unit at the main field. The operator also has to decide when to abandon the field and decommission the production unit, by taking into account future production as well as equipment life time and operational cost. Once the field is abandoned restarting production is in most cases not realistic.

Real option valuation has been applied to petroleum projects for a long time as they have many attributes that make them suitable for this valuation. The projects often involve a large initial investment, the output is a risky and easily traded commodity, and management have many choices available, both related to timing, production technology and size. Siegel et al. (1988) assess investing in offshore petroleum leases. Cortazar and Schwartz (1998) use a Monte Carlo model to find the optimal timing of investing in a field with a set production rate that declines exponentially and with varying but known operating costs. With this predetermined production rate the value of the field becomes a function of the oil price, which is modeled as a two-factor model where the spot price follows a geometric Brownian motion and the convenience yield follows a mean reverting process. McCardle and Smith (1998) consider the timing of investment, the option to abandon and to tie in surrounding fields. Both prices and production rates are modeled as stochastic processes, where the price follows a geometric Brownian motion. Ekern (1988) use a ROV model to value the development of satellite fields and adding incremental capacity using a binomial model. Lund (1999) consider an offshore field development by using a case from the North Sea field Heidrun. The model used is a dynamic programming model, and take into account the uncertainty regarding both reservoir size and well rates, in addition to the oil price. The paper models the price as a geometric Brownian motion, and use a binomial valuation model to find the optimal size of the production rig and investment timing. Dias et al. (2003) use Monte Carlo simulations together with non-linear optimization to find an optimal development strategy for oil fields when considering three mutually excluding alternatives. Chorn and Shokhor (2006) combine dynamic programming and real options valuation to value investment opportunities related to petroleum exploration. Dias (2004) provides a more thorough review of real option valuation related to petroleum E&P.

We focus on real options related to the FPSO technology. An FPSO is a floating unit that can process and store oil, much like a traditional production platform. Our focus on FPSOs are due to several reasons. Firstly, they can be used over a large range of different depths, since the cost of an FPSOs does not vary dramatically with water depth, as opposed to conventional fixed platforms (Ronalds 2005). As new fields often will be further from shore, the ability to produce at large depths will be important. Another advantage is that many wells can be used with this solution. This makes it suitable to connect the FPSO with new fields as a way to keep the production rate at a high level longer. That way it is also possible to extract oil from fields that otherwise would have been unprofitable. It is also possible to move the FPSO to an other field if this is desired, which can be an option at smaller fields. Another advantage that makes the FPSO versatile is the possibility to expand the production capacity and process capacity due to the large area available on the deck of an FPSO (Dueas Diez 2009). Kemp and Stephen (2001) found some limitations to the FPSO technology. One of them is a capacity limit which limits the field size that can be developed with a FPSO. Even so, they conclude that FPSOs are the preferred technology for medium sized fields and a good choice for deep-water projects.

In this work we take into account both price risk and technical risk in a ROV model based on the Least Squares Monte Carlo algorithm presented in (Longstaff and Schwartz 2001). We do not consider the problem of initial investment, as it has been considered both in petroleum production and other industries before. Also, Lund (1999) found that the value of deferring an investment is generally low in petroleum production since profits are high and new information rarely change the investment decision. Instead we focus on decisions being made as the field is in production, where we focus on expanding production by tying in surrounding satellite fields as well as the option

Table 1: Returns and initial market weights for selected oil companies

Company	Return	Market Weight
Exxon Mobile	8 %	26 %
Royal Dutch Shell	4 %	25 %
Conoco Phillips	10 %	13 %
Chevron	10%	10 %
Repsol	5 %	10 %
Total	10 %	6 %
BP	4 %	5 %
Eni	13 %	4 %

to abandon the field early and selling the production equipment.

In Section 2 we present the data uses in this work, and discuss its properties. In Section 3 we investigate which price model is most suited to capture the long term behavior of the oil price, and select a model appropriate for real option valuation. In Section 4 we study the option to expand the production with a tie-in field as well as abandoning the field early. We then apply these models in a case study of two such real options. Finally, in Section 5 we conclude and offer suggestions for further work.

2 Data

To study the long term behavior of the oil price and to find a suitable time series model we have used the real price of crude oil denominated in 2008 USD from EcoWin-Reuter (2009). The series can be seen in Figure 1. The series consist of the US average price in the years from 1861 to 1944, then Arabian Light posted at Ras Tanura from 1945 to 1985, and Brent spot since 1985 to today. It has 148 annual observations going back to 1860. We could have used more high-frequency data for the latter years, but these are not available prior to 1946 for monthly data and 1977 for daily data. In order to avoid mixing the different series, only the annual observations have been considered. The properties of the annual time series will be studied further in Section 3

To find an estimate for an appropriate rate of return for an E&P project we use total returns for a selection of E&P oil company stocks. Total returns are defined as the annual stock price growth together with dividends. To find the average return we use the relative market weights for the selected companies at the start of the estimation period, and multiply this with the company total return. Market weights are calculated as each company's market cap divided by the market cap of all companies in the selection at the start of the estimation period. We have used information from EcoWin-Reuter (2009), for a period from 11.11.1999 to 11.11.2009. This can be seen in Table 1. For the last 10 years we find that the market weighted average return is 7.4%. We compare this with what other researches have used previously. Castro et al. (2002) use a risk-adjusted rate of 8% while Lund (1999) and Cortazar and Schwartz (1998) use 7%. Brashear et al. (2000) states that for the E&P industry the average return on net asset is 7%. Thus, a risk-adjusted rate of 7.4% seem to be reasonable estimate for risk adjusted return from a petroleum production project.

To obtain risk-neutral growth and the oil lease rate we use forward prices, $F_{0,T}$, from Wall Street Journal (2009) for Light Crude Oil, as seen in Figure 2. The series have contracts for each month till December 2014, and semiannual contracts expiring as late as December 2017. We find an estimated risk neutral long term growth, g, of 2.70%. This is higher than the estimated real growth rate from Section 3.4 that was found by using past price data. This could indicate that the expected growth going forward is higher than the historic average. The longest duration for the forward contracts used is 8 years, but we assume that the growth indicated for these is a good estimate for the growth in our 20-year period. We can use the expected growth together with an estimate for the risk free rate to find the oil lease rate, δ, by using eq. (1) from (McDonald 2005). This puts the oil lease rate at 1.6%.

$$\delta_l = r_f - \frac{1}{T} * ln\frac{F_{0,T}}{S_0} = r_f - g \tag{1}$$

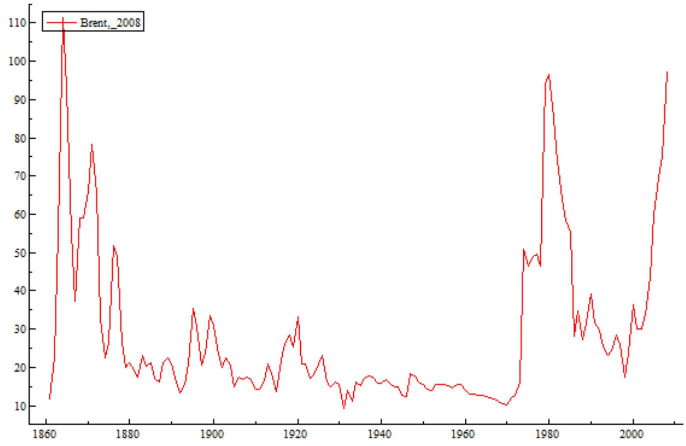

Figure 1: Real price adjusted Brent spot price, USD 2008

Another way of estimating the lease rate is from eq. (2) found in (Dixit and Pindyck 1994). Using the average return for an E&P company as an estimate for the project return we find a lease rate of 3.1%, which is higher than what was found when using forward prices. This indicates that the returns from a production project and the whole of an E&P company might not be perfectly correlated, and that using company wide discount rates to value such projects might be inappropriate. This contradicts our previous conclusion. We use the lease rate obtained from forward prices in eq. (1), since it does not require any assumptions about the appropriate expected return, μ.

$$\delta_l = \mu - r_f \tag{2}$$

For a market-based estimate for the volatility, we have used implied volatility from options quoted at ICE (2009) for Brent oil options. The series have options expiring at 17 different dates with several options set to expire at each date. The longest time to expiry is 3 years. We have used an average value over all strike prices available to find a mean implied volatility for each date. The implied volatility is falling with longer expiration time, implying that the 3-year forecast might not be valid for the long-term real options. Even so, we use the implied volatility for the 3-year option as the oil price long term volatility, with an implied volatility of 29.5%. This is quite a lot higher than the historic average for the last 148 years, but close to the volatility in the last 40 years found in Section 3.4 of 28.8%. It is also higher than the 20% that Pindyck (1998) found when estimating volatility from historical data. Costa Lima and Suslick (2006) refer to Pindyck (1998) and also argue that the volatility has been stable around 20%. We use the implied volatility as an estimate for the long term volatility. One reason for the difference between the market view and the conclusions of Costa Lima and Suslick (2006) and Pindyck (1998) could be the increase in oil price volatility in the last years.

To estimate the USD-denominated risk free rate, we have used 20-year US Treasury bonds from EcoWin-Reuter (2009) as an estimator for the risk free rate. As seen in Figure 4, the yield for a 20-year US Treasury bond is 4.3%.

3 Time series analysis of the oil price

In this section we will review two common stochastic price processes, and find the one best suited for modeling the price of oil for the real option study.

3.1 Introduction

One of the most significant factors in valuing a potential oil field is the price of oil. Like the price of other tradeable items the oil price is governed by supply and demand. The theoretical ideal model would take into account all the factors that affect supply and demand and produce a forecast of the oil price based on this information (Pindyck 1998). Several such models have been developed, among others the Hubbert model of supply (Hubbert 1956) and

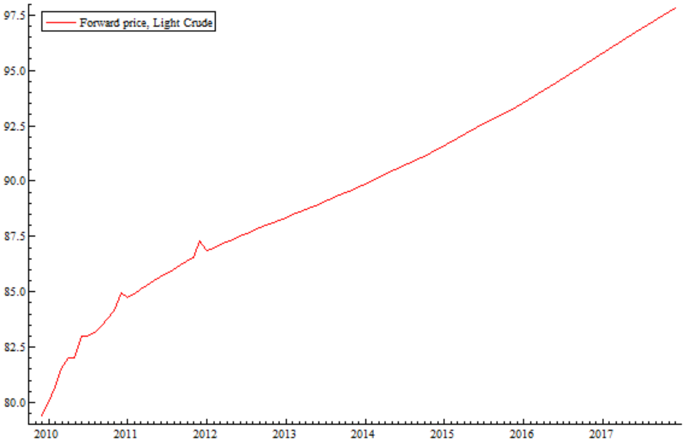

Figure 2: Forward price light crude oil

Figure 3: Implied volatility

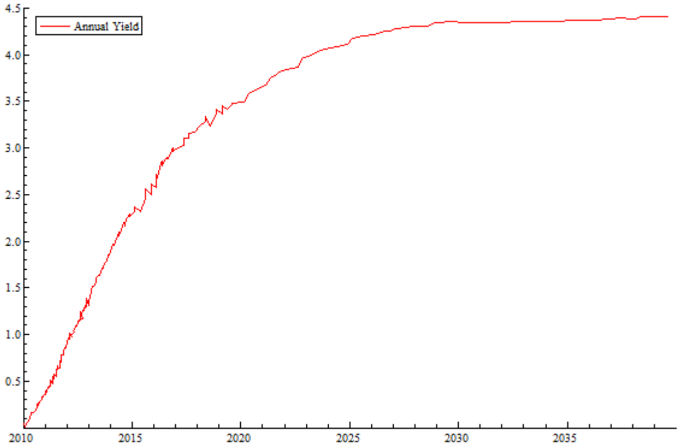

Figure 4: Yield curve, US Treasury Bonds

the LOPEC model (Rehrl and Friedrich 2006). These model the price development by looking at the underlying factors that drive supply and to some extent demand. There are two major obstacles for implementing such a model for generating long term forecasts. First, identifying all of the factors affecting the oil price is in itself a difficult task. Second, producing good forecasts for all of these factors might be just as difficult as producing a forecast for the oil price.

A time series model thus seems like an attractive alternative model formulation. In a time series model one attempts to model and predict the variable using only past information about the development of the variable. Following a decade of stable prices, the last few years have seen high volatility compared to previous years, and has fluctuated from USD 51 at the start of 2007 to the peak of USD 130 in late may 2008 before going down to USD 34 at the start of 2009 (EcoWin-Reuter 2009). This makes modeling difficult, since it is difficult to know if the mean level and the volatility of the oil price will revert to its long term average or if a structural break has occurred. This can make predictions made by forward-looking market prices, like those discussed in Section 2, more attractive than a time series analysis based on backward-looking estimates.

In this section we consider two common stochastic processes used to model commodity prices, geometric Brownian motion and mean reversion. These are not usually regarded as time series models, but can be modeled as such.

3.2 Different stochastic processes for the oil price

3.2.1 Geometrical Brownian Motion

One of the most common stochastic processes for commodities and other financial assets is the Geometrical Brownian Motion (GBM) (McDonald 2005) :

$$\frac{dP}{P} = \alpha dT + \sigma dZ \tag{3}$$

Here the price is denoted by P, α is the expected growth rate and σ the instantaneous standard deviation.
This can also be expressed as:

$$\frac{dP}{P} = (\mu - \delta)dT + \sigma dZ \tag{4}$$

Here $(\mu - \delta)$ has replaced α as the drift term. μ represents the expected return, and δ the convenience yield.

The GBM is an It process where the drift and volatility change with the stock price. Another property of a price following (3) is that it becomes lognormally distributed if dZ is normally distributed:

$$dln(P) = (\mu - \frac{\sigma^2}{2})dt + \sigma dZ \tag{5}$$

Since the log of P is normal, P follows a lognormal distribution. To find a closed form expression for P, one can integrate (3) to obtain:

$$P(t) = P_0 * e^{(\mu - \frac{\sigma^2}{2})t + \sigma Z} \tag{6}$$

The expected value of P is:

$$E(P(t)) = P_0 * e^{\alpha t} \tag{7}$$

and the standard deviation:

$$\sigma_t = P_0 * e^{\alpha t} * \sqrt{e^{\sigma^2 t} - 1} \tag{8}$$

To estimate the parameters in the GBM based on past oil prices the estimators from (Luenberger 1998, p. 303-310) can be used. These are given in Section A.3. Alternatively, one can use risk-neutral estimators as described in Section 2.

A good quality of a lognormal distribution with incremental Δt is that the variable can never become negative, which is a true property for prices.

3.2.2 Mean reverting process

Several papers have argued that commodity prices can be modeled as mean reverting price process, among others (Pindyck 1998). The fundamental argument supporting mean-reversion in commodity prices is that if the price is lower than the long term equilibrium, high-cost producers will exit the market which reduces supply. Similarly, if the price is high more producers will enter which will put a downward pressure on the price. This process will tend to revert the price to a mean level. The model is also known as an Ornstein-Uhlenbeck process (McDonald 2005):

$$dP = \lambda(\alpha - P)dT + \sigma dZ \tag{9}$$

λ denotes the mean-reversion parameter. If the price P should rise above the expected value α the drift term, $\lambda(\alpha$ - P), will become negative and vice versa if the price should fall below the mean level. The price will revert toward the mean level, α, in the long run . If λ is large the reversion will be quicker.

The expected value of any future time is:

$$E(P_t) = \alpha + (P_0 - \alpha)e^{-\lambda * t} \tag{10}$$

P_0 is the current price. The standard deviation of P_t follows:

$$\sigma_P^2 = \frac{\sigma^2}{2\lambda}(1 - e^{-2\lambda t}) \tag{11}$$

To estimate values for the model parameters we use the procedure described in Section A.4. A negative issue with the mean reverting model is that the price can become negative. The model can also be extended by including a drift term, e.g to correct for increasing marginal cost of extraction. This might be a more fitting model for petroleum extraction, since the least expensive resources are developed first and as these are depleted more expensive fields take their place. This will cause an increase in marginal cost over time.

3.3 Model analysis

By testing whether different characteristics are present when creating a model for a price process, it is possible to see if there are some models that can be excluded or if others can be favored. Pindyck (1998) suggest that the oil price could follow a mean-reverting process with a very low mean reversion rate and that few errors will be introduced when modeling the price as a GBM for medium time horizons. We examine if the requirements and assumptions for these two models are met before deciding on one of the models. In order to test how sensitive the parameters are to the length of the time series three periods have been selected. The first consider the full time series of 148 years, the second the last 100 years, and the third the last 40 years. The first observation is lost to calculate returns.

3.3.1 Durbin-Watson test results

Table 2: DW GBM

Periods	Test statistic	d_L	d_U	Comments
147	1.67	1.52	1.56	Do not Reject H_0
100	2.12	1.52	1.56	Do not Reject H_0
40	1.95	1.25	1.34	Do not Reject H_0

In Table 2 and 3 d_L represents the lower critical value and d_U the upper critical value at 1% significance level.

Table 3: DW Mean reverting

Periods	Test statistic	d_L	d_U	Comments
147	1.47	1.52	1.56	Reject H_0
100	1.80	1.52	1.56	Do not Reject H_0
40	1.71	1.25	1.34	Do not Reject H_0

As can be seen in Table 2 we cannot prove autocorrelation in the residuals for any the of periods. This is in accordance with the assumptions for a GBM.

For the mean reverting model Table 3 shows that we do not find evidence of autocorrelation in the residuals for the series with 100 and 40 observations. There is significant autocorrelation in the case of 147 periods, however. This might imply that the model do not sufficiently capture the trends in the data set. As discussed in Section A.2.3, this can cause the standard error estimates to be wrong, and probably lower than the true standard errors.

3.3.2 Unit root test results

Using (36) and (37) to find the Dickey-Fuller test statistic, we first test whether the returns are stationary and second we test if the price itself is stationary:

Table 4: DF Returns

Periods	Test statistic	Critical value	Comments
147	-10.54	-4.01	Reject H_0
100	-10.59	-4.04	Reject H_0
40	-5.93	-4.24	Reject H_0

Table 5: DF Price

Periods	Test statistic	Critical value	Comments
147	-2.40	-4.01	Do not Reject H_0
100	-1.45	-4.04	Do not Reject H_0
40	-1.13	-4.24	Do not Reject H_0

Where the critical value is given on a 1% significance level.

As showed in Table 4 H_0 is rejected for all return series, which means that there is no unit root in the return series and that the returns are stationary. This is according to the assumptions for the GBM-model, since it assumes that the returns are stationary.

We see in Table 5 that H_0 is not rejected for the price series. We cannot reject that the price is non-stationary. This is in accordance for for the GBM model where the price is assumed to be non-stationary, but not for the mean reverting process where the price is assumed to be reverting back to a mean value and to be stationary. The price data fail to reject H_0 even at a 10% level.

3.3.3 Jarque-Bera test results

The critical value is given at a 1% significance level.

From Table 6 and Table 7 we see that for GBM residuals and the mean reverting residuals H_0 is rejected on a 1% significance level. This indicates that one must reject that the residuals are normally distributed. This is not surprising as oil returns rarely follow a normal distribution, but this is often disregarded when modeling the series.

Table 6: Jarque-Bera GBM residuals

Periods	Test statistic	Critical value	Comments
147	59.21	9.21	Reject H_0
100	161.87	9.21	Reject H_0
40	36.54	9.21	Reject H_0

Table 7: Jarque-Bera Mean reverting residuals

Periods	Test statistic	Critical value	Comments
147	669.78	9.21	Reject H_0
100	802.39	9.21	Reject H_0
40	48.53	9.21	Reject H_0

When the residuals not are normally distributed, it could be better to use a different type of distribution such as the student-t distribution or a skewed student-t distribution.

3.3.4 Jarque-Bera test results when removing outliers

As can be seen in Figure 5 a couple of extreme outliers may be the cause of the normality rejection in this case. Note that the figure without removing outliers has higher values on the axis. By using dummy variable to exclude a few of the extreme values it may make the residuals normally distributed. By looking at the extreme outliers we observe what year they appeared in, and then investigate if there is any justification to remove them. Removing outliers should only be done with sufficient argumentation for why the observation is not representative for the series, since it removes information. It can also make the results prone to data mining.

In this case using dummy variable for the extreme residuals in the years 1974, 1979 and in 1986 can to some extent be justified. The two first observations are connected to two different oil crises which made the oil price soar, the last is when the price plummeted because of low demands as a result of the high prices from the two oil crises. There also seems to be a shift in the oil price behavior from year 1880, with less extreme variation than in the years before. We therefore start the series in year 1880 and correct for the three extreme observations using 129 periods. The results are showed in Figure 6. We then do the Jarque-Bera test again. The results are found in Table 8 and 9.

Table 8: Jarque-Bera Mean reverting residuals when removing outliers

Periods	Test statistic	Critical value	Comments
129	103.15	9.21	Reject H_0
100	90.01	9.21	Reject H_0
40	9.89	9.21	Reject H_0

For the mean reverting series, we observe that the residuals are closer to a normal distribution than before, but we still reject the normality hypothesis. Only with 40 periods is it close to not being rejected, but this can be related to having few observations to base a conclusion on. The unit root test should according to Dixit and Pindyck (1994) not be performed on data sets with less than 30 observations. Based on this we conclude that the residuals from the mean-reverting model do not follow a normal distribution.

By removing the same outliers and using the same assumptions on the GBM case we obtain the residuals seen in Figure 7.

In this case we see that the results are different. We cannot reject that the residuals from the GBM model is normally distributed. From the Jarque-Bera results, GBM is looking more promising than a mean reverting model. For the mean reverting case, it could be better to use another distribution model, such as the student-t distribution.

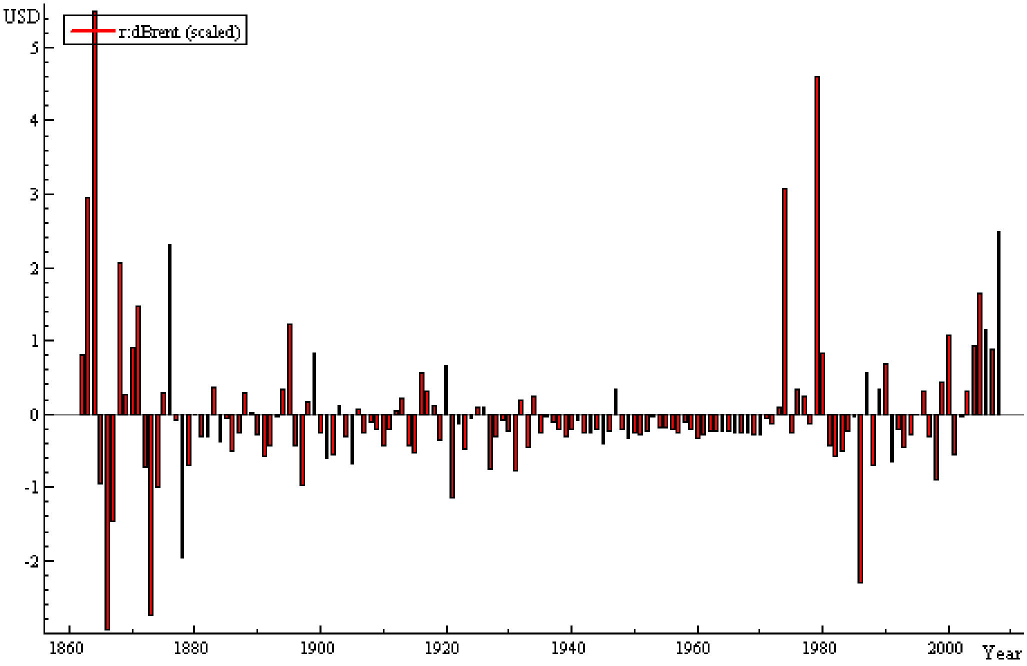

Figure 5: Residuals to mean reverting model

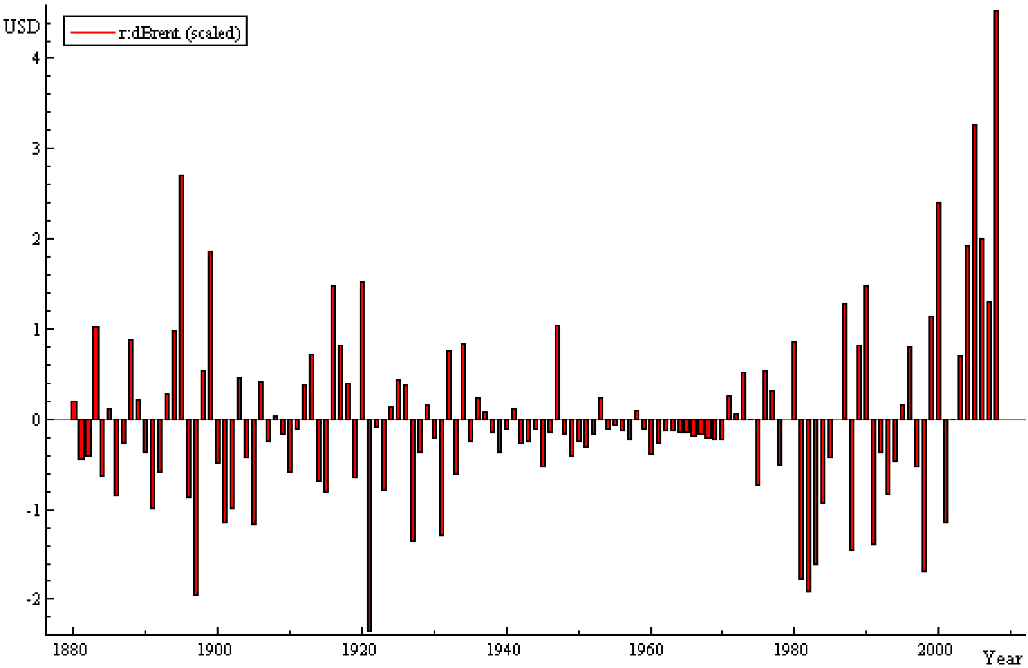

Figure 6: Residuals to mean reverting model when removing outliers

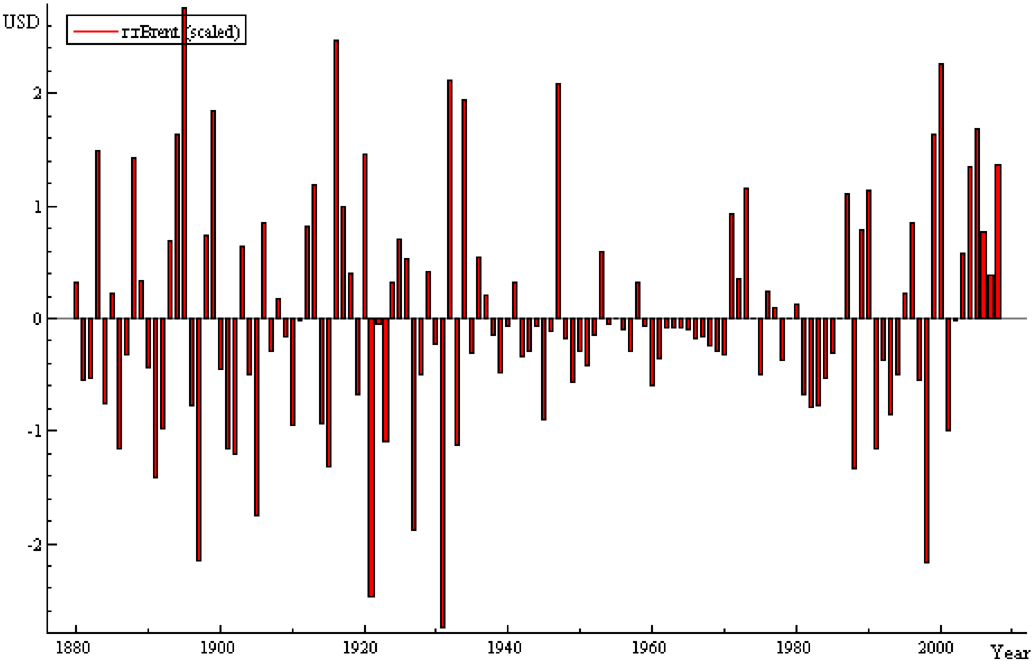

Figure 7: Residuals to GBM model when removing outliers

Table 9: Jarque-Bera GBM residuals when removing outliers

Periods	Test statistic	Critical value	Comments
129	2.60	9.21	Do not reject H_0
100	3.87	9.21	Do not reject H_0
40	0.11	9.21	Do not reject H_0

3.4 Modeling

3.4.1 Geometrical Brownian Motion

To estimate the parameters in the GBM model equation (46) from Section A.3 is used. To evaluate how sensitive the estimators are to the length of the period, four estimators are produced with periods of 40, 100, 129 and 147 years. The estimators from the 129 periods have had the outliers removed. The results are seen in Table 10:

Table 10: GBM Estimators

Period	$\overline{r}$	s	$\widehat{\alpha}$	$\widehat{\sigma}$
147	0.0507	0.2695	0.0870	0.2695
100	0.0175	0.2296	0.0428	0.2296
40	0.0517	0.2877	0.0931	0.2877
129	0.0031	0.1877	0.0207	0.1877

Here $\overline{r}$ is the expected or mean change of the oil price from one period until the next, s is the standard deviation of the oil price in dollars.

From these results we see that the estimator values are sensitive to the period used. It is also interesting to see how the results are affected by removing the most volatile period early in the series as well as the three outliers. This has the effect of lowering the expected return from around 5 % to almost zero. Based on the results from the 129 period case with the three most extreme outliers removed a Monte Carlo simulation is conducted to obtain random price realizations. With 100 000 random realizations the 20-year forecast has an expected value of USD 90.3 and a standard deviation of USD 90.3 measured in USD 2008, based on a current price of USD 60. Using eq.(7) and eq.(8) we find an expected value of USD 90.77 and a standard deviation of USD 91.81. We see that the analytical results are close to the numerical solutions.

3.4.2 Mean Reverting Process

To estimate the parameters in the mean reverting model equation (50) from Section A.4 are used. To check if the values are sensitive to the length of the period four estimators are produced, as in the previous section. The results are shown in Table 11:

Table 11: Mean-Reverting Estimators

Period	$\bar{r}$	$\overline{P}$	s	$\widehat{\alpha}$	$\widehat{\sigma}$	λ
147	0.0553	28.667	20.874	33.614	7.353	0.0493
100	0.0482	26.736	20.094	56.314	5.475	0.0116
40	0.1079	24.443	41.686	63.409	8.342	0.0424
129	0.0413	25.568	18.077	30.280	5.082	0.0135

$\overline{P}$ is the mean value of the 2008-equivalent oil price.

It is worth noting that some of the estimators changes significantly if we use 40 periods instead of 147 periods or 100 periods, $\bar{r}$ and the standard deviation more than double in value. This indicate that the mean level could have increased. Based on these results a Monte Carlo simulation is conducted to obtain random price realizations. With 100 000 random realizations and using 129 periods the 20-year forecast has a mean value of USD 52.86 and a standard deviation of USD 20.05, based on a current price of USD 60. Using eq.(10) and eq.(11) we find an expected value of USD 52.97 and a standard deviation of USD 19.98. The analytical solution is close to the numerical solution. Both expected value and volatility is significantly lower than for the GBM model. The price development can be seen in Figure 8a

Table 12: Mean reversion forecast and goodness of fit

Period	Mean price	std price	R^2
147	43.26	22.03	0.039801
100	59.23	22.13	0.003623
40	61.91	26.32	0.029591
129	52.86	20.05	0.557098

Table 12 shows the 20-year forecast of the mean and standard deviation for the price. We see that the expected price increases if we use fewer periods in the model, due to a higher mean price level in recent years. This indicate that a mean-reverting model with drift might be a good alternative. The same conclusion is also valid for the standard deviation but it increases significantly more if we use 40 periods instead of 147 or 100 periods. The R^2 statistic is very close to zero for 147, 100 and 40 periods. But R^2 increases to 0.557 for 129 periods when the the outliers are removed. This is a good indication that it is necessary to remove the outliers in order to get a good model that fits the data.

Comparing the forecasts from the mean reverting model with those from the geometric Brownian motion we see that the volatility is much lower in the mean reverting case. One can also note in Figure 8bthat the mean reverting model does at some times produce a price below zero, something that is avoided in the GBM formulation. By inspecting the forecast in Figure 8b or by noting the low value of λ, one can see that the speed of reversion back to a mean value is very low in this case. This indicates, together with non-stationarity in the series, that the mean reversion model might be wrong for this series. Several other papers conclude differently, among others Pilipovic

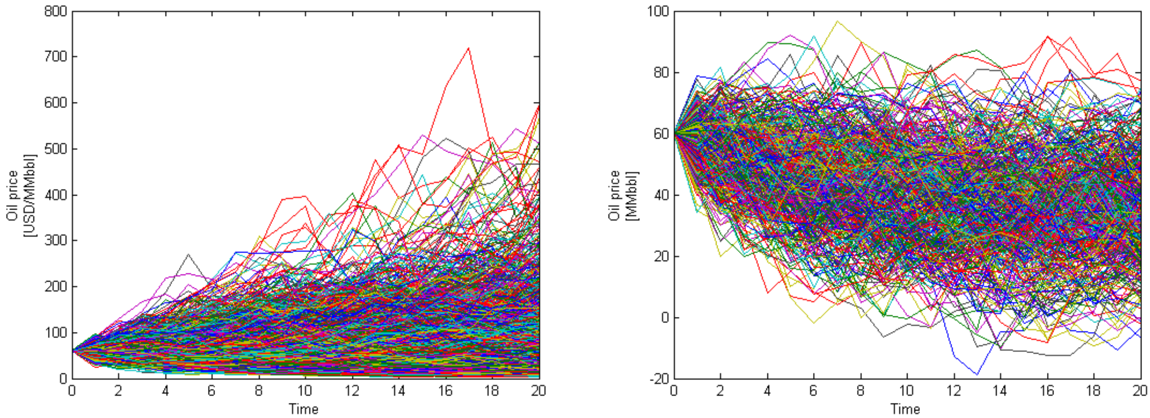

(a) Geometric Brownian motion

(b) Mean reversion

Figure 8: Simulated Monte Carlo forecasts

(1998). We believe that the main reason for our different conclusion is the large increase both in the oil price and its volatility since 2000, data points that earlier works of obvious reasons have not considered.

3.4.3 Model Conclusions

Based on this analysis we choose to model the oil price as a geometric Brownian motion. We have identified a unit-root in the price process even when looking at the full 147 year period, a finding that contradicts the mean reverting model. This is the main reason for not proceeding with the mean reverting model. We also find that the residuals of the geometric Brownian motion can be normally distributed after removing the most extreme outliers, where as the residuals from the mean reverting model are not. Other models, like the two-factor model in (Schwartz 1997), have not been considered here but could be a good alternative to a pure geometrical Brownian motion model in order to model more realistical price behavior. We have chosen not to consider this model in order to make the ROV model more easily understandable.

The forecast produced by the GBM is not very accurate, as the volatility grows with time. This is not necessarily a negative property, as it has been very difficult to make good predictions about the future price of oil. It is also sensitive to the choices made with regards to excluding outliers in regressing the return. We argue that not taking the first 20 years of the series into account is a valid assumption, since the market conditions in the early years of the market is not representative for current prices. The exclusion of the three most extreme returns reduce the expected return with 1% and lowers the volatility with 4 %, indicating that these three returns have a strong influence on the regression.

The forecasted value is also very dependent on which price point the forecast is produced from, since the GBM only specifies an expected long term growth. With an expected growth, α, of 2 % the expected price 20 years from now is 50 % higher than the starting price. Thus, the starting point becomes very important. This is somewhat problematic, since the volatility have been very high the last few years, making the forecast very dependent on the date the forecast is made. One way to make reduce the sensitivity of the starting date is to use an average price value, for instance the average value of the price through the year.

The residuals of the returns do not follow a normal distribution. Some of this non-normality might stem from outliers, and we do not reject normality in the returns when removing the three most extreme outliers the residuals. Even so, it might be preferable to model the process with another distribution than the normal distribution. We disregard this in our further work, however, and estimate the price increase using a normal distribution.

4 Real Option valuation

4.1 Flexibilities in petroleum production

In this section we consider two cases where the operator has flexibility, and develop valuation models for this flexibility. We do not consider the case of initial investment timing. We have chosen to consider two real options relevant for FPSOs, since FPSOs are considered to be a flexible alternative with regards to expanding production and other issues important in frontier fields.

4.1.1 The value of including satellite fields

In some situations the operator knows of a smaller and nearby field that can be produced through the main production platform. These smaller fields will often have higher per-barrel costs due to economies of scale, and are thus more interesting to consider in a real option model than ordinary fields since they are not necessarily economical to develop. Typically, such fields will not be large enough to warrant an independent platform but it can be profitable to tie the fields to existing platforms. Tying in a small field will increase the produceable reserves connected to the platform, but require an investment. The deterministic NPV of tying in such a satellite field can be calculated by using the reservoir model presented in Section 4.2 and valuing the incremental production from the satellite, given the capacity constrains and the time of connection. Given that the increased costs by adding the satellite are fixed, the value of extra production will only vary with the price of oil, which is stochastic, and the time of connection. If the satellite field is connected before the production declines, it will not increase the production from the platform until the main field is off its plateau, since the plateau is given by the platform's maximum production rate. Further, if the satellite is connected near the end of the platform's life time, much of the extra fields reserves will be left in the ground unless one extends the lifetime of the platform, which might not be possible depending on the availability of infrastructure etc. Developing a satellite field can require a large initial investment, and it is assumed that any extra operational costs are included in the investment cost. Since these are modeled as deterministic cash flows the NPV of the future costs are simply added to the investment. Thus, the value of being able to include a satellite field takes the form of a call option to acquire the extra production by paying the investment cost.

The increase in production is the difference between the line and the dotted line in Figure 10. We can calculate the net present value of increased production when connecting the tie-in at time t by using Equation 12:

$$NPV_{S,t} = S * \sum_{i=t}^{T} Prod_i^{\Delta} * e^{-\delta * i} - I \tag{12}$$

S represents the price of oil, $Prod_i^{\Delta}$ the production from the satellite in period i, δ the convenience yield and I the present value of the investment and operational costs.

4.1.2 The timing dimension of including a tie-in field

The process of valuing a project with a fixed end date is different than for an ordinary stock. Even with uncertainty in the output, one is certain that the tie-in will be worthless at the time the main platform is decommissioned. In the case study we have used a production profile from (Robinson 2009) to calibrate the model from (Lund 1999) in order to get a representative production profile. The profile from (Robinson 2009) can be seen to the left in Figure 9, and the simulated one to the right. We see that the simulated profile is similar to the one from the case in (Robinson 2009), but that the decline in the case data does not follow a smooth exponential decline. Still, we consider the model to be sufficient as the main characteristics of the production profile are captured.

4.1.3 The value of early shut down

One of the advantages of an FPSO versus other production technologies is that it can be moved if the value of the remaining production is low, and the FPSO is not near the end of its life. This can be the case if the true field reserves are lower than estimated. To model this we have used the same price and reservoir model as in the expansion case, but now it is the whole project value that is relevant. Thus the value of ending the production prematurely can be calculated by using Equation 13:

$$NPV_{S,t} = K_t - (S * \sum_{i=t}^{T} Prod_i * e^{-\delta * i} - C_i e^{-r*i}) \tag{13}$$

This states that the value of decommissioning the field early is the income from selling the FPSO, K_t, less the future expected profit, stated as remaining production less the operational costs, C_i. It is assumed that it is possible to sell the FPSO either to another project or another company for a positive price. We have assumed that the FPSO depreciates linearly and that the income from a sale follows this value. We assume that the FPSO has a planned lifetime of two times the field lifetime. Then the strike will take the form:

$$K_t = K_0 * \frac{2T - t}{2T} \tag{14}$$

4.1.4 The effect of uncertain production

In Section 4.2 we model the production uncertainty as a mean reverting process. Unlike a Brownian motion, the expected value of a mean reverting process at time t is dependent on both its current value and its equilibrium value.

$$E(\gamma) = \alpha + (\gamma_0 - \alpha)e^{-\lambda t} \tag{15}$$

where γ represents the production index, α the mean index level, and λ the mean-reversion speed. From this we can derive an expression for the expected present value of the production, given the current value of the production index γ.

$$PV = \int_{t=0}^{T} S * (\alpha + (\gamma_0 - \alpha)e^{-\lambda t}) * Prod_t * e^{-\delta t} dt \tag{16}$$

$$PV = \int_{t=0}^{T} S * (\alpha * Prod_t * e^{-\lambda t} dt + \int_{t=0}^{T} S * \alpha * \frac{\gamma_0 - \alpha}{\alpha} * Prod_t * e^{-(\delta + \lambda)t} \tag{17}$$

$$PV = PV_{0,\delta} + \frac{\gamma_0 - \alpha}{\alpha} * PV_{0,\delta+\lambda} \tag{18}$$

We see that the effect on expected value given uncertainty in production in the last term of Equation (18). The current difference between γ and its equilibrium value α is dampened by λ, and this can be expressed in the present value calculation as a increased rate of return. Now we can consider $\lambda's$ effect on the present value. If λ is zero, the mean reversion effect is non-existent. Then the expected value of λ is simply its current value, which is the same as a Brownian motion. Also if λ approaches infinity the effect of a difference between γ and α will only effect the value in this period, as γ will revert to its equilibrium value in the next period.

4.2 Reservoir model

The field production profile is useful when valuing real options, since it provides information on volume and time of production. A realistic model of reservoir performance is challenging to create and to calculate, because of the need to model many parameters in a 3D-setting with many non-linear relations. In this work a simple zero-dimensional model from (Lund 1999) is used. This models the reservoir as a tank with a uniform fluid and with uniform properties in the whole reservoir. Thus, it does not account for differences in permeability in different areas or local differences in pressure caused by the well flow as the areas surrounding the producing wells empties. It is however a simple model which has great computational advantages compared to a more complex reservoir model, and it does reflect the form of reservoir production profiles of several types of petroleum fields (Lund 1999).

The reservoir pressure follows the following relation:

$$P_{w,t} = P_{w,0} - \frac{R_0 - R_t}{R_0} * (P_{w,0} - P_{min}) \tag{19}$$

The reservoir pressure provides the maximum well flow, which decays exponentially with time with continuous production if there are no other constraints on the well flow. The maximum well rate is based on the capacity of

Table 13: Reservoir Parameters

$P_{w,0}$	-	Initial reservoir pressure
$P_{w,t}$	-	Reservoir pressure at time t
P_{min}	-	Abandonment pressure
R_0	-	Initial reservoir volume
R_t	-	Reservoir volume at time t
$q_{r,t}$	-	Maximum reservoir depletion rate at time t
q_w	-	Maximum well rate
q_{max}	-	Maximum capacity, or plateau production
$q_{ramp-up,t}$	-	Maximum production during field development
N_t	-	Number of wells producing at time t

the wells installed.

$$q_{r,t} = N_t * q_w * \frac{P_{w,t} - P_{min}}{P_{w,0} - P_{min}} \tag{20}$$

Together, equations (19) and (20) becomes the simple equation

$$q_{r,t} = N_t * q_w * \frac{R_t}{R_0} \tag{21}$$

This is the maximum production from the field, given that there is no water injection or other types of pressure maintenance performed. It is rarely optimal to construct the production unit so that it can produce at the maximum rate $q_{r,t}$, because of high investment costs. When the field has a maximum processing capacity that is lower than the field maximum production, the production profile will have a flat region where the production is equal to the capacity maximum. This level is called the plateau production. The optimal plateau level is mainly a function of investment cost, production and required rate of return, since it is a trade off between investment cost and the ability to get the oil quickly out of the ground. There might also be technical reasons to limit the capacity. We have included a ramp-up period of three years, which is similar to the case found in (Robinson 2009) as seen at the left in Figure 9. During this ramp-up period we have assumed that the production grows linearly to capacity maximum over the three year period. The background for such a ramp-up period is among other topics well drilling. It will not be possible to drill all wells at the same time, and connecting the streams to the platform will also require some time.

The actual production thus becomes the minimum of $q_{r,t}$, q_{max} and $q_{ramp-up,t}$. The resulting production can be seen in Figure 9b. The model does not follow the profile from (Robinson 2009) exactly, but it does capture the main features.

4.2.1 Production Profile with a tie-in field

To model the increase in production by a tie-in satellite field, the new reserves, R_{new} are added to the initial reserves. This increase both the initial reserves, R_0 and the reserves at the connection time, R_t. The effect of this increase is dependent on when the new field is built. If the satellite is connected before the field goes into decline, then the plateau production will be maintained longer as seen in Figure 10a.

4.2.2 Uncertainty in production

Production volumes are often uncertain, as wells can produce more or less than planned. Lund (1999) models this by a changing well capacity. The well capacity is modeled as a simple stochastic function, where the well can either have a high or low well rate. The probability of one of the wells changing regime from a high rate to a low or opposite is 0.1 per period of 6 months. Each well capacity will be highly random, but with a large number of wells the process resemble a mean reverting stochastic process. The variance of the field production will be very dependent on the number of wells connected to the field. McCardle and Smith (1998) take a different approach by modeling the decline rate as a geometric Brownian motion. This might be appropriate when the field is in decline, but it does not take into account the effect of the production capacity limit and it does not clarify which

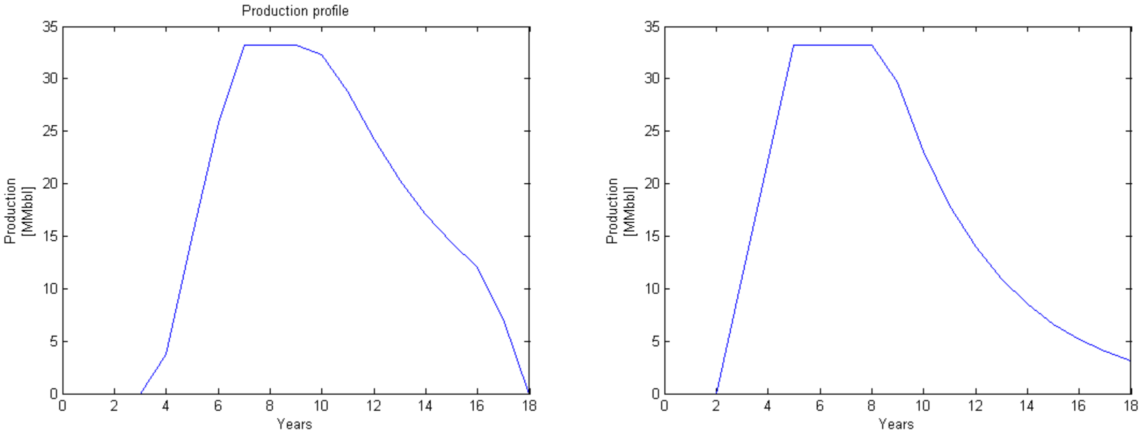

(a) Robinson (2009)

(b) Our model

Figure 9: Production profile from Robinson (2009) compared to our model

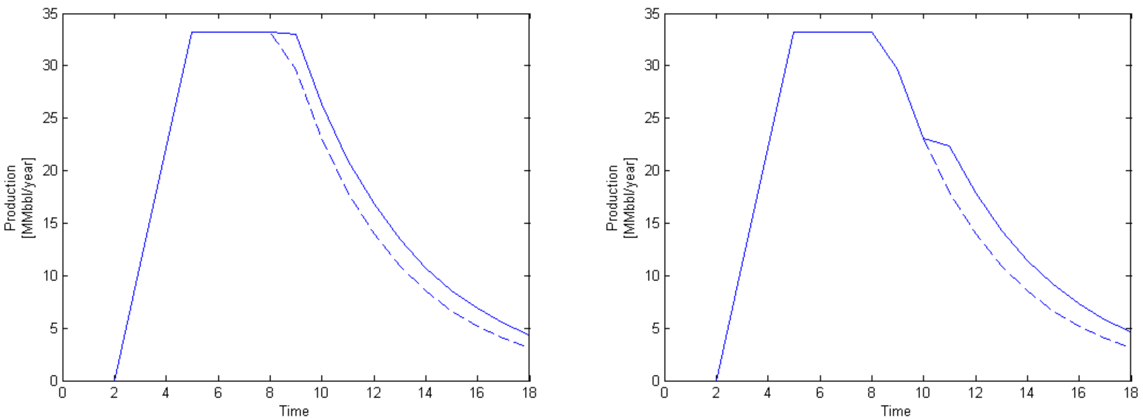

(a) t=5

(b) t=10

Figure 10: Production profiles with tie-in at t=5 and t=10

fundamental property that varies. We consider changing well rates as the main source of uncertainty, as in the the switching model in (Lund 1999). We do not model each well individually, however, instead we consider the whole field production by assuming a number of wells. This is implemented as a production factor for the whole field, γ_t, as a mean reverting process. We believe that this aggregate production factor is more versatile than the model in (Lund 1999), as operators can create historic production factors from current and previous fields and easily take into account other risk factors like technology development or unscheduled maintenance. The production factor follows:

$$\gamma_t = \gamma_{t-\Delta T} + \lambda(\alpha - \gamma_{t-\Delta T})dT + \sigma dZ \tag{22}$$

where γ_t is the well production index at time t, and λ, α and σ are mean reversion parameters from the regression. The parameter values can be seen in Table 14.

4.3 Valuation framework

There are mainly two ways of calculating the present value of future cash flows. One solution is using risk-adjusted rates of return and real expected growth rates. The other is risk-neutral pricing.

Table 14: Capacity mean reversion parameters

Parameter	Parameter Value
α	0.665
λ	0.218
σ	0.050

The rate of return used in the valuation of real options have a significant influence on both optimal exercise policy and option value. Especially with long term valuations, like many real options, a slight change in the rate of return can make a substantial difference due to the compounding effect. Using a correct discount rate is thus important to obtain correct results. We found in Section 2 that a discount rate between 7 and 8% is a reasonable estimate.

Another procedure of obtaining a valuation is to price the cash flows using other securities with similar risk profiles that are traded in the market. By replacing the real price growth with the risk-neutral price growth obtained from traded forward-contracts, one can use the risk-free rate to obtain the value of the project and connected options. This treats risk in a consistent manner compared to the market, avoiding biases that can occur otherwise (Laughton et al. 2008). This is commonly called risk-neutral valuation. Since all parameters are estimated from financial markets, which are assumed to be efficient, this leads to a correct valuation of the project. We describe the relevant parameters and how to obtain them in Section 2.

Using risk-adjusted rates has the advantage of being familiar to decision-makers in most firms today, and is perhaps the most intuitive of the two approaches. We do however choose to use risk-neutral pricing, since this ties the valuation of the risky cash flows directly to observed prices of this risk and should thus give the correct valuation. The risk-neutral method is also the most common approach when valuing options. One issue with using risk-neutral pricing is that the risk-neutral method can underestimate capital costs when risk of default is present (Almeida and Philippon 2007). This can lead to inaccurate valuations when the cost of distress is high. This was the case with the current financial crisis, where the risk-free rates went down but the cost of capital for firms increased. Thus, the risk neutral valuation would advice firms to invest more in a time where firms' capital costs increased, which is clearly the wrong advice. However, in more stable conditions the distortions related to the risk of default should be low, specially when considering large petroleum companies.

4.4 Case Study

4.4.1 Input data

In this section, we use the model developed in previous sections to value two real options connected to an offshore oil project with the Least Squares Monte Carlo algorithm developed in (Longstaff and Schwartz 2001). An overview of the algorithm can be found in Appendix A.1. We use risk neutral pricing.

Table 15: Financial parameters

S_0	-	Current oil price	-	USD 60
r_f	-	Risk free rate of return	-	4.3%
μ	-	Risk neutral growth rate	-	2.7%
σ	-	Annual volatility oil price	-	29.5%
K_E	-	Expansion option strike/Total cost tie-in field	-	MUSD 600
K_A	-	Early decommissioning option strike/Initial FPSO sales price	-	MUSD 500

For a more detailed explanation of the timing and distribution of investments in an offshore petroleum field produced with an FPSO see (Robinson 2009).

4.4.2 Expansion option

The option to invest in a tie-in field takes the form of a call option, as discussed in Section 4.1.1. To acquire this option the operator might have to invest in extra deck-space or other forms of extra capacity today, denoted C_{tie-in}.

Table 16: FPSO and reservoir parameters

R_0	-	Initial reservoir reserves	-	300 MMbbl
R_{tie-in}	-	Initial tie-in reserves	-	15 MMbbl
q_w	-	Maximum well production	-	66 MMbbl/yr
q_p	-	Platform production capacity	-	33.17 MMbbl/yr
T	-	Field life time	-	18
I_{FPSO}	-	FPSO Investments	-	MUSD 561.5
I_{Tot}	-	Total Investments	-	MUSD 2,228
$T_{Ramp-up}$	-	Production Ramp-up time	-	3 Years
σ_{tie-in}	-	Uncertainty in tie-in reserves	-	10%

Table 17: Monte Carlo parameters

N	-	Number of realizations	-	100 000
M	-	Number of time points	-	100

This will be the cost of getting the real option, and should not be confused with K_E which is the investment needed when the tie-in is connected. Using the input data in the previous section and taking the price growth into account, we find that the maximum static NPV is obtained at T=8 which is the last year of plateau production. However, after deducting investment costs the NPV is MUSD -344 at the optimal investment time. This corresponds to a NPV of MUSD -262 at t=0. In a deterministic setting, it does not pay off to produce the satellite and based on this the operator should not invest in excess capacity in order to have the opportunity.

When we add price uncertainty the answer changes. By valuing the investment opportunity as an American call option on the incremental production, the option to invest is estimated to be worth MUSD 150. This implies that if the investment needed today, C_{tie-in}, is less than MUSD 150, the operator should invest in order to have the option. This helps explain why operators frequently invest in extra capacity, since having the opportunity of producing nearby satellite fields creates valuable real options.

Adding further uncertainty by introducing uncertainty in production, the option value is still in the same range as before with an option value of MUSD 161. The lower contribution is not surprising, as the variation in production is lower compared to price variation and the production follows a mean reverting process rather than a Brownian motion.

4.4.3 Sensitivity analysis

As we can see from Figure 11a, the option value increase with increasing initial oil price. Unlike a static NPV calculation the option value increases nonlinearly with low initial oil prices, but the growth becomes linear at higher prices. This is natural, as the tie-in is almost certain to be developed at high prices, and the extra value from the option is low. In this case, the option value is almost equal to a static NPV. However, unlike the static NPV the option value is never negative. Because the operator has the choice but not the obligation to develop the tie-in, it will never be developed if it has a negative NPV.

Another important variable is oil price volatility, and the option sensitivity to this variable can be seen in Figure 11b. The option does not have any significant value for volatilities below 5% per year, and this confirms the conclusion that the project would not have positive NPV in a static valuation method. That the value of a project should increase with larger volatility is contrary to the common view. The crucial difference between real option valuation and a discounted cash flow approach is that the project owner has the option to not exercise the option. Thus the owner is protected from the case where the price falls, since the satellite field will not be developed in this case. High volatility increases the value because it increases the probability of a very high payoff, without increasing the probability of a large loss. However, higher volatility will increase the optimal exercise price and delay the investment time. This is because one needs to have a price high above the break-even price to be certain that the price will not drop to a level where the project has a negative NPV when the volatility is high. Also, we observe that the volatility has less effect on the option value than the initial oil price.

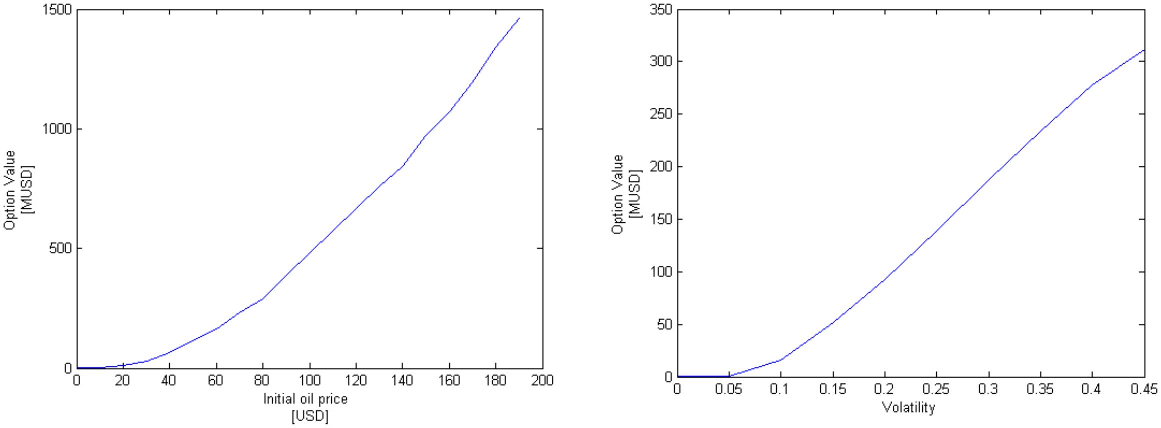

(a) Initial oil price

(b) Oil price volatility

Figure 11: Expansion option value sensitivities

4.4.4 Early decommissioning option

The opportunity of decommissioning the field prematurely could be a response to lower production volume than expected, or very low oil prices. The operational costs of an oil project are often low compared to the investment cost, and the value of being able to prematurely abandon the field is believed to be low. Abandoning the field at T=18 and selling the FPSO has in the deterministic case a NPV of MUSD 59.2.

When only including price uncertainty the option value is MUSD 5.4, after deducting the deterministic payoff at the last period. This is about 0.25% of the initial investment. Adding uncertainty in the reservoir reserves, we obtain an option value of MUSD 5.6. We conclude that the option of abandoning the field prematurely is not very valuable, and that the flexibility of being able to move an FPSO can be disregarded when choosing production technology.

4.4.5 Sensitivity analysis

Since the decommissioning option is similar to a put option, we expect the option value to decrease with rising oil prices. This is also the case, as can be seen in Figure 12a. Unlike a regular put, the option is worth more than the strike price as the oil price approaches zero. This is because as the project is abandoned the operator also avoids the operating costs. The option value of abandoning is high when the oil price is low, but since the project as a whole will have a negative NPV it will not be built in the first place. Also, we have assumed that the value of an FPSO is deterministic. A more realistic assumption would be that the sales price is positively correlated with the oil price, as few new projects will be initiated if the price is low. This will further reduce the value of early decommissioning. For initial prices close to todays price the option value is negligible compared to the investment. The option value is sensitive to the price volatility, as seen in Figure 12b. If the price volatility should continue to increase in the future, decommissioning options could become valuable.

5 Conclusion

In this paper we study the flexibility related to investment timing in offshore oil exploration and production. The oil price is the main source of risk that influence the value of real options related to the project. We find that a geometric Brownian motion is an adequate model for the long term oil price, and that it is better suited than a mean reverting model. It is shown that the option to abandon by moving the FPSO is not significant compared to the cost of developing the field. The option to expand the production by adding new fields adds value and the value of making initial investments in order to be able to connect such satellite fields in the future can be large even when the current NPV from the satellite fields are negative. Both options increase in value when faced with

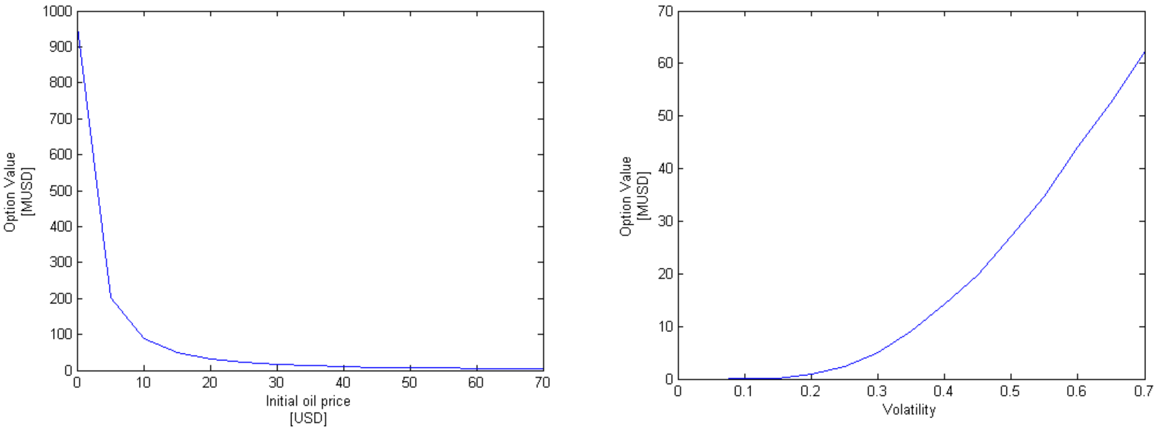

(a) Initial oil price

(b) Oil price volatility

Figure 12: Abandonment option value sensitivities

increased volatility.

For further work, exploring if more advanced price models leads to different option valuations would be an interesting extension. This includes studying the effect on option value from drawing from a skewed student-t distribution closer resembling the observed returns. Another extension related to the option value framework would be to introduce a stochastic process governing when and if a tie-in field is found. This would be more general than our assumption that the operator knows from the start if there is a nearby field.

Acknowledgements

We would like to thank our academic supervisor professor Stein-Erik Fleten for his valuable help and advice. Further, we would like to thank research fellow Vidar Gunnerud and Marta Dueas Diez from Repsol YPF for helping us understand the issues related to petroleum production and their advice and encouragement, as well as the Center for Integrated Operations in the Petroleum Industry for support. All errors are solely our responsibility.

References

H. Almeida and T. Philippon. The risk-adjusted cost of financial distress. *Journal of Finance*, 62(6):2557–2586, 2007.

J. P. Brashear, A. B. Becker, and Faulder D. D. Where have all the profits gone? Or, evaluating risk and return of E&P projects. *SPE Annual Technical Conference and Exhibition, 1-4 October 2000, Dallas, Texas*, 2000.

C. Brooks. *Introductory Econometrics for Finance*. Cambridge, 2008.

G. T. Castro, C. K. Morook, and S. N. Bordalo. Decision-making process for a deepwater production system considering enviromental, technological and financial risks. Technical report, SPE Annual Technical Conference in San Antonio, Texas, September 2002.

L.G. Chorn and S. Shokhor. Real options for risk management in petroleum development investments. *Energy Economics*, 28 (4):489 – 505, 2006.

G. Cortazar and E. S. Schwartz. Monte Carlo evaluation model of an undeveloped oil field. *Journal of Energy Finance & Development*, 3(1):73–84, 1998.

G. Cortazar, M. Gravet, and J Urzua. The valuation of multidimensional American real options using the LSM simulation method. *Computers & Operations Research*, 35(1):113 – 129, 2008.

G. A. Costa Lima and S. B. Suslick. Estimation of volatility of selected oil production projects. *Journal of Petroleum Science and Engineering*, 54(3-4):129 –139, 2006.

M. A. G. Dias. Valuation of exploration and production assets: an overview of real options models. *Journal of Petroleum Science and Engineering*, 44(1-2):93 – 114, 2004.

M. A. G. Dias, J. G. L. Lazo, M. A. C. Pacheco, and M. M. B. R. Vellasco. Real option decision rules for oil field development under market uncertainty using genetic algorithms and Monte Carlo simulation. *7th Annual Real Options Conference, Washington DC*, 2003.

D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366):427–431, 1979.

A. K. Dixit and R. S. Pindyck. *Investment under uncertainty*. Princeton University Press, 1994.

M. Dueas Diez. Telephone conference with Marta Dueas Diez, Repsol YPF, November 2009.

J. Durbin and G. S. Watson. Testing for serial correlation in least squares regression, II. *Biometrika*, 38(1-2):159–177, 1951.

EcoWin-Reuter. Ecowin reuter database. Retrived 02.09.2009, 2009.

S. Ekern. A option pricing approach to evaluating petroleum projects. *Energy Economics*, 10(2):91–99, 1988.

A. Gamba. Real options: a Monte Carlo approach. *EFA 2002 Berlin Meetings Presented Paper*, 2003.

P. Glasserman and B. Yu. Number of paths versus number of basis functions in american option pricing. *The Annals of Applied Probability*, 14(4):2090–2119, 2004.

M. K. Hubbert. Nuclear energy and the fossile fuels. Technical report, American Petroleum Institute Drilling and Production Practice, Spring Meeting, San Antonio, Texas, March 1956.

ICE. Daily volumes for ice brent crude options. Retrived 11.11.2009, 2009. URL `https://www.theice.com/marketdata/reports/`.

C. M. Jarque and A. K. Bera. Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 6(3):255–259, 1980.

Wall Street Journal. Markets data center. Retrived 11.09.2009, 2009. URL `http://online.wsj.com/`.

A. Kemp and L. Stephen. The economics of petroleum exploration and development west of Scotland. *Continental Shelf Research*, 21(8-10):1095 – 1120, 2001.

D. Lamper and S. Howison. Monte Carlo valuation of American options. OFRC working papers series, Oxford Financial Research Centre, 2003.

D. Laughton, R. Guerrero, and D. R. Lessard. Real asset valuation: A back-to-basics approach. *Journal of Applied Corporate Finance*, 20(2):46 – 65, 2008.

F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least-squares approach. *Rev. Financ. Stud.*, 14(1):113–147, 2001.

D. Luenberger. *Investment Science*. Oxford University Press, 1998.

M. W. Lund. Real options in offshore oil field development projects. In *3rd Annual Real Options Conference, Leiden*, 1999.

K. F. McCardle and J. E. Smith. Valuing oil properties: Integrating option pricing and decision analysis approaches. *Operations Research*, 46(2):198 – 217, 1998.

R. L. McDonald. *Derivatives Markets*. Pearson/Addison-Wesley, 2005.

D. Pilipovic. *Energy Risk; Valuing and Managing Energy Derivatives*. McGraw-Hill, New York, 1998.

R. S. Pindyck. The long-run evolution of energy prices. Technical report, 1998.

T. Rehrl and R. Friedrich. Modelling long-term oil price and extraction with a Hubbert approach: The LOPEC model. *Energy Policy*, 34(15):2413 – 2428, 2006.

R. Robinson. The economic impact of early production planning (EPP) on offshore frontier developments. Technical report, Offshore Mediterranean Conference and Exhibition in Ravenna, Italy, March 2009.

B. F. Ronalds. Applicability ranges for offshore oil and gas production facilities. *Marine Structures*, 18(3):251 – 263, 2005.

E. S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52 (3):923–73, 1997.

Y. Shimamura. FPSO/FSO: State of the art. *Journal of Marine Science and Technology*, 7(2):59–70, 2002.

D. R. Siegel, J. L. Smith, and J. L Paddock. Valuing offshore oil with option pricing models. *Midland Corporate Finance Journal*, (5):22–30, 1988.

J. E. Smith. Alternative approaches for solving real-options problems. 2(2):89–102, 2005.

A Appendix

A.1 Appendix 1 - The Longstaff-Schwartz model

A.1.1 Introduction to the model

A popular way of valuing options is using simulation-based methods. Simulation procedures were early developed to price European derivatives, but because of the complexity of American style options simulation procedures for these options have been developed later (Lamper and Howison 2003). The main issue with valuing American style claims is defining the optimal exercise boundary, since this requires a backward looking algorithm and the original Monte Carlo simulation are forward-looking (Gamba 2003). Longstaff and Schwartz (2001) developed a procedure to value American put and call options by Monte Carlo simulations. The procedure uses a least squares regression for each time-step to estimate the continuation value of the option in each time step. If the value of exercising the option is higher than the estimated gain of continuing the option, the option is exercised in that time step. This approach is useful in multi factor situations where the analytical solution is more complex. Gamba (2003) extended the procedure presented by Longstaff and Schwartz (2001) to value more specific real options. Cortazar et al. (2008) has extended the approach to cover multi factor risk models more suitable for long-term commodity real options. This approach is similar to the dynamic programming approach, in that it uses a predetermined discount rate to value the underlying project. It has an advantage in that its complexity does not grow with the number of uncertainties. Thus, it should be a good alternative for petroleum projects, where one can have uncertainty in i.e investment cost, volume produced and the output price. The complexity does increase with the complexity of the options and the number of choices (Smith 2005).

A.1.2 Valuation algorithm

The LSM algorithm uses a least square regression to approximate the expected continuation value of the option. This value is conditional on the current value of the option. The algorithm starts at $t = T$ by calculating the exercise value, $V_{n,T} = max(S - K, 0)$. The option is exercised if and only if the value of exercising is positive. This provides the input for the continuation value regression in the previous period. Here the value of keeping the option alive calculated by approximating the continuation value, $F_{n,T-1}$, by $F_{n,T-1} = a_{T-1} * X_{n-1,T}$, where X is a set of basis functions and a_t represent the regression coefficients from Equations 25. The basis function coefficients are found solving the regression which minimize the sum of squares in the following regression,

$$Y_t = a_t * X_t \tag{23}$$

$$Y_{n,t} = max(S_{n,t} - K, 0) \tag{24}$$

$$X = [X^1_{n,t}, X^2_{n,t}, .., X^j_{n,t}] \tag{25}$$

By only including the realizations where the option exercise is positive in the regression, the accuracy of the algorithm is increased significantly (Longstaff and Schwartz 2001).

A.1.3 Choice of basis functions

One of the elements the model builder needs to take into account is the set of functions that the continuation value will be regressed on. (Longstaff and Schwartz 2001) suggests several different basis functions, like Laguerre polynomials or even simple polynomials or trigonometric functions. After deciding upon a set of basis functions, the model builder must decide on the number of basis functions. (Longstaff and Schwartz 2001) tests the robustness of the LSM algorithm to the choice of basis function(s), and finds that it is very robust, with regards to the choice of polynomials. They also prove that if $N \rightarrow$ inf and with a large number of basis functions the LSM algorithm converges to the true option value, where N is the number of paths in the Monte Carlo simulations. In our implementation we have chosen to use three Laguerre polynomials seen in Equations (26), (27) and (28), which is similar to those used in (Longstaff and Schwartz 2001). Glasserman and Yu (2004) show that the number of polynomials must grow at most with $\sqrt{log(N)}$.

$$X^1 = (-S+1) \tag{26}$$

$$X^2 = \frac{1}{2}(S^2 - 4*S+2) \tag{27}$$

$$X^3 = \frac{1}{6}(S^3 + 9*S^2 - 18*S+6) \tag{28}$$

A.1.4 Scenario realizations

To produce input to the model we used the estimators from Table 10 and Table 11 from page 11 to create N realizations of the oil price for a T year horizon. Each time period is divided into M parts, so that $\Delta T = \frac{T}{M}$. When the price is modeled as a geometric Brownian motion the matrix is populated by drawing M*N draws from a standard normal distribution, $dZ \sim N(0,1)$.

$$S_{t,i} = S_{t-1,i} * (\mu - \frac{1}{2} * \sigma^2) * \Delta T + \sigma * dZ, \forall t \in M \geq 2, \forall i \in N \tag{29}$$

Each of the realizations is then multiplied with the value function, either Equation 13 or Equation 12
This produces the following $M*N$ matrix.

$$\begin{pmatrix} V_{0,1} & V_{1,1} & .. & .. & V_{M,1} \\ V_{0,2} & V_{1,0} & .. & .. & V_{M,2} \\ .. & .. & .. & .. & .. \\ .. & .. & .. & .. & .. \\ V_{0,N} & V_{1,N} & .. & .. & V_{M,N} \end{pmatrix}$$

A.2 Model analysis

A.2.1 Stationarity

Determining if the series is stationary or not is very important when determining its properties. The series is strictly stationary if its statistical distribution is the same at all times. It is weakly stationary if it satisfies the three equations (Brooks 2008):

$$E(y_t) = \mu \tag{30}$$

$$E(y_t - \mu) * (y_t - \mu) = \sigma^2 < \infty \tag{31}$$

$$E(y_{t_1} - \mu) * (y_{t_2} - \mu) = \gamma_{t_2 - t_1} \tag{32}$$

The three equations state that the series should have a constant mean and variance, and in addition have constant autocorrelations.

A.2.2 Goodness of fit statistic, R^2

Goodness of fit statistic is statistic that describes how well a model fits to the data given. The most common goodness of fit statistic is the R^2. R^2 can be described as the square of the correlation between the values of the dependent variable and the corresponding fitted value from the model (Brooks 2008). R^2 lies between 0 and 1. If R^2 is close to 1 the model fits the data well, if it is close to zero then the model do not provide a good explanation for the variation in the data.

R^2 is given by:

$$R^2 = \frac{ESS}{TSS} = \frac{TSS - RSS}{TSS} = 1 - \frac{RSS}{TSS} \tag{33}$$

Here TSS is the sum of squared deviations from the mean, RSS is the sum of squared residuals from the regression, and ESS is the difference between the two.

A.2.3 Durbin-Watson test for autocorrelation

An assumption used in linear regression is that the covariance between the error terms over time is zero, stating that the errors should be uncorrelated with another. If they are not, they are said to be autocorrelated. With autocorrelation the estimators are still unbiased, but the estimated standard errors for the parameters will be too small. This will make estimators seem more accurate than they really are, thus making hypothesis testing difficult, as well as inflating the R^2 statistic.

The Durbin-Watson (DW) is a test for first order autocorrelation which tests if there are correlation between an error and the previous error value (Durbin and Watson 1951). The test has the null hypothesis that there is no autocorrelation, $H_0 : \rho = 0$, and the alternative hypothesis $H_1 : \rho \neq 0$.

The test statistic can be calculated as:

$$DW = \frac{\sum_{t=2}^{T} (\hat{u}_t - \hat{u}_{t-1})^2}{\sum_{t=2}^{T} \hat{u}_t{}^2} \tag{34}$$

$\hat{u}_t$ denotes the the residual from the regression in period t. The DW test does not follow a standard statistical distribution. DW has 2 critical values, an upper value and a lower value. There is also an intermediate region where the null hypothesis can neither be rejected nor not rejected (Brooks 2008, p. 165). At the region between the lower critical value and "4 - lower critical" the hypothesis cannot be rejected, there is no clear evidence of autocorrelation. If the test statistic falls in the region between the lower and upper critical regions the test is inconclusive.

A.2.4 Unit Root Test for stationarity

A unit root test can be used to see if the series is a stationary process or a non-stationary process. It is important to see if the oil price is stationary or not, because this will affect the series behavior and properties. For a stationary series the value of a variable is expected to revert to a mean value, which is not the case for a non-stationary series. A well known example of a non-stationary process is the random walk with or without drift. To test if the price of oil is a mean reverting process or a random walk, one can use a unit root test. However, for short time series a unit root test will often fail to reject the hypothesis of a random walk, even if the underlying process is mean reverting.

To test if a series contain a unit root, meaning that it is non-stationary, one can use the Dickey-Fuller test from (Dickey and Fuller 1979). The test has a null hypothesis H_0 : The series contains a unit root in the following regression:

$$\Delta y_t = \psi y_{t-1} + u_t \tag{35}$$

which is the same as $\psi = 0$ against the one sided alternative $\psi < 0$, or H_1:series is stationary.

The model for the unit root test is:

$$\Delta y_t = \psi y_{t-1} + \mu + \lambda t + u_t \tag{36}$$

and the test statistics for the DF tests are defined as

$$\text{test statistic} = \frac{\hat{\psi}}{S\hat{E}(\hat{\psi})} \tag{37}$$

The test statistics follows a non-standard distribution since the null hypothesis is that the series is non-stationary (Brooks 2008). The critical values are derived from simulation experiments, and can be found in (Brooks 2008).

If the test statistic is more negative than the critical value, the null hypothesis is rejected. This test is only valid if u_t is white noise.

A.2.5 Jarque-Bera test for normality

A normal distribution is characterized only by its mean and its variance. The Jarque-Bera test (JB) uses this property of a normal distribution to test if the distribution is truly normal (Brooks 2008, p. 161). A normal distribution is symmetric around its mean value. This is formally stated by requiring that the skewness, the extent to which the distribution is non-symmetric, is zero. A normal distribution also has a kurtosis of 3, which is a measure of how fat the tails of the distribution are.

The JB-test tests if the coefficient of skewness and the coefficient of excess kurtosis are both zero. u represents the residuals, and σ^2 the variance, the skewness and kurtosis can be expressed as (Jarque and Bera 1980):

$$Skewness = b_1 = \frac{E[u^3]}{(\sigma^2)^{3/2}} \tag{38}$$

$$Kurtosis = b_2 = \frac{E[u^4]}{(\sigma^2)^2} \tag{39}$$

The Jarque-Bera test statistic is:

$$W = T\left[\frac{b_1^2}{6} + \frac{(b_2-3)^2}{24}\right] \tag{40}$$

where T is the sample size. To test if the distribution is normal, one first assume a null hypothesis that the distribution is symmetric and have an excess-kurtosis, $b_2 - 3$, equal to zero. If the test statistic W is greater than the critical value of a $\chi^2(2)$, the null hypothesis of normality is rejected.

Often a couple of extreme outliers may be the cause of a normality rejection. One way to improve the chance of the errors to be normally distributed is to remove the outliers by introducing dummy variables for the outlier observations(Brooks 2008, p. 165).

A.3 GBM parameter estimators

In using discrete time steps dT becomes Δ_t.

$$r = \frac{ln(P_t)}{ln(P_{t-1})} \tag{41}$$

$$E(r) = (\mu - \frac{\sigma^2}{2})\Delta_t \tag{42}$$

$$Var(r) = \sigma^2\Delta_t \tag{43}$$

$$s_r = \sqrt{\frac{1}{n-1}\sum_{t=0}^{T}(r_t - \overline{r})^2} \tag{44}$$

$$\widehat{\sigma} = \frac{s_r}{\Delta_t} \tag{45}$$

$$\widehat{\mu} = \frac{\overline{r}}{\Delta_t} + \frac{s_r^2}{2\Delta_t} \tag{46}$$

A.4 Mean reverting parameter estimators

Then to estimate (9) one can use the regression:

$$P_t - P_{t-1} = a + bP_{t-1} + \varepsilon_t \tag{47}$$

and calculate

$$\hat{\alpha} = \frac{-\hat{a}}{\hat{b}} \tag{48}$$

$$\hat{\lambda} = -log(1+\hat{b}) \tag{49}$$

$$\hat{\sigma} = \sigma_{\varepsilon}\sqrt{\frac{2log(1+\hat{b})}{(1+\hat{b})^2-1}} \tag{50}$$