Stockholm School of Economics / Università Commerciale Luigi Bocconi
MSc in Finance Double Degree Program

Master Thesis

Information content and pricing of options in a jump-diffusion setting

Kristaps Vasiļjevs

Consists of:

Pricing of Asian options in a jump-diffusion setting

and

The importance of selective market information in calibration and option pricing

May, 2016

Abstract

It is believed that information that is incorporated within market prices is accurate and useful. To evaluate this, I, first of all, calibrate the Merton jump-diffusion model to oil options over the period from 2009 to 2015. I show that the retrieved parameters capture market events properly and appropriately. This is most evident during turbulent periods and is most pronounced for parameters retrieved from options with shorter maturities. In addition, the results indicate that in practice the parameters are of little use in predicting future prices of the underlying. Secondly, I price oil Asian options based on the retrieved parameters and compare these prices with the ones observed in the market. The results show that the relative pricing error fluctuates noticeably, but always stays within acceptable levels. In addition, when Asian options are priced using parameters that are retrieved from options with similar moneyness, the relative pricing error is substantially improved.

Table of contents

1 Introduction 4
2 Literature review 7
2.1 Oil market 7
2.2 Dynamics 7
2.2.1 Development of models 8
2.2.2 Merton jump-diffusion model 9
2.3 Applications of Merton jump-diffusion model 10
2.3.1 Implied information 11
2.3.2 Pricing 12
3 Methodology 13
3.1 Calibration 14
3.1.1 Pricing mechanism 14
3.1.2 Inverse problem 15
3.1.3 The choice of weights 16
3.1.4 Jump intensity 17
3.2 Pricing 18
3.2.1 Asian option 18
3.2.2 Monte Carlo simulation 19
4 Data description 20
4.1 Calibration dataset 20
4.2 Pricing dataset 20
4.3 Adjustments 21
5 Analysis and discussion 21
5.1 Oil returns 21
5.2 Implied information 22
5.2.1 Calibration and economic significance 23
5.2.2 Forecasting 28
5.3 Pricing Asian options 31
5.3.1 Separate calibration and pricing 34
5.4 Sensitivity analysis 36
6 Conclusions 40
7 References 42
8 Appendix: sensitivity surfaces 45

1 Introduction

In the finance industry one of the most significant and hardest decisions is the selection of the most appropriate model for analysis as it gives the foundation on which all recommendations are based. Hence, an incorrect model can result in flawed decisions with far-stretching consequences as the financial crisis of 2008 showed. A significant factor that made the collapse possible in 2008 was rating agencies, which employed models that strongly underestimated the risk of mortgage backed securities, therefore, provided an inaccurate view on the market.

In derivatives pricing a model will determine how the price of the underlying asset evolves over time; hence, a model will indirectly determine the price of a derivative. In option pricing the most widely used model is the Black and Scholes (1973) model, which assumes a log-normal diffusion with a constant volatility. The empirical literature strongly rejects pure diffusion models for modelling asset prices. This is primary because implausibly high levels of volatility must be assumed in order for a pure diffusion model to capture sharp changes in returns. Hence, academics and practitioners have developed various extensions, one of which is adding a jump component to the diffusion term. The empirical literature (see, for instance, Bakshi, Cao and Chen, 1997; Broadie, Chernov and Johannes, 2007) provides a strong support in including jumps because jump-based models are able to capture the leptokurtic nature of asset returns. In particular, Kearney, Murphy and Cummins (2015) show that the Merton jump-diffusion (MJD) model, the model used in the thesis, very well captures the dynamics of oil returns.

The MJD model has four parameters that must be estimated. One possibility is to select these parameters based on one's own view; however, this is subject to a substantial bias. Another approach is to retrieve the parameters by calibrating the MJD model to the market so that option prices implied by the model coincide with market prices. Nevertheless, prior to turning to pricing derivatives, one must make sure that these retrieved parameters are correct, i.e. since the parameters are retrieved from market information, they must correspond to what happened in the market at that point in time. In addition, since changes in the market are occasionally caused by self-fulfilling expectations, the retrieved parameters can be useful in predicting future prices of the underlying. With this research I contribute to the

existing discussion of information content in the market and answer the following first research question: *Do market prices contain valuable information?*

To answer the first research question, I employ a dataset that consists of all quoted call option prices on NYMEX over all available strikes written on oil futures from October 2009 to May 2015. This dataset is accompanied by corresponding futures contracts and interest rates.

The empirical work that looks whether parameters retrieved from the market explain what happened in the market at that point in time is scarce. The closest research is done by Murphy and Ronn (2015), who calibrate the MJD model to oil option prices over the period from 2007 to 2012. They find that the retrieved parameters strongly agree with the situation in the market. However, their analysis focuses mainly on one event, the Arab Spring. In the empirical research more attention has been devoted to see whether retrieved parameters are also useful for predicting the future. Kearney *et al*. (2015) also look at the oil market using the MJD model and find signals of events before they occur; however, they do not analyse whether these signals can be used by practitioners.

I find that the retrieved parameters very well explain what happens in the market. In particular, the average jump process amplitude increases (decreases) with an escalating (declining) oil price while the diffusion volatility and the standard deviation of jump process amplitude rises whenever there is a pronounced volatility in oil returns. Additionally, the parameters capture also smaller changes in oil price, and all movements in the parameters are in line with theory. Regarding the predictability, I find that a relatively simple strategy based on the retrieved diffusion volatility results in a positive, but marginal gain before transaction costs. Hence, the results do not provide evidence that the future can be predicted from the past market information.

Once reasonable parameters are estimated, they can be used in the MJD framework to price derivatives. In order to be able to evaluate whether the MJD model prices options correctly, comparison must be made between estimated prices and market prices. However, considering that the parameters are retrieved from oil options, they are valid only for pricing other derivatives on oil. I contribute to the existing discussion on pricing derivatives and answer the following second research question: *Is information implied in market prices useful for pricing exotic derivatives?*

To answer the second research question, I employ a dataset that consists of all quoted Asian call option prices on NYMEX over all available strikes written on oil

futures from December 2011 to May 2015. The reason for pricing particularly Asian options is that Asian options have a great practical use and they are the most liquid exotic oil derivative quoted on the market.

To my knowledge, there is no empirical research that price exotic options using retrieved parameters and compare their prices with the market. In order to evaluate the performance of pricing, I look at the relative pricing error between estimated prices and prices quoted on the market. I find that the relative pricing error fluctuates over time, but within acceptable levels. The fluctuations are most pronounced for options with shorter maturity and are most evident when there is alleviated volatility in the market. In particular, I show that market prices always lie within the range of 95% confidence interval around the estimated prices of in-the-money (ITM), at-the-money (ATM) and out-of-the-money (OTM) options. In addition, I show that the relative pricing error is substantially improved when options are priced using parameters that characterise these options. That is, the relative pricing error is smaller if ITM options are priced using the parameters retrieved from ITM options. Finally, through sensitivity analysis I show that in the current market environment the diffusion volatility has the highest impact on the price of an Asian option. Therefore, the diffusion volatility must be most carefully examined before the MJD model is applied in the practice.

My work contributes to the existing body of literature in the following four ways: (i) retrieved parameters capture market events properly and appropriately; (ii) weak-form market efficiency cannot be rejected; (iii) a model calibrated to the market is useful for pricing exotic derivatives; (iv) pricing performance is improved when calibration is based on selected market information.

The paper proceeds as follows. In section 2 I present relevant literature review and hypotheses I seek to find evidence for. In section 3 I describe methodology used in the research and the motivation for choosing the particular approach. Section 4 presents my dataset with applied adjustments. In section 5 I describe the results of the study and discuss their implications. I conclude with section 6.

2 Literature review

In this section I give a short overview of the oil market, highlighting an important difference as compared to capital markets. It is followed by a brief review of models used in pricing. Particular attention is paid to the model employed, namely, the Merton jump-diffusion model. Afterwards, I review papers that have previously analysed information content obtained from calibration. Finally, I investigate papers that have looked at the pricing issue.

2.1 Oil market

On a given day more than a million light sweet crude oil futures and option contracts are traded, making oil the most frequently traded commodity (CME Group, 2016a). The importance of oil cannot be overstretched. Not only oil is the single most important source of income for some countries, but also it has a profound impact on daily lives of most societies. Therefore, policy makers follow oil price very carefully before making decisions. For example, since oil has a weighty impact on inflation, oil price is an important consideration for central banks before changing key rates, as this decision has even more profound impact on the economy (Kearney *et al*., 2015).

An important difference between commodity markets and capital markets is that in the former there is not only a fear of a price drop, but also a fear of prices going upwards. As a result, the implied volatility curve for oil is distinctively U-shaped, while, for example, the implied volatility curve for S&P 500 index is characterized by a “sneer”, where implied volatilities decrease with increasing strikes (Kearney *et al*., 2015). The observed phenomenon in the oil market is due to oil being highly price inelastic. This implies that even a small deviation from demand and supply equilibrium requires a large change in price for the market to clear (Askari and Kirchene, 2008).

2.2 Dynamics

Taking into account the specifics of the oil market, it should be expected that modelling oil price dynamics by assuming a log-normal diffusion with a constant volatility as in the Black and Scholes (1973) model is incorrect. If market participants would follow the Black and Scholes assumptions, the implied volatility curve for a given asset would be flat, i.e. the implied volatility would be the same for every strike. However, as pointed out in the previous sub-section, the implied volatility curve takes

distinctive shapes. It is thus in general necessary to input a larger volatility for a lower strike as compared to the ATM strike for the Black-Scholes formula to hold. Looking from different perspective, since the price of an option increases with volatility, prices for options that have low strikes are larger than theory suggests. The existence of this observed pattern in the markets is attributed to a fear of large downward movements: market participants purchase OTM put options to insure themselves from a drop in the price of an asset. This phenomenon is observed in commodity markets, as well as, in capital markets. In addition, since oil market participants also fear of a price hike, OTM call options also trade at higher prices than suggested by theory. The latter effect is not pronounced in capital markets. As a result, the implied volatility curve in capital markets is decreasing with increasing strikes and is characterised by a “sneer”, while the implied volatility curve in the oil market is referred as a “smile” (Andersen and Andreasen, 2000).

2.2.1 Development of models

Given that the Black and Scholes model is flawed as it is unable to capture the presence of a volatility smile or skew, various extensions to the model have been developed over time. These extensions can be clustered in three groups.

The first approach is to stay in the Black-Scholes framework with one-factor diffusion, but to set volatility to be a deterministic function of time and price of the underlying asset. This method was first developed by Derman and Kani (1994), Dupire (1994), and Rubinstein (1994) and is referred to as a deterministic volatility function (DVF). This approach has gained popularity because in principle it allows for a precise fit to observed market prices while retaining market completeness (Andersen and Andreasen, 2000). However, DVF models experience several shortcomings. The DVF framework predicts that the implied volatility smile will fade in the near future while in practice the curvature is rather stationary over time. Also, the implied volatility function must take rather dramatic form to fit short-term options. This poses significant challenges for pricing as well as for hedging. Therefore, it is quite comprehensible that the empirical evidence rather strongly rejects the DVF approach in modelling (e.g., Aït-Sahalia, Wang, and Yared, 2001; Dumas, Fleming and Whaley, 1998).

Another approach is to set volatility to be a random process itself resulting in stochastic volatility models. Typically, in this approach volatility is modelled as a

mean reverting diffusion process that is correlated with the dynamics of the underlying asset (e.g., Heston, 1993; Hull and White, 1987). Since stochastic models are multi-factor in nature, a multitude of volatility curves can be generated that can also resemble the observed smiles and skews in practice. In addition, the empirical literature supports the existence of stochastic volatility in asset returns. However, usually implausibly high correlations between volatility and returns are required for stochastic models to generate volatility curves that match the observed ones in the market (e.g. Bakshi *et al*., 1997; Bates, 2000). Additionally, stochastic volatility models are difficult to work with from estimation perspective and they do not allow for perfect option hedging (Andersen and Andreasen, 2000).

The third approach in modelling asset prices and by extension also in modelling volatility smile is to add jumps to the Black-Scholes diffusion framework. These models are called jump-diffusion models and, in contrast to the Black-Scholes framework, have discontinuous paths. Two well-known jump-diffusion models are the Kou (2002) double exponential model and the Merton (1976) model.

2.2.2 Merton jump-diffusion model

The Merton (1976) jump-diffusion model is one of the first extensions of the Black-Scholes model and it has gained interest not only from academics, but also from practitioners for several reasons.

First of all, in order for a diffusion model that follows Brownian motion to generate a large movement in the price of the underlying in a short time period, volatility must be set to an improbable level. However, if the price of the underlying is allowed to jump, then even in a short time span a reasonable probability exists for a large price swing. Building on this intuition, Merton (1976) developed a model, where diffusion behaves like Brownian motion with a constant drift and volatility, while jumps follow Poisson process. The model generates asset prices that are more aligned with the empirical distribution. In particular, the MJD outperforms continuous diffusion models during turbulent periods, as the MJD results in fatter tails (Kearney *et al*., 2015).

Secondly, if the dynamics in the market would follow a continuous diffusion process, then the market would be deemed complete and, consequently, every terminal payoff could be perfectly replicated. This would mean that financial derivatives, as options, are redundant from theoretical perspective. However, if jumps are added to the market, as it is in practice, then hedging just with the underlying and

money market account is strongly inferior as compared to when also options are added (Tankov and Voltchkova, 2009).

Thirdly, the MJD model also allows to quantify and price jumps. This is very useful from risk management perspective when calculating value at risk or the expected shortfall (Tankov and Voltchkova, 2009).

There is substantial empirical evidence that suggests including jumps in modelling asset prices. Bakshi *et al*. (1997) argue that models with jumps can easily generate variability and asymmetry in returns that capture the curvature for close to maturity options that stochastic models struggle with. Similar results that acknowledge the outperformance of jump models or the importance of adding them to supplement stochastic models are also obtained by e.g. Bates (1996), Broadie *et al.* (2007), Chernov and Ghysels (2000), and Das and Foresi (1996). In particular, Broadie *et al*. (2007) conclude that adding jumps improves fit to the data by as much as a half when compared to a simple stochastic volatility model.

There is also empirical support in adding jumps for modelling particularly oil prices. Askari and Kirchene (2008) analyse the oil market from 2002 to 2006 and find that the MJD model very well captures the empirical return distribution of oil returns that is characterised by a negative skew and excess kurtosis. They find that substantially larger part of variation in oil returns is explained by the jump rather than the diffusion process. Hence, they conclude that a model that is governed by jumps is an adequate model to capture the dynamics in the oil market. Cont and Tankov (2004) also indicate that changes in oil price that are caused by discontinuity are the most important aspect to be considered when modelling. This reasoning leads to the first hypothesis: *a jump component must be included when modelling returns in the oil market.*

2.3 Applications of Merton jump-diffusion model

Once an appropriate model is selected, in this research the MJD model, there are several applications of the model in practice, two of which are of particular interest. The first application is to use the model to infer information from the market. This is strongly associated with the calibration problem as calibration is done on market prices and the retrieved parameters not only convey information regarding concurrent market sentiment, but also can be useful for predicting. The second application is for

pricing and hedging. In this case one can use the retrieved parameters obtained from calibration or can input own view.

2.3.1 Implied information

Option prices can be viewed as the market's expectations of the discounted future possible outcomes for the underlying asset. This is particularly appropriate for a highly liquid option market, as that of oil (McIntyre and Jackson, 2009). Therefore, by inversely applying a pricing model based on selected dynamics to observed option market prices that have the same underlying and maturity, but varying strikes, it is possible to retrieve parameters that are based on market's view. This procedure is called calibration. For example, the Black-Scholes option pricing formula has only one unknown, namely, volatility. Since all other parameters are known, the unknown is very easily retrieved by an inverse recursion. The volatility that equates the Black-Scholes formula with observed market price is called the implied volatility and as a concept is extensively used by practitioners. Also, the implied volatility can carry predictive power. If the volatility for deep ITM call options is larger than for deep OTM call options, this can be considered as an indication that the market places higher probability of a price drop of the underlying. However, in the case of the MJD model, there are four unknown parameters and hence the calibration is rather complex. This issue is looked in a greater detail in the methodology section.

Rather limited amount of empirical work exists that looks how parameters retrieved from the market vary over time and whether they convey some information. Murphy and Ronn (2015) calibrate the MJD model to oil option prices over the period from 2007 to 2012. They look at the evolution of the parameters over time and conclude that the retrieved parameters agree with the situation in the market at that point in time. In particular, they analyse period around the Arab Spring when there was a pronounced turbulence in the market. However, they do not look at how the retrieved parameters interact with each other or whether these parameters also convey information in non-severe environment. Taking into account this gap, I put forward the second hypothesis: *retrieved parameters explain market events over time.*

More attention in the empirical research has been devoted to see whether the retrieved parameters are also useful for predicting the future. Kearney *et al*. (2015) also calibrate the MJD model to the oil market and estimate the slope and skewness

of ITM, ATM and OTM option volatility. They find that during periods of oil price increase, the change in slope of OTM call volatility is increasing because market participants buy OTM calls to hedge from large price spikes while the prices of OTM puts relatively decrease. The opposite is observed when the price of oil drops because OTM puts become more valuable as they offer insurance. They also find that already before there is a change in the trend of oil price, changes in the slope of the option volatility are taking opposite sign, indicating that there is some predictability in the market. Similar results that indicate that the retrieved values contain predictive power has also been obtained by Bates (1991), Rappoport and White (1994), and Yan (2011). In particular, Bates (1991) looks at the prices of OTM calls and OTM puts on S&P 500 futures to test whether the October 1987 crash was expected by the market. He finds that already during the price upsurge in August 1987, the price for 4% OTM puts were around 25% more than for the corresponding calls. In addition, he fits the MJD model to obtain the implied probability density function. He finds that the function starts displaying fat tails in June, which are most pronounced in August. This coincides with OTM puts being more expensive than the corresponding calls. Hence, he concludes that the crash was expected with the expectations reaching peak two months before the crash. Similarly, Rappoport and White (1994) conclude that the crash of 1929 was expected with the expectations being highest just before the crash. It is important to note that the crashes of 1929 and 1987 were very extreme and negative events. McIntyre and Jackson (2009) try to answer whether the market can also predict smaller movements in price. Basing the analysis on the implied returns distribution obtained from S&P 500 index options, they conclude that there is virtually no predictability in the market from 1999 to 2002.

While the results suggest that there might be some predictability, none of the empirical research to my knowledge has examined whether a strategy built on parameters estimated from the market can be applied in practice. If this is possible, then it has serious implications to market efficiency, as it is based on a fact that the past information cannot predict the future. This line of reasoning leads to the third hypothesis: *retrieved parameters do not contain predictive power.*

2.3.2 Pricing

There are two highly important considerations when pricing options: choice of a model and a method. Andersen and Andreasen (2000) show that for path dependent

exotic options, as Asian options priced in this thesis, price is very sensitive to the assumed model dynamics. In particular, they show that jump-diffusion models and DVF models return theoretical prices that are significantly different. While this difference is mostly pronounced for barrier options, it is also highly significant for Asian options. This again highlights the importance of choosing the right model for pricing.

Monte Carlo method has all the needed properties, but is very slow and this is troublesome when a model is built for trading. Consequently, a substantial empirical work in option pricing field has been devoted to develop methods that rapidly converge to the price estimate by Monte Carlo simulation (see, for instance, Bayraktar and Xing, 2011; Christoffersen, Jacobs and Ornthanalai, 2012). However, in order to increase speed, assumptions must be made that impact precision. Therefore, Monte Carlo method will be applied in this thesis.

After an extensive research I have found no empirical analysis that price exotic options and compare their values with market prices. In my view, the main reason for this situation is data availability. Since retrieved parameters from call and put options are valid only for the same underlying and for the same maturity, comparison can only be made with exotic options. But exotic options are usually traded over-the-counter and data for these types of transactions is very scarce. The next best approach where to apply pricing mechanism is hedging. Kearney *et al*. (2015) compare the hedging performance using the Black-Scholes and the MJD model. They find that hedging strategy based on the MJD is superior to the Black-Scholes strategy as the former results in 8% less implementation cost. Andersen and Andreasen (2000), and He, Kennedy, Coleman, Forsyth, Li and Vetzal (2006) also highlight that the MJD model improves hedging performance. This rationale leads to the fourth hypothesis: *the MJD model prices Asian options correctly.*

3 Methodology

The inclusion of a jump component adds three new parameters that must be estimated: jump amplitude, jump variance and jump intensity. The estimation is done by calibrating the MJD model so that its price is equal to the one observed in the market. The methodology how to calibrate the model will be shown in the first part of this section. Once the parameters are retrieved, they can be used to price other

derivatives, as exotic options, on the same asset. This will be covered in the second part of this section.

3.1 Calibration

Calibration is an inverse problem associated with pricing a derivative. It tries to back out parameters, which generate model prices equal to market prices. Therefore, before turning to calibration, one must understand the dynamics of a price process and the mechanism of how a price of a derivative is obtained.

3.1.1 Pricing mechanism

In the MJD model, the stochastic differential equation for the price process of the underlying asset F is given by

$$dF(t) = F(t^-)\big((\mu - q)dt + \sigma dW(t) + dJ(t)\big)\,, \tag{1}$$

where μ is an instantaneous return, q is a convenience yield, σ is an instantaneous volatility, W is a one-dimensional Brownian motion and J is a jump process defined as follows:

$$J(t) = \sum_{j=1}^{N(t)} \big(Y_j - 1\big)\,, \tag{2}$$

where $\big(Y_j - 1\big)$ is the size of j^{th} jump and $N(t)$ is a counting process, which is assumed to be a Poisson process. W, N and Y_j are assumed to be independent. The solution of the stochastic differential equation (1) is given by

$$F(T) = F(0)e^{\left(\mu - q - \frac{1}{2}\sigma^2\right)T + \sigma W(T)} \prod_{j=1}^{N(T)} Y_j\,. \tag{3}$$

Equation (3) has infinite equivalent martingale measures (EMMs); however, Merton (1976) assumes that jump risk is specific, i.e. can be diversified away. This results in selecting a particular EMM for which equation (1) in terms of risk-neutral valuation results in

$$F(T) = F(0)e^{\left(r - q - k\lambda - \frac{1}{2}\sigma^2\right)T + \sigma W^Q(T)} \prod_{j=1}^{N(T)} Y_j\,, \tag{4}$$

where r is a risk-free rate, λ is a jump intensity, i.e. the number of jumps per annum, and k is the average jump process amplitude, i.e. the percentage change in price due to jump process. Since the underlying asset is an oil futures contract and futures are

marked to market on a daily basis, $r = q$ (Björk, 2009). In particular, k is defined as follows:

$$k = e^{\left(\mu + \frac{1}{2}\delta^2\right)} - 1\,, \tag{5}$$

where δ^2 is the variance of jump process amplitude. Therefore, the price of a call option is

$$\begin{aligned} C_0 &= E^Q[(F(T) - X)^+ e^{-rT}] \\ &= \sum_{n=0}^{\infty} e^{-\lambda' T} \frac{(\lambda' T)^n}{n!} E^Q \left[\left(F(0) e^{\left(r_n - q - \frac{1}{2}\sigma_n^2\right)T + \sigma_n \sqrt{T} Z} - X \right)^+ e^{-r_n T} \right], \end{aligned} \tag{6}$$

where $\lambda' = \lambda(1 + k)$, $r_n = r - \lambda k + \frac{n \ln(1+k)}{T}$, $\sigma_n^2 = \sigma^2 + \frac{n\delta^2}{T}$, Z is a standard normal random variable and n the is the number of iterations. Usually 10 iterations are enough for full convergence (Glasserman, 2004). Moreover, the expectation part in equation (6) is the simple Black-Scholes formula with a dividend yield q, risk-free rate r_n and volatility σ_n:

$$BScall(F(0), X, \sigma, T, r, q) = E^Q \left[\left(F(0) e^{\left(r_n - q - \frac{1}{2}\sigma_n^2\right)T + \sigma_n \sqrt{T} Z} - X \right)^+ e^{-r_n T} \right]. \tag{7}$$

Therefore, the price of a call option in the Merton jump-diffusion model in semi-closed form is

$$C_0 = \sum_{n=0}^{\infty} e^{-\lambda' T} \frac{(\lambda' T)^n}{n!} BScall(F(0), X, \sigma, T, r, q)\,. \tag{8}$$

3.1.2 Inverse problem

A model is deemed to be calibrated if price of an instrument, as estimated by the model, corresponds to price quoted on the market. An important underlying assumption is that market quotes price of an instrument correctly, which is not always the case (Coval, Jurek and Stafford, 2009; He *et al*., 2006). However, for exchange traded derivatives as plain vanilla options on a highly liquid asset as oil, the market should be close to perfect weak-form efficiency (McIntyre and Jackson, 2009).

In a practical situation, where prices of call options are not available for all strikes, it is not possible to obtain the parameters by directly inversing the formula. Instead, one can minimize the quadratic pricing error over all the available option prices in order to obtain the vector of parameters:

$$\Phi(\theta) = \arg\inf \sum_{i=1}^{n} w_i \left(C_i^{obs} - C^{\theta}(X_i) \right)^2 . \tag{9}$$

The nonlinear least squares method in equation (9) is widely applied by researchers (e.g. Bakshi *et al*., 1997; Bates, 1996; Dumas *et al*., 1998). The problem with purely minimizing the quadratic pricing error is that the function has multiple local minima and hence is highly sensitive to starting values (Giot, Laurent and Petitjean, 2010). This leads to a calibration problem that is ill-posed. Cont and Tankov (2004) propose to add a regularization parameter based on relative entropy to equation (9). The regularization method imposes a convexity on the function; hence, there is a unique solution. However, the method also adds error because the regularization parameter cannot be purely estimated from data. In addition, Cont and Tankov (2004) argue that regularization method should provide calibration improvement particularly for near the maturity options. However, Matsuda (2005) shows that for options with expiry less than a year there is virtually no difference in retrieved parameters when calibrating with or without relative entropy. Matsuda (2005) also argue that due to the bias imposed by regularization, the best method is to calibrate without regularization over wide set of starting values and select those parameters that have the lowest quadratic pricing error.

3.1.3 The choice of weights

The relative weight w_i is another important consideration. Some researchers (e.g. Christoffersen *et al*., 2012) set the relative weight as the inverse of the absolute bid-ask spread difference. In that way the most liquid options, which are options with strikes close to the price of the underlying asset, have larger impact on minimization than the ones that are deep ITM or deep OTM. However, weighting by bid-ask spread possesses several concerns (Cont and Tankov, 2004). First of all, retrieved parameters are mostly valid for pricing ATM options as there might be important information implied in deep OTM and deep ITM options that is missed. Secondly, bid-ask spreads are not always available. Cont and Tankov (2004) suggest setting the relative weight to inverse of squared vega:

$$w_i = \frac{1}{vega^2\left(\sigma_i^{imp}\right)}, \tag{10}$$

where option's vega is calculated on its implied volatility as quoted by the market. The motivation is that options, which are deep ITM and deep OTM, are less sensitive to change in parameters and therefore those options need larger relative weighting.

By applying this method each option will carry almost the same impact on retrieved parameters.

3.1.4 Jump intensity

Murphy and Ronn (2015) raise an attention on jump intensity when calibrating the MJD model. The expression $r_n = r - \lambda k + \frac{n \ln(1+k)}{T}$ shows that λ and k emerge in multiplicative fashion and are closely dependant. The relationship indicates that for a given r_n a large λ implies a low k and vice versa. Therefore, when tracking the parameters over time, it would be possible to witness a point in time with large λ and small k that is succeeded by unjustifiably small λ and large k. The two states emerge from having imperfect set of available prices and hence it is needed to set bounds on one of the two parameters (He et al., 2006). Since large λ results in many, but small jumps, bounding λ in large values would mean that the jump component covers the diffusion part. Therefore Kearney et al. (2015) and Matsuda (2005) suggest setting λ to small values. Murphy and Ronn (2015) go further and fix $\lambda = 0.3$ over time. The selection of a particular jump intensity is motivated by estimating the frequency of jumps of oil futures and arguing that k is the second most important parameter after σ. Also, they test a range of values of λ and concluded that $\lambda = 0.3$ has one of the lowest errors in the sample.

I employ the same analysis and the results are reported in table 1. The number of jumps in oil futures is calculated by counting observations when the absolute return was more than n times the standard deviation of returns over the last 252 days. The results in table 1 show that there is a relatively small change in the number

Table 1. Jump intensity and size of oil futures. A change in oil price is defined as a jump when its absolute value is larger than n times the standard deviation of changes in oil price over the past 252 days. λ is the average annual number of jumps. |k| is the average annual absolute jump amplitude. Jump statistics are calculated for futures with expiry of 2, 12 and 30 months over time period from October 2009 to May 2015.

	2 months		12 months		30 months	
n	λ	$\lvert k \rvert$	λ	$\lvert k \rvert$	λ	$\lvert k \rvert$
1	64.48	2.79%	59.10	2.34%	58.39	1.88%
2	13.79	4.69%	13.79	3.80%	13.07	3.05%
3	4.30	5.94%	3.58	5.22%	4.66	3.88%
4	1.07	8.30%	1.25	6.34%	1.79	4.22%
5	0.36	9.83%	0.72	6.50%	0.72	5.43%
6	0.18	10.72%	0.36	6.66%	0.36	7.30%

of jumps per annum when n is equal to 4 or 5 for futures with two month expiry and 5 or 6 for longer dated futures. In order for jump sizes k to be comparable also across different expiration periods, I select $\lambda = 0.5$ for all calibration periods.

3.2 Pricing

Only in very few cases and for basic derivatives prices can be obtained also via closed form pricing formulas (e.g. the Black-Scholes formula for call and put options, the Margrabe formula for an exchange option). In addition, these closed form solutions are based on stringent assumptions, which when relaxed disrupt the pricing mechanism. Therefore, in practice, derivatives such as path-dependent options are priced using Monte Carlo simulations.

3.2.1 Asian option

An Asian option is a path-dependent option whose payoff is based on the average price of the underlying asset or on the average value of the strike over the option's lifetime. The former is referred to as a fixed strike Asian option, while the latter as a floating strike Asian option. Since the fixed strike Asian option is more widely used in practice and also the one implemented in the thesis, for simplicity it will be referred to as an Asian option. Also, only Asian call options will be considered. The motivation to price an Asian option steams from its practicality. In contrast to a plain-vanilla call option, the price of an Asian option is almost immune to the manipulation of the underlying as well as it is less affected by a sudden increase in market turmoil. Also, because of its increase in popularity, Asian options for some highly liquid assets as oil are quoted on the market.

The general formula for a price of an Asian call option is as follows:

$$C_0 = E^Q[(A(T) - X)^+ e^{-rT}], \tag{11}$$

where $A(T)$ is the average price of the underlying given by

$$A(T) = \frac{1}{m}\sum_{i=1}^{m} F(t_i). \tag{12}$$

m specifies the number of cut-off points over which the price of the underlying is averaged. While the size of m must be selected depending on the frequency of data as well as the time to expiry, its impact on the precision for Asian option is marginal in contrast to a barrier option where m should be infinitesimal.

The description of an Asian option so far refers to a setting where the asset price is averaged over the period from valuation date till option's expiry. This is a standard setting and it will be used when performing sensitivity analysis. However, in order to make a comparison between estimated Asian option prices and the market quoted prices, a small adjustment in pricing formula must be made. Since the prices for Asian options quoted on NYMEX are calculated by averaging asset prices only over the last month, the asset prices will be simulated on a daily basis, but equation (12) will be applied only over the last month till option's expiry (CME Group, 2016b).

3.2.2 Monte Carlo simulation

The most frequent tool used in pricing derivatives is a Monte Carlo simulation (Glasserman, 2004). Intuitively, the method randomly samples from all the possible outcomes and the law of large numbers ensures that we obtain the expected value. In practice, this implies simulating a random variable that acts as a shock to the evolution of the price of the underlying.

Since closed form solutions are derived from risk-neutral dynamics by taking the expected value from the imposed distribution, then any such solution can be achieved by using a Monte Carlo simulation on these risk-neutral dynamics. For the Merton's jump-diffusion model the risk-neutral dynamics are reported in equation (4). However, taking into account that Asian options are path-dependent, equation (4) must be discretized by dividing time into many small steps as follows:

$$F(t_{i+1}) = F(t_i)e^{\left(r-q-k\lambda-\frac{1}{2}\sigma^2\right)dt+\sigma\sqrt{dt}Z}\prod_{j=N(t_i)+1}^{N(t_{i+1})} Y_j\,, \tag{13}$$

where dt is a time step defined as $dt = \frac{T}{m}$ and Z is a standard normal variable, referred to as a shock.

Combining equation (13) with (11) leads us to the price of an Asian call option for one simulation:

$$C_s = \left(\frac{1}{m}\sum_{i=1}^{m} F(t_i)e^{\left(r-q-k\lambda-\frac{1}{2}\sigma^2\right)dt+\sigma\sqrt{dt}Z}\prod_{j=N(t_i)+1}^{N(t_{i+1})} Y_j - X\right)^{+} e^{-rT}\,, \tag{14}$$

The final price for an Asian call option in the Merton jump-diffusion model using Monte Carlo simulations is given by

$$C_0 = \frac{1}{sims}\sum_{s=1}^{sims} C_s\,. \tag{15}$$

4 Data description

In the following section I discuss the data employed in the analysis with a focus on two main datasets that are used for calibration and for comparing estimated prices with market ones.

4.1 Calibration dataset

As for the absolute majority of commodities traded on an exchange, oil is quoted only in terms of futures contracts. Since futures contracts are standardized, in case of oil futures this implies that there are only twelve expiration dates available in each given year. This means that exactly one month before expiration of the closest futures contract, it is not possible to buy a ten day or a fifty day futures contract on an exchange while it is possible to buy a contract with expiry in a month, two months or even in sixty months.

For each of the oil futures contract available in the market, there is an imperfect set of options written on it, whose time to expiry coincide with the expiry of the underlying futures contract. Therefore, the first main dataset consists of all quoted call option prices over all available strikes written on oil futures from October 2009 to May 2015. Since in the analysis I will examine and compare options with 2-month, 12-month and 30-month expiration periods, the corresponding expirations are from January 2010 to November 2017.

The data is taken from NYMEX and compromises all prices available for oil call options on Bloomberg Terminal and Thomson Reuters DataStream. The limited historical price availability is the reason why a particular date is selected as the starting point.

4.2 Pricing dataset

Asian options have witnessed their growth in popularity relatively recently. As a result, prices on NYMEX are available only for expiration dates starting from February 2012. Similar to call options on oil, Asian options also expire only twelve times in a year. In addition, for a given underlying futures contract the number of Asian options available in the market is on average for only twenty strikes. Therefore, the second main dataset consists of all quoted Asian call option prices over all available strikes written on oil futures from December 2011 to May 2015.

4.3 Adjustments

Similar to Kearney *et al*. (2015), I perform few adjustments to all option prices to ensure that results are not severely affected by illiquidity. First of all, since a plain-vanilla call option price and an Asian call option price decreases as strike price increases, an option is eliminated if its quoted price is larger than for an option with a lower strike. Secondly, an option must have a minimum quoted price of 0.05 dollars. Both of these adjustments only affect deep OTM call prices. Thirdly, options, for which moneyness, defined as strike over the price of futures contract, is lower than 0.5 or larger than 2.0, are also eliminated. All these adjustments are carried out to ensure that illiquidity, which is most pronounced at extreme moneyness levels, does not disrupt calibration and pricing results.

5 Analysis and discussion

In this section I present and discuss the results of the analysis. I start by providing a brief overview of oil returns. Afterwards I examine whether information obtained from calibration corresponds to what happened in the market at that point in time and whether this information can be used to predict the future. Next, I look at how prices obtained using retrieved parameters compare to market prices for Asian options. Finally, I present sensitivity analysis for a hypothetical Asian option.

5.1 Oil returns

For many models, including the Black-Scholes model, the underpinning assumption is that returns are normally distributed. Figure 1 presents historical probability density for oil returns from October 2009 to May 2015 and probability density as implied by the normal distribution. As it can be seen in figure 1, the historical oil return distribution is more peaked with fatter tails as compared to the normal distribution; the historical oil return distribution is leptokurtic. While skewness for the historical oil return distribution is very close to the one implied by the normal distribution (0), kurtosis is much larger as implied by the normal distribution (3). Not surprisingly, Jarque-Bera test for normality very strongly rejects that the historical oil returns are normally distributed. The results obtained here are nothing unusual as most asset returns are characterised by positive excess kurtosis. In addition, it is typical to see a

return distribution that is highly negatively skewed (see, for instance, Askari and Kirchene, 2008; Christoffersen *et al*., 2012).

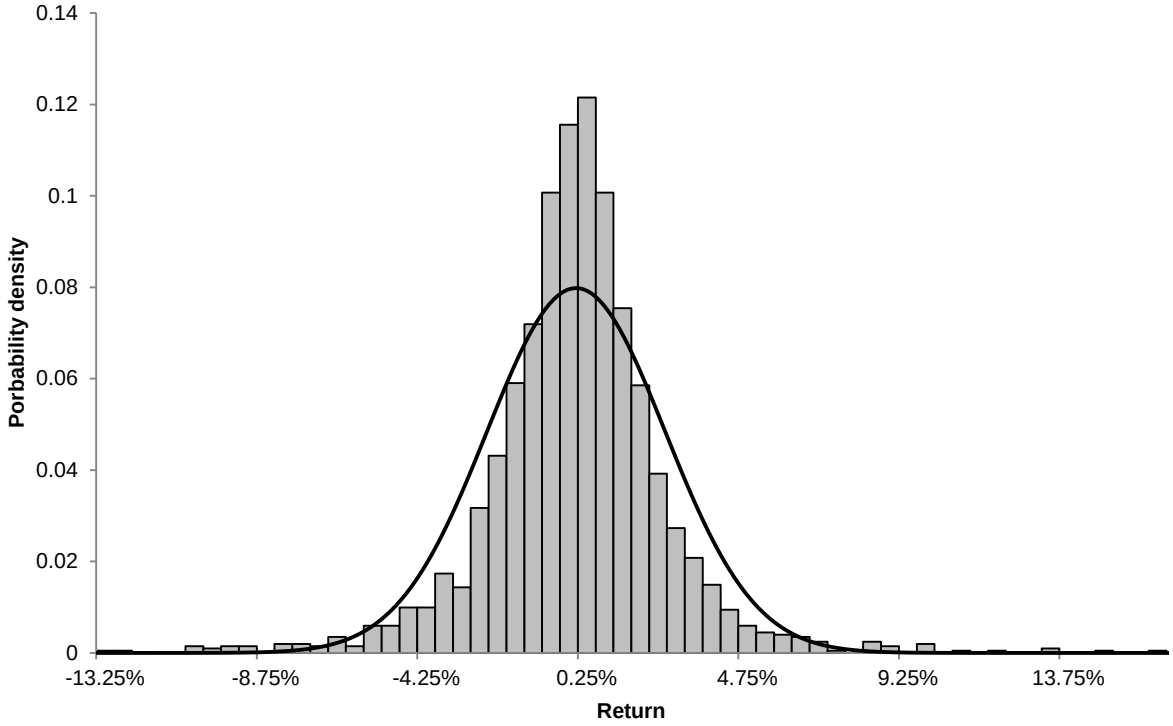


Figure 1. Probability density for historical oil returns. Historical oil returns are calculated from oil futures with a maximum of 1 month till expiry from October 2009 to May 2015. Bars represent probability density for the historical oil returns while solid represents probability density as implied by the normal distribution. Descriptive statistics for the historical oil returns: mean = -0.05%, standard deviation = 2.5%, skewness = 0.12, kurtosis = 8.12, Jarque-Bera normality statistic = 2202, p-value = 0.

5.2 Implied information

After investors make trades, their views are reflected into prices. For example, currently there is a heated discussion whether Britain will leave the European Union, commonly referred to as Brexit. Popular opinion is that if Britain leaves, its economy will be affected negatively. If an investor believes that Brexit will happen, she should, for example, sell British Pound in favour of Euro. This trade will have a direct impact on the currency pair by pushing British Pound lower. Looking at this trade inversely, from changes in the price one can extract the aggregate view of the market and use this information to forecast and to price assets. Essentially, this is calibration as market data must be fitted to a model.

5.2.1 Calibration and economic significance

In order to evaluate whether the retrieved parameters correspond to what happened in the market, I first discuss the main events that occurred in the oil market from October 2009 to May 2015. The events with the most notable effect on oil price are presented in table 2.

Table 2. The main events in the oil market. Brief and indicative list of few of the main main events that occurred in the oil market from October 2009 to May 2015. The impact on change in oil price is indicitaed by arrows.

Date	Event	Impact
Beginning of 2011	The Arab Spring	↑
Middle of 2011	Concern over next recession	↓
Beginning of 2012	Sanctions on Iran	↑
Middle of 2012	EU debt crisis	↓
Middle of 2013	Political crisis in Egypt	↑
End of 2014	Oversupply	↓
Beginning of 2015	New equilibrium	-

The period was mostly driven by political crises in the Middle East countries and economic crises in the European Union countries. Figure 2 shows the development of oil futures prices with maturity of 2 months, 12 months and 30 months. Apart from political and economic crises, another phenomenon emerged. Starting from the middle of 2013 till the middle of 2014, oil futures with shorter maturities traded higher than those with longer maturities. The reason for this observance is due to convenience yield, which is negatively related to inventory levels (Glasserman, 2004). During this period inventory levels were low. In addition, political crisis in Egypt threatened to disrupt oil supply even more, resulting in even lower inventory levels. This resulted in a large convince yield that translated into lower oil futures prices for longer maturities through spot-future parity. The occurrence in the beginning of 2015 is opposite as inventory levels were high and consequently the yield was low.

The retrieved parameters coincide with what happened in the oil market. In figure 3 the average jump process amplitude k, as calibrated from options with expiry of 2 months, 12 months and 30 months, is depicted over time. The first observation is that k_2 fluctuates more than k_{12} and k_{30}. This is as expected because events usually fade away in a longer time period because political and economic crises get resolved. The second observation is a spike in k_2 in March 2011 when the Arab Spring took place.

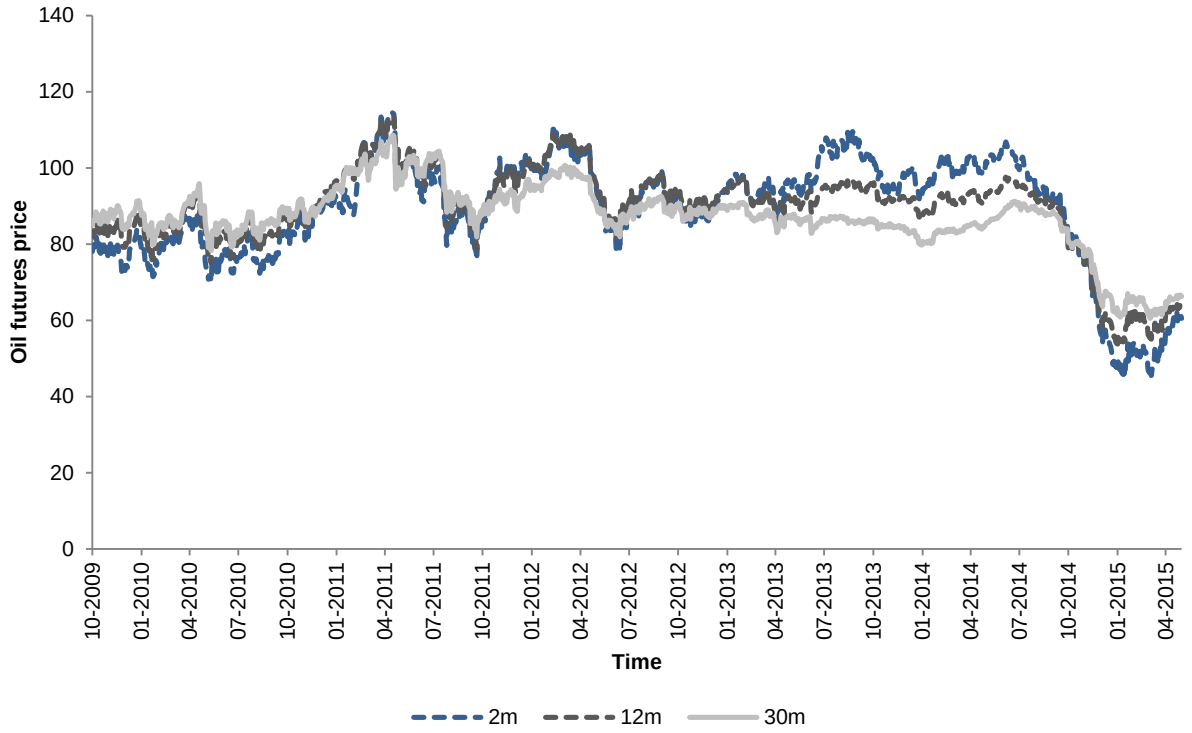


Figure 2. Oil futures price dynamics. Generic oil futures prices with 2, 12 and 30 months to expiry from October 2009 to May 2015.

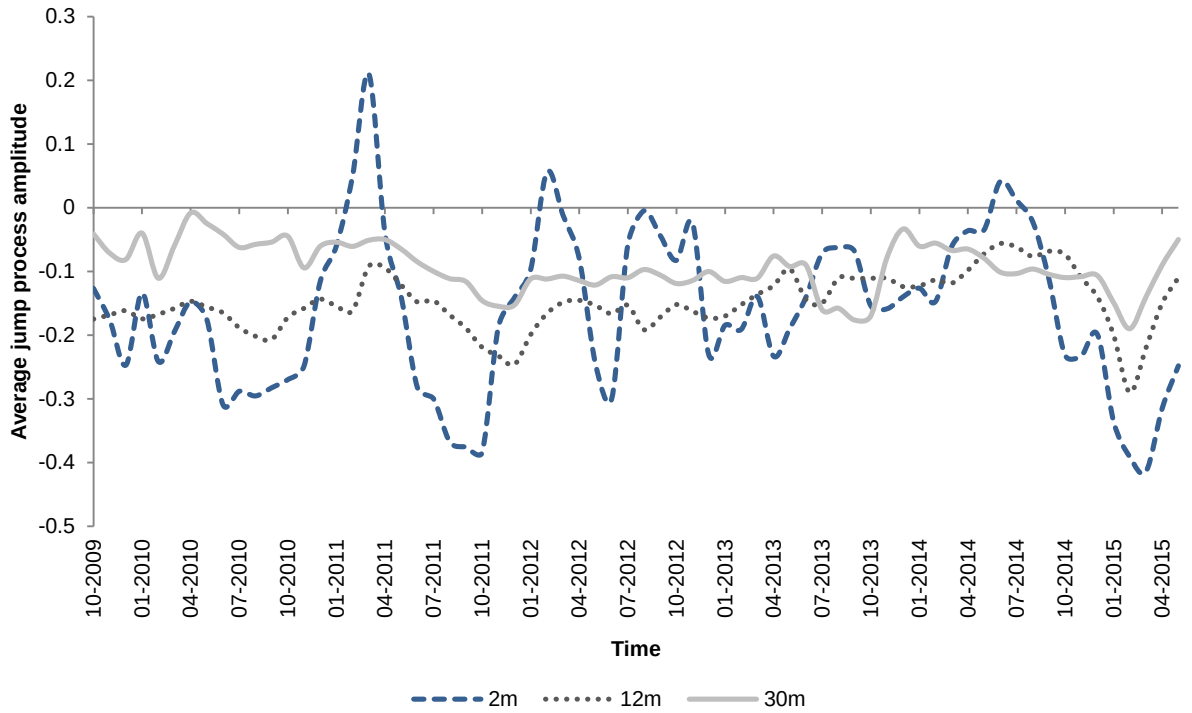


Figure 3. Average jump process amplitude dynamics. The dynamics of the average jump process amplitude as calibrated from options available in the market with constant maturity of 2, 12 and 30 months from October 2009 to May 2015.

Since Arab countries are the lead exporters of oil and the political crisis resulted in decreased exports from Libya with more possible disruptions, the average jump process amplitude was positive as suggested by theory. The third observation is a large drop in k_2 and a smaller drop in k_{12} and k_{30} in the end of 2014. This period was characterised by a decreasing demand and an increasing supply. Since oil is an inelastic asset, a large drop in price should happen for the market to clear and hence the significant decrease in the average jump process amplitude is observed. Also, by comparing figure 2 with figure 3 it is clearly evident that the average jump process amplitude fluctuates in line with oil futures prices. This is particularly evident when looking at the main events from table 2.

The retrieved values of the diffusion volatility σ also very well describe what has happened in the market over time. It is especially evident when comparing with oil returns (figure 4). As can be seen in figure 5, σ_2 is always larger than σ_{12} and σ_{12} is always larger than σ_{30}. This is in line with expectations because events usually have short term impact and, while the retrieved diffusion volatility from longer maturities is also affected, the impact is evidently smaller. An interesting observation is that while the average jump process amplitude significantly increased in the beginning of 2011, there was no increase in the diffusion volatility. However, there was a large increase in the diffusion volatility in middle of 2011 when it became clearer that Greece will not default on its debt. The fact that the average jump process amplitude conveys different information than the diffusion volatility is important for both pricing and forecasting. When comparing figure 4 with figure 5, it is clearly evident that the diffusion volatility capture changes in oil futures returns.

The standard deviation of jump process amplitude δ also tracks market events. Development of δ is shown in figure 6 and it resembles more how the diffusion volatility has fluctuated over time than the average jump process amplitude. Two important findings can be drawn when comparing δ and σ fluctuations over time. First of all, δ fluctuates significantly more over time. Secondly, δ also for longer maturities capture short-lived events in the market while σ for longer maturities remain rather unchanged. These findings show the importance of modelling oil returns using a jump-diffusion model, such as the MJD, instead of using a pure diffusion model as the Black-Scholes or even a stochastic volatility model without jumps. For instance, consider the time period around the middle of 2012 when there was a sharp

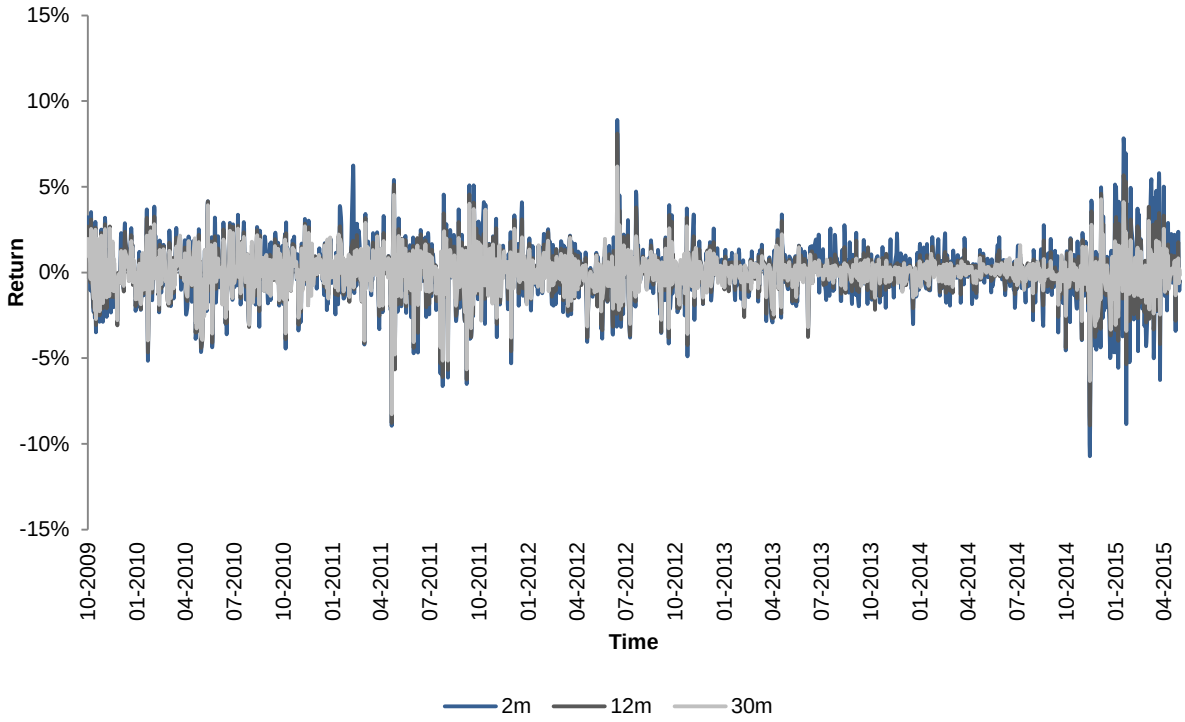


Figure 4. Oil futures return dynamics. Generic oil futures returns with 2, 12 and 30 months to expiry from October 2009 to May 2015.

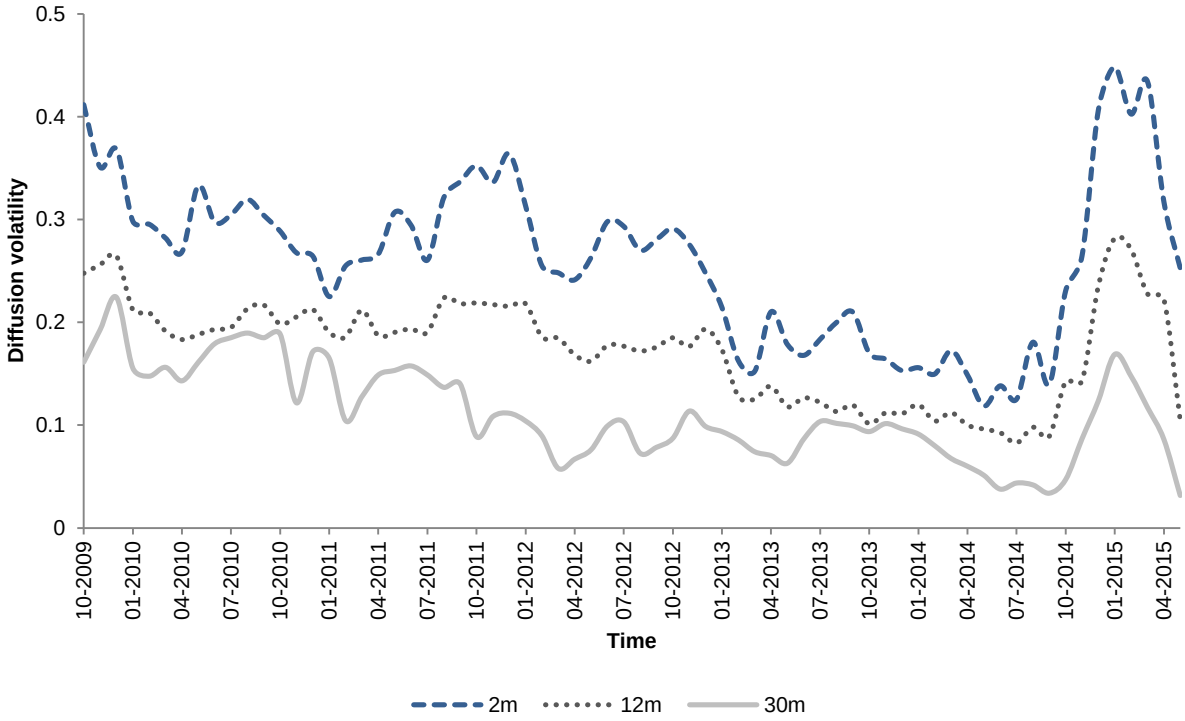


Figure 5. Diffusion volatility dynamics. The dynamics of the difusion volatility as calibrated from options available in the market with a constant maturity of 2, 12 and 30 months from October 2009 to May 2015.

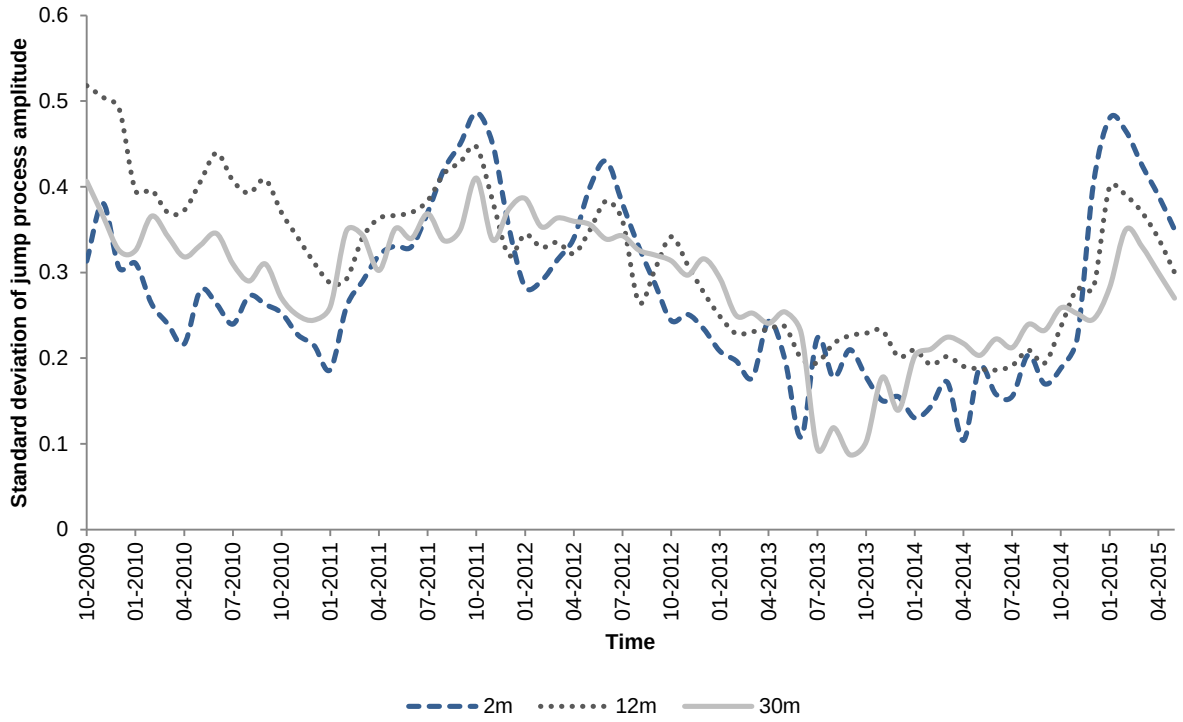


Figure 6. Standard deviation of jump process amplitude dynamics. The dynamics of the standard deviation of jump process amplitude as calibrated from options available in the market with a constant maturity of 2, 12 and 30 months from October 2009 to May 2015.

decrease in oil price followed by a fast recovery. Since this event was not fundamental and could be characterised as a jump, in theory it should be captured by jump components, leaving the diffusion volatility virtually unchanged. The retrieved values from options with two months of expiry were as follows: a large drop in the average jump amplitude followed by a sharp increase, a minor increase in the diffusion volatility and a large jump in the standard deviation of jump process amplitude. However, if a model without jumps would be employed instead, then unusually large volatility must be used to capture this fluctuation in oil price. Since option price increases with volatility, the price of an option would be relatively large due to increased volatility. However, fluctuations in returns afterwards decayed making a large increase in volatility unjustified. Hence, if models without jumps are employed, option prices would be notably mispriced. As such, while small changes in the diffusion volatility are reasonable, large changes in the diffusion volatility should be driven by fundamental events as oversupply of oil that started in the end of 2014. As a result, the first hypothesis that *a jump component must be included when modelling returns in the oil market* cannot be rejected. This finding is in line with the

empirical literature (see, for instance, Bates, 1996; Broadie *et al*., 2007; Chernov and Ghysels, 2000; Das and Foresi, 1996) that also acknowledges the importance of adding jumps in modelling asset prices.

All in all, the average jump process amplitude k, the diffusion volatility σ and the standard deviation of jump process amplitude δ very well explain what happened in the market at that point in time. Hence, the second hypothesis that *retrieved parameters explain market events over time* cannot be rejected. This is in line with results obtained by Murphy and Ronn (2015). In a similar fashion they calibrate the MJD model to oil prices over the period from 2007 to 2012. Since my analysis starts from the end of 2009, direct comparison can be made only for two years. In general, the results are very similar regarding the direction of movement in the retrieved parameters. The biggest difference is that in their work the retrieved parameters fluctuate even more. For example, in the end of 2011 δ_2 even reached value as high as 1.5 as compared to 0.5 obtained in my analysis. There are several reasons for this observation. First of all, they calibrate δ_2 from options that have maturity between one month and two months while I keep maturity always constant, which in case of δ_2 is two months. Secondly, they apply equal weight to all options while I weight each option so that it has an equal contribution when calibrating. Implicitly, I put more weight on options away from ATM strike as compared to Murphy and Ronn (2015). Since there is more variability in options with short maturity and in more liquid options that are around ATM strike, the fact that their retrieved values show more variability is not surprising. However, in rather calm periods the retrieved values are very comparable.

5.2.2 Forecasting

The analysis so far has focused on whether retrieved values explain what happened in the oil market at that time. However, most practitioners are more interested in whether retrieved values can tell something about the future as well. In order to test this, I adopt a simplified methodology from Kearney *et al*. (2015). To implement this, I first need to obtain ITM, ATM and OTM volatilities and then I can examine whether changes in the ends of the implied volatility curve can provide valuable predictive information.

I split the dataset for options with two months to expiry in three parts: ITM options, ATM options and OTM options. A two month period is chosen because of

liquidity reasons. Since I work only with call options, the ITM option dataset consists of calls with moneyness below 0.8, the ATM option dataset consists of calls with moneyness between 0.9 and 1.1, and the OTM option dataset consists of calls with moneyness above 1.2. Afterwards, to obtain the retrieved parameters I calibrate each dataset separately. The results for the diffusion volatility are shown in figure 7. The results are similar to the ones that are obtained by calibrating to a full dataset; however, few new important findings emerge. ITM diffusion volatility is all the time

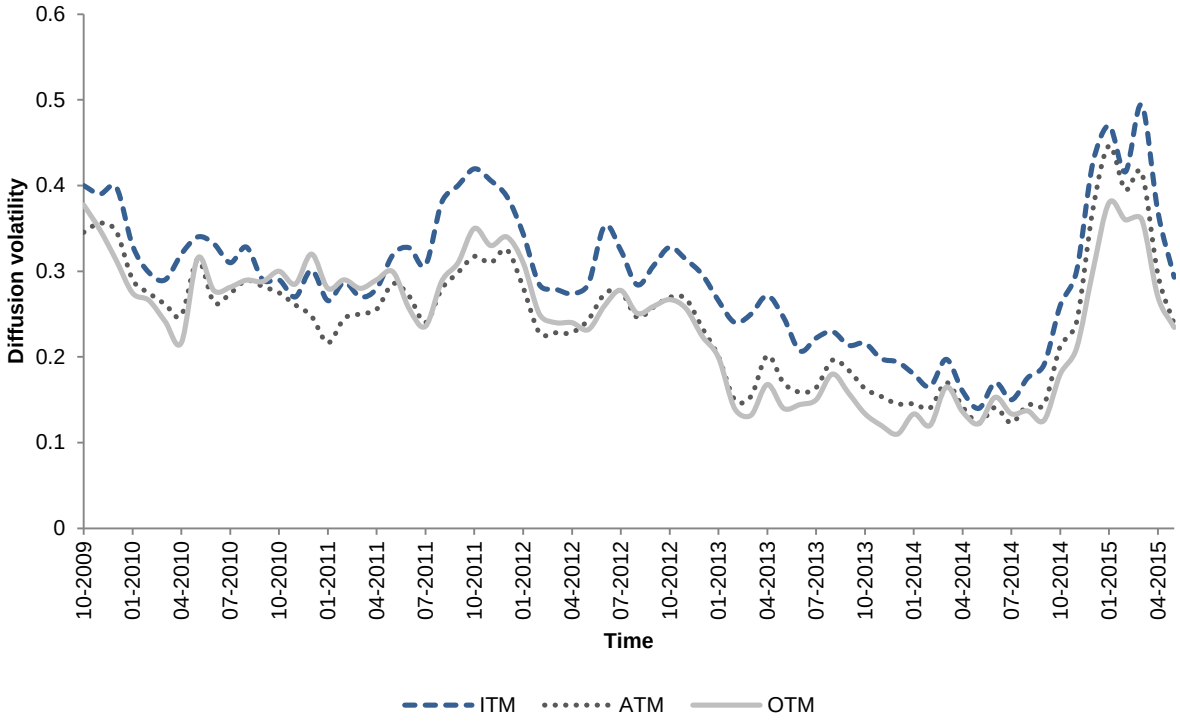


Figure 7. Diffusion volatility dynamics for ITM, ATM and OTM options. The dynamics of the difusion volatility as calibrated separately from ITM, ATM and OTM call options available in the market with a constant maturity of 2 months from October 2009 to May 2015.

above ATM diffusion volatility while OTM diffusion volatility is more frequently below ATM diffusion volatility than above. This shows that the implied volatility curve is more frequently characterised as a skew rather than a smile. It is important to note that in practice the implied volatility refers to a volatility that is extracted using the Black-Scholes formula; however, as results show, also the implied volatility obtained from the MJD model takes distinct shapes that are in line with events in the oil market. For example, consider the period from the end of 2010 till the middle of 2011 when oil prices were surging. During this period OTM diffusion volatility overtook ITM

diffusion volatility as market participant bought OTM calls to protect themselves from an even larger increase in oil price. This is in contrast to what happened in the end of 2014 when oil prices were sinking. In order to hedge from an even larger drop in oil price, market participants bought OTM puts and, as a result, ITM call diffusion volatility was significantly higher than that of ATM and ITM.

In order to test whether changes in ITM and OTM diffusion volatility have a predictive power, I employ a simply strategy. This strategy consists of two steps. First, an oil futures contract is sold whenever in the previous lag there was a positive change in ITM diffusion volatility and this change was larger than that of ATM and OTM diffusion volatilities. Second, an oil futures contract is bought whenever in the previous lag there was a positive change in OTM diffusion volatility and this change was larger than that of ITM and ATM diffusion volatilities. Since there might be a delayed reaction, I look at a month and a two months lag. The results of this strategy are presented in table 3.

Table 3. Performance of the strategy. Average return and total return are reported over the number of trades (months) when the strategy resulted in gain or loss based on information one month ago (t-1) or two months ago (t-2). Also, the net result is shown. The strategy is employed over the period from October 2009 to May 2015.

	Gain			Loss			Net		
Lag	Average	Total	Trades	Average	Total	Trades	Average	Total	Trades
t-1	10.35%	155.23%	15	-8.92%	-107.06%	12	1.78%	48.17%	27
t-2	13.09%	170.12%	13	-5.86%	-82.03%	14	3.26%	88.09%	27

As it can be seen, this strategy results in a positive net return. However, since this return is not accounted for transaction costs and in total 27 trades are performed, the overall net result would be margin, if positive at all. While the results provide some predictability based on a very simplistic model, they are of very little use to practitioners. Hence, the third hypothesis that *retrieved parameters do not contain predictive power* cannot be rejected.

The conclusion obtained here is very similar to that of McIntyre and Jackson (2009), who looked at implied returns distribution obtained from S&P 500 index options and found that that there is virtually no predictability in the market from 1999 to 2002. However, the conclusion obtained in my analysis is in contrast with Bates (1991), who analyse the October 1987 crash, and Rappoport and White (1994), who look at the crash of 1927. Both studies find predictive evidence. However, it is

important to note that these crashes were very extreme events. If I apply my strategy only on the event when oil prices sunk in the end of 2014 due to oversupply, the strategy provides very strong positive performance; however, it does not predict the drop beforehand.

5.3 Pricing Asian options

Once retrieved values are obtained, the usual string in the empirical literature is to use them to create a hedging model, which is compared with a benchmark hedging model (see, for instance, Andersen and Andreasen, 2000; Kearney *et al*., 2015). In contrast, I use the retrieved parameters to price Asian options. For each Asian option available in the market, I estimate its price using the retrieved parameters and then calculate its absolute relative pricing error. The results for the average values are shown in figure 8. There are two important observations. First of all, the error for options with shorter maturity fluctuates more than for options with longer maturities. This could be expected because events usually have a short-term effect and hence they have rather marginal impact in the longer term. This is particularly evident in the

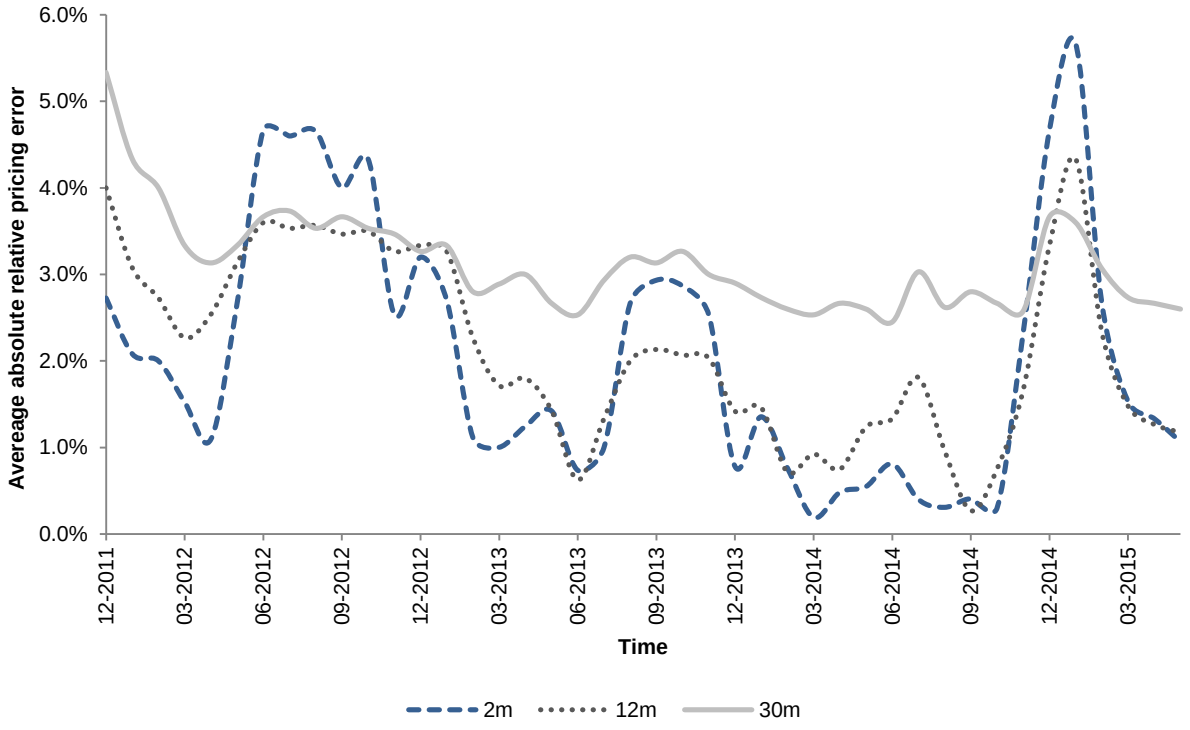


Figure 8. Average absolute relative pricing error dynamics. For a single option absolute relative pricing error is defined as the absolute difference between its market price and its estimated price over its market price. All available Asian options are considered and their estimated prices are calculated based on the retrieved parameters from all available options. Options with a constant maturity of 2, 12 and 30 months from December 2011 to May 2015 are shown.

middle of 2012 and 2013, and in the end of 2014 when events resulted in a substantially increased error for options with a shorter maturity that rapidly decayed after these events. Secondly, in absence of events the error for options with a shorter maturity is notably smaller. A reasonable explanation is liquidity that affects both estimated prices and market prices. There are substantially more options with a maturity of 2 months than with a maturity of 30 months and there is also smaller trading activity with longer maturity options. As a result, prices quoted in the market for options with longer maturity might be imprecise, as well as the precision of estimated prices could be negatively affected because of biased retrieved parameters.

The overall results in figure 8 indicate that the absolute relative pricing error, while not excessive even during periods of market turmoil, is still considerable. Hence, to evaluate whether the error is within an acceptable range, I estimate prices for three options and their confidence intervals, and compare whether market prices lie within these confidence intervals. I do this only for options with a maturity of two months because of the liquidity reasons mentioned in the previous paragraph. The

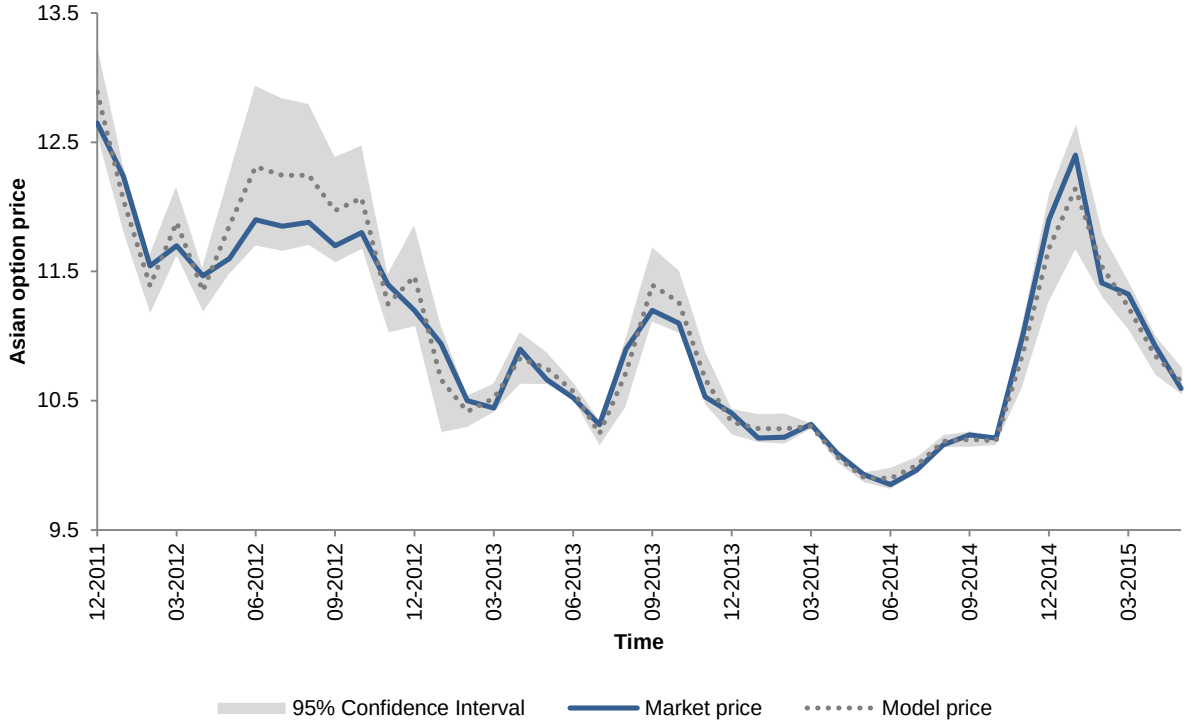


Figure 9. ITM Asian option price dynamics. Dynamics of ITM Asian option market price (solid line) and its estimated price based on the retrieved parameters from all available options (dotted line) from December 2011 to May 2015. The grey area represents 95% confidence interval around the estimated price.

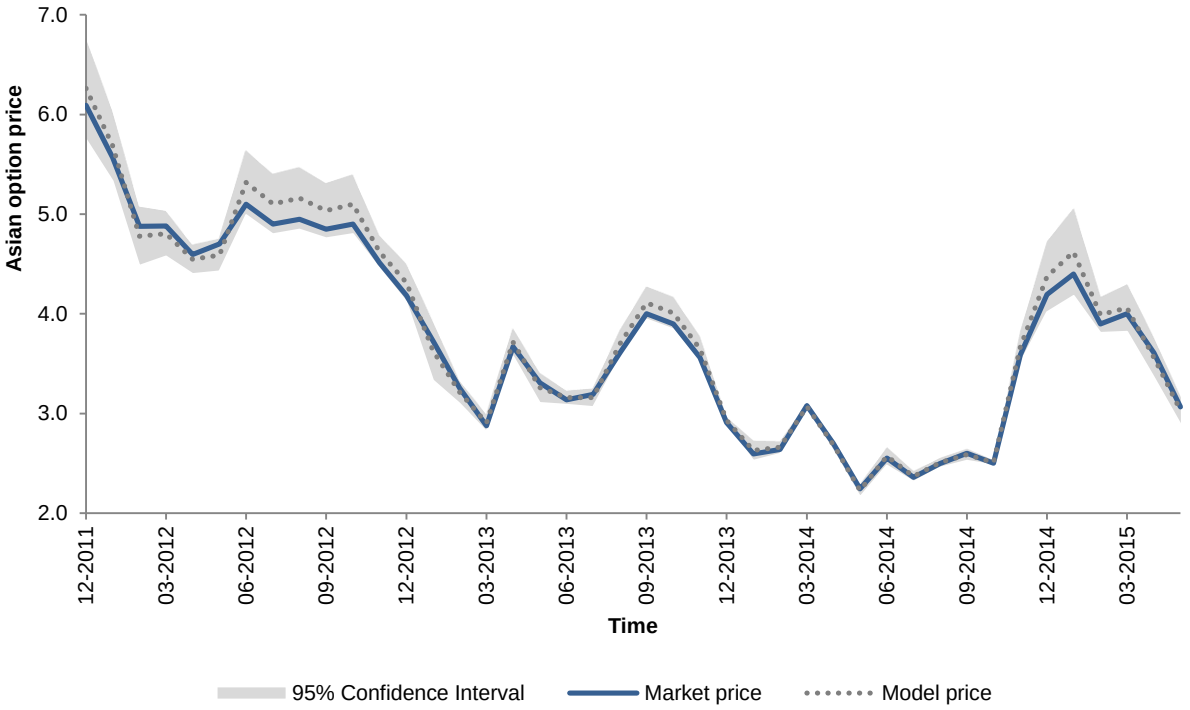


Figure 10. ATM Asian option price dynamics. Dynamics of ATM Asian option market price (solid line) and its estimated price based on the retrieved parameters from all available options (dotted line) from December 2011 to May 2015. The grey area represents 95% confidence interval around the estimated price.

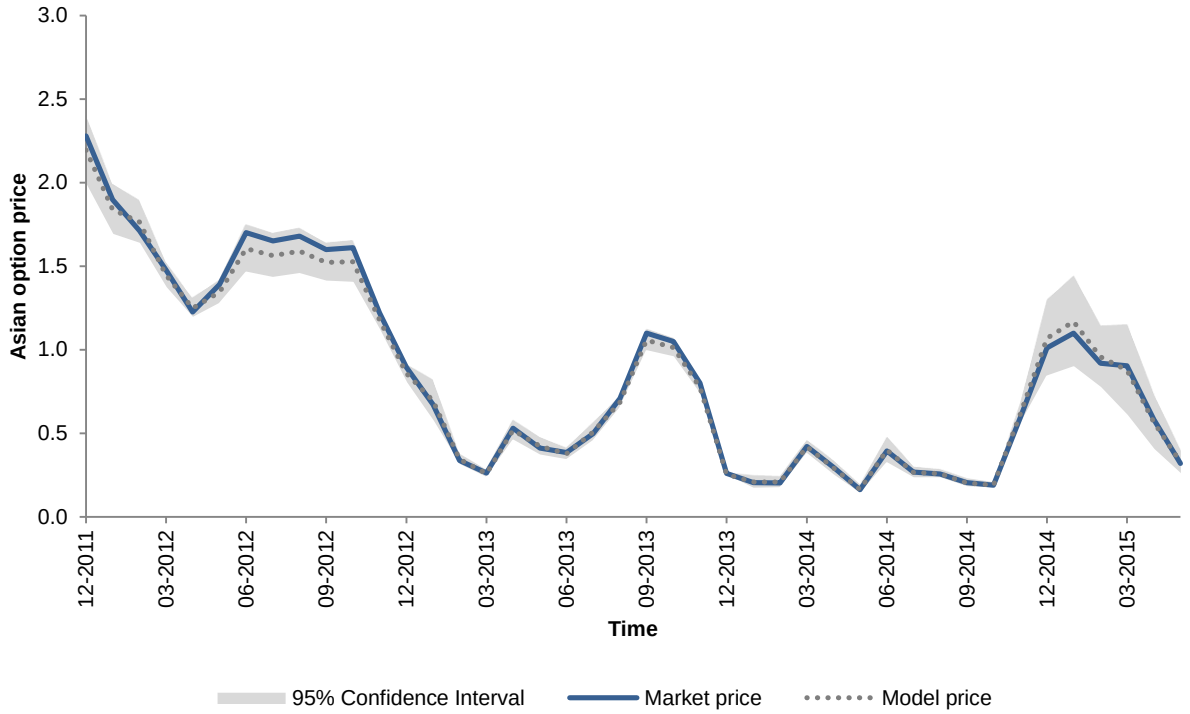


Figure 11. OTM Asian option price dynamics. Dynamics of OTM Asian option market price (solid line) and its estimated price based on the retrieved parameters from all available options (dotted line) from December 2011 to May 2015. The grey area represents 95% confidence interval around the estimated price.

results in the case of an ITM option, whose strike is 10 USD lower than the ATM strike, are shown in figure 9. The results for an ATM option are reported in figure 10, and the results for an OTM option, whose strike is 10 USD larger than the ATM strike, are depicted in figure 11. In all of the three cases market prices are always within the confidence interval. Hence, the fourth hypothesis that *the MJD model prices Asian options correctly* cannot be rejected.

5.3.1 Separate calibration and pricing

The analysis so far has been based on prices that are estimated from retrieved parameters based on all the available options. That is, the estimated price in figure 9, which is for an ITM option, is calculated based on retrieved parameters that were obtained from a full spectrum of options available at that time, not only ITM options. In this subparagraph, I compare the performance of ITM, ATM and OTM options, whose prices are estimated only from parameters obtained from ITM, ATM and OTM options, respectively. To better grasp the idea consider figure 12. The solid blue line corresponds to the absolute relative pricing error only between ITM market prices and their estimated prices based on parameters retrieved only from calibrating ITM options. In contrast, the dotted grey line compares the error between only ITM market prices and their estimated prices based on parameters retrieved from calibrating all options. The results in figures 13 and 14 are obtained in the same manner only for ATM and OTM options.

There are several valuable findings. First of all, in all cases the error is almost always lower when prices are estimated using parameters retrieved from a set of options. Secondly, increased levels of error correspond to periods of surge in the volatility of oil returns. Thirdly, consider only the error when all options are used. In the case of ITM options in figure 12, the level of error in the end of 2014 is lower than the error during the event that started in the middle of 2012. The opposite is true for OTM options in figure 14. A reasonable explanation is as follows. The event in the middle of 2012 resulted in a spike in oil price while the event in the end of 2014 resulted in a drop in oil price. Since during price spikes market participants buy OTM calls to protect themselves from an even larger increase in oil price, both market prices and prices used in calibration contain more OTM information. Hence, in the case of OTM options, the error based on all options during a price spike should be of similar levels than that based only on OTM options. The results show that this is the

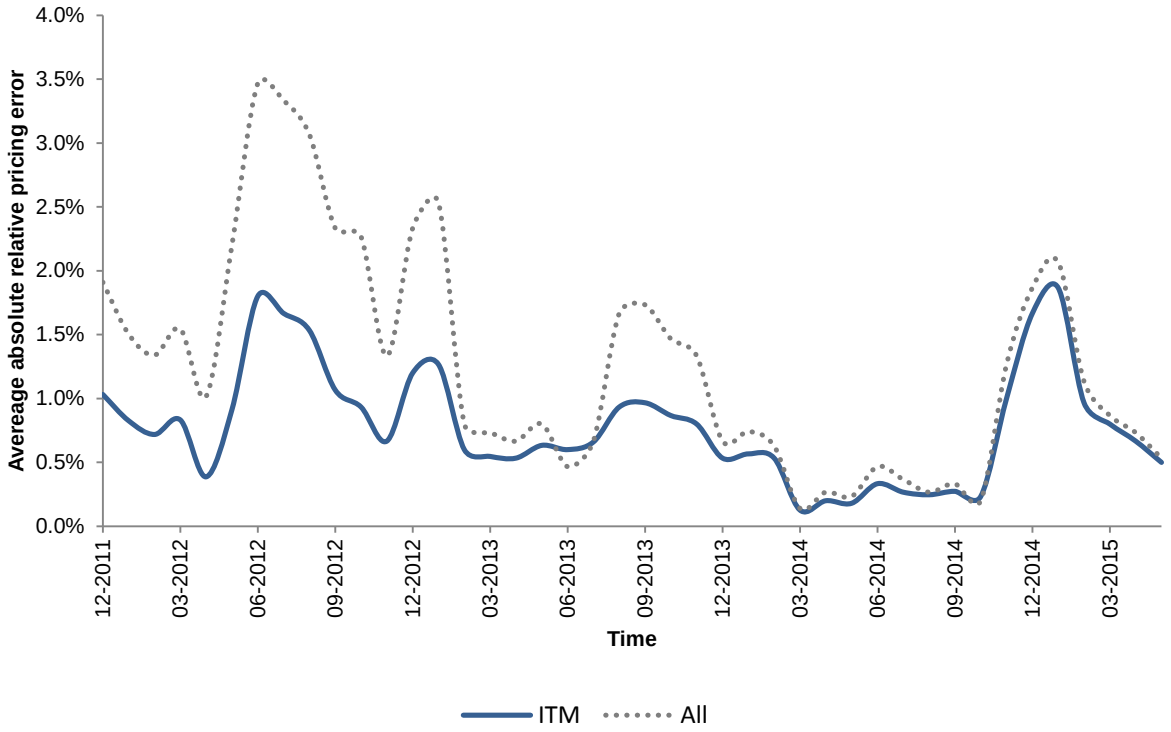


Figure 12. Average absolute relative pricing error dynamics for ITM options. Only ITM Asian options are considered. Their estimated prices are calculated based on the retrieved parameters from all available options (dotted line) and the retrieved parameters from only ITM options (solid line). Dynamics are shown from December 2011 to May 2015.

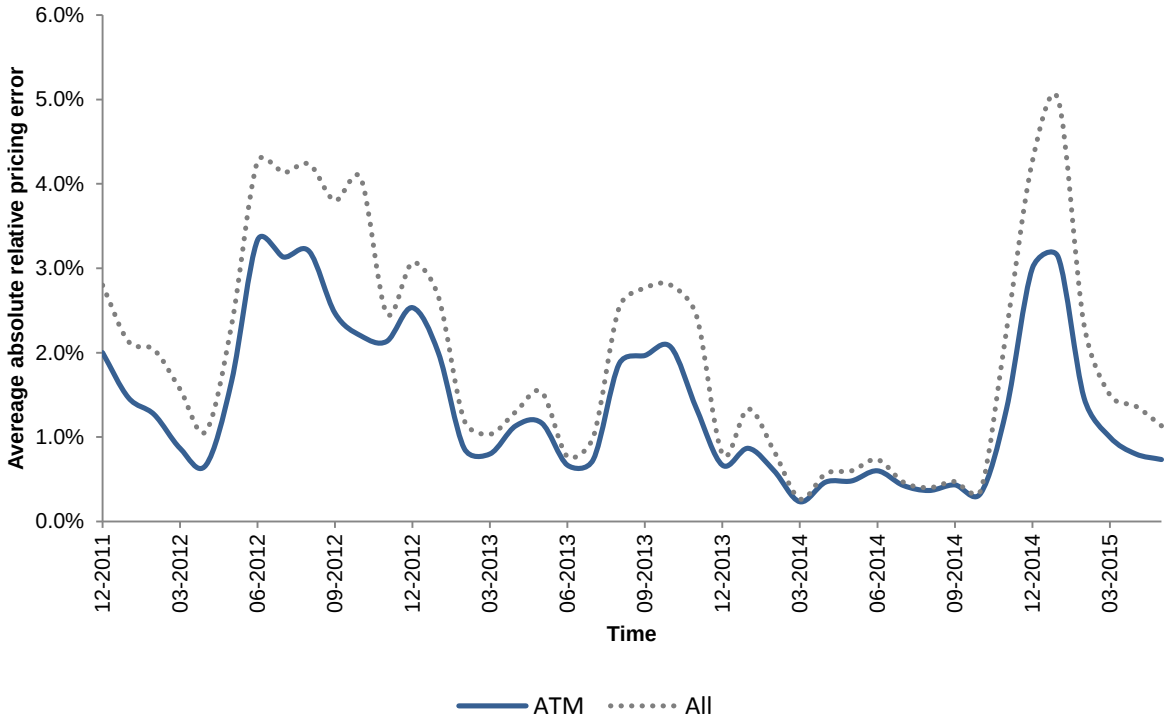


Figure 13. Average absolute relative pricing error dynamics for ATM options. Only ATM Asian options are considered. Their estimated prices are calculated based on the retrieved parameters from all available options (dotted line) and the retrieved parameters from only ATM options (solid line). Dynamics are shown from December 2011 to May 2015.

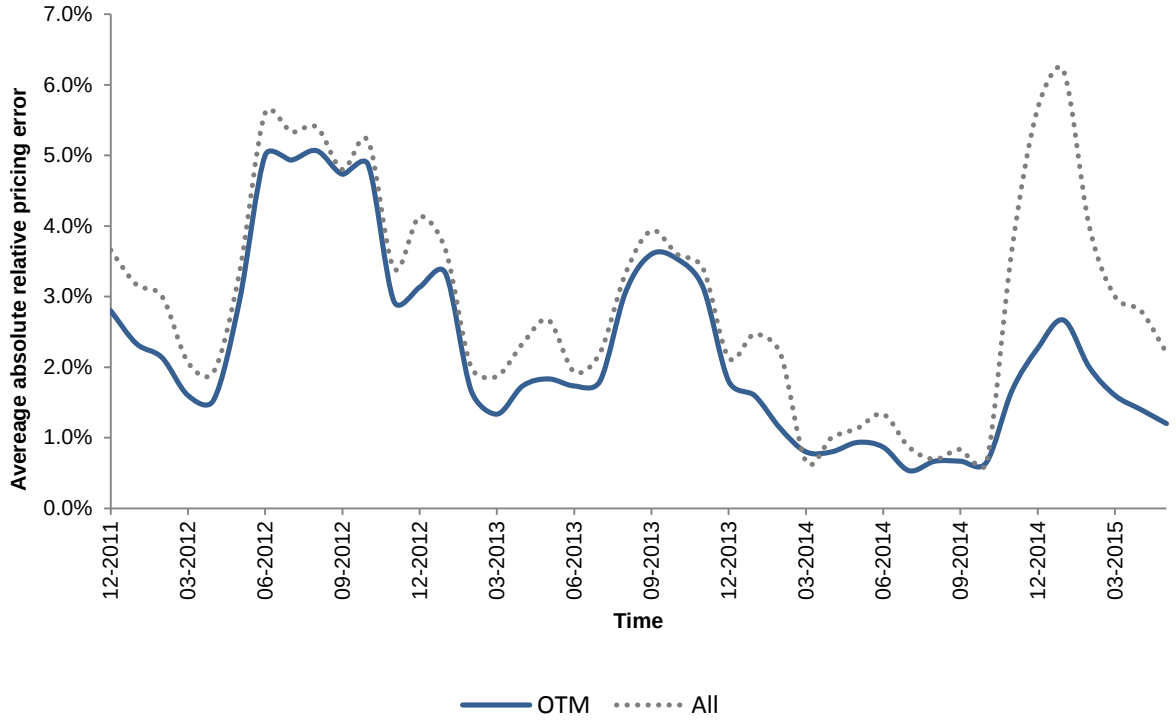


Figure 14. Average absolute relative pricing error dynamics for OTM options. Only OTM Asian options are considered. Their estimated prices are calculated based on the retrieved parameters from all available options (dotted line) and the retrieved parameters from only OTM options (solid line). Dynamics are shown from December 2011 to May 2015.

case as the error based only on OTM options, which is the blue solid line in figure 14, is very close to the error based on all options. However, in the case of ITM options, the error based on only ITM options is considerably lower because there is different, more valid information for pricing ITM options. The results are opposite when considering a sink in the oil price in the end of 2014. During the drop, there is more ITM information and hence there is substantial improvement in OTM error when using only OTM information rather than all information.

These results suggest that because of information bias, options should not be priced using parameters retrieved from all options in the market, but rather from its characteristic subset.

5.4 Sensitivity analysis

Sensitivity analysis shows how the price of a derivative varies when a particular parameter is changed, keeping other parameters constant. Hence, sensitivity analysis helps to identify which parameters must be more closely examined before

they are put in a model. The analysis here is carried out on a hypothetical Asian option, which is characterised by the following initial parameters: $\sigma = 0.2,\ \lambda = 0.5,\ \delta = 0.3,\ \mu = -0.25, F_0 = 50,\ r = q = 0.01,\ t = 1.$ These parameters are chosen because they very closely represent situation in the market in May 2015. In addition, the hypothetical Asian option is priced by taking the average over the whole path, not only the last month.

Prior to looking at how changes in the key variables, which are retrieved from the market, influence Asian option price, the relationship between time to maturity and strike price is examined. The results are shown in figure 15 and 16. In figure 15 the typical sensitivity surface is depicted, while in figure 16 the relative surface is shown. Relative surface depicts changes in price against a price that is obtained by the lowest level of parameter for a certain strike, in this case time to maturity of zero.

The results show that for deep OTM options, price increases with time to maturity while for deep ITM options the opposite is true. This is because Asian options are averaged over time and due to exponential compounding there is an implicitly higher probability for an Asian option to gain value rather than to lose it. Hence, in this particular setting, for deep ITM options there is almost a non-existent probability to result in out-of-the money. Therefore, the option price when time to maturity is five years is almost equal to the option price when time to maturity is equal to zero, adjusted for time value of money. However, in case of deep OTM options there is a larger probability that the option will end up in the money when there is more time. This phenomenon is most pronounced for an ATM option as the relative surface graph shows.

Similar surface plots are also constructed for each of the key parameters: σ, μ, k, δ, and λ. It is important to note that k is a function of μ and δ. Hence, the sensitivity of option price to k can be measured by changing μ or δ, or both. However, since the rationale behind k is closer to μ, only μ will be changed to measure sensitivity. The results are reported in the appendix. The results for σ and λ are very close to those for t with a similar explanation. The major difference is that in case of σ and λ relative surface is always positive as increase in the parameters increase the range of possible Asian option prices, but since the downside is limited, the net effect is positive. Similarly, the relative surface is also positive for δ. The particular shape is because δ influences compounding exponentially. More interesting results are for μ

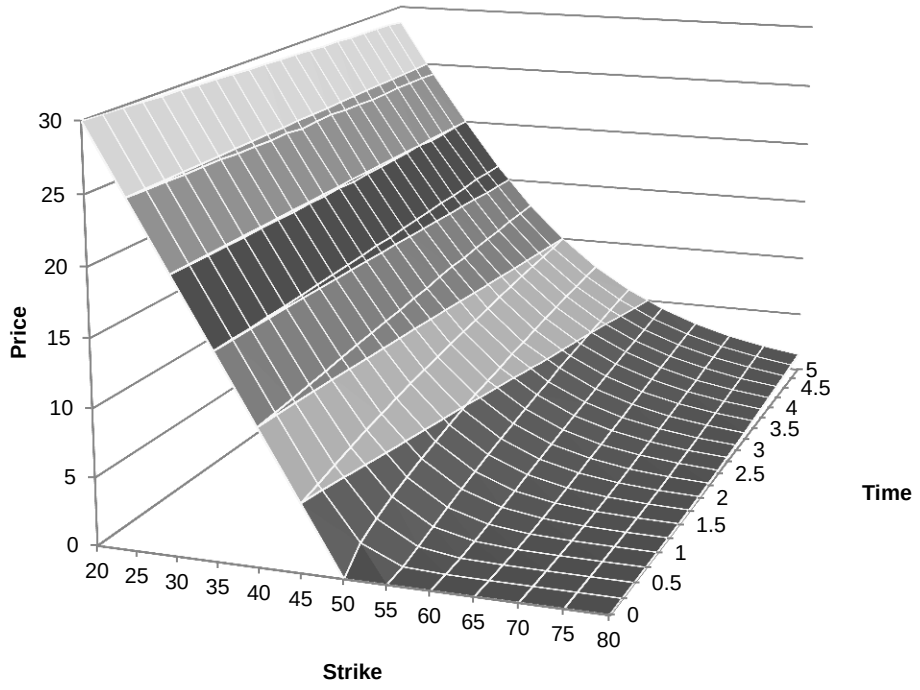


Figure 15. Time to maturity surface of an Asian option. A sensitivity surface of t against strike while keeping other variables constant. $\sigma = 0.2,\ \lambda = 0.5,\ \delta = 0.3,\ \mu = -0.25, F_0 = 50,\ r = q = 0.01,$ $t = 1.$

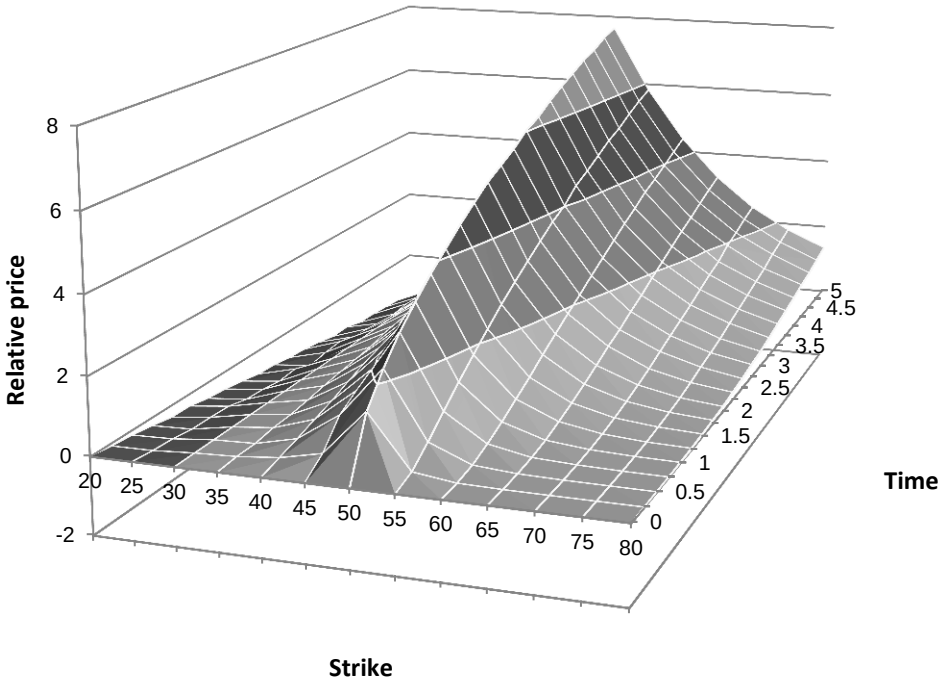


Figure 16. Time to maturity relative surface of an Asian option. A relative sensitivity surface of t against strike. The relative price is measured against the price obtained by the lowest value of t. Other variables are kept constant. $\sigma = 0.2, \lambda = 0.5, \delta = 0.3, \mu = -0.25, F_0 = 50, r = q = 0.01, t = 1.$

and k. Both relative surfaces have similar tendencies as expected from the construction. The most interesting observation is that a decrease in the value of μ and k can also lead to increase in option price. In particular, this is most prevalent around the ATM strike for values of μ and k that are observed in the market most.

The analysis so far has examined how changes in parameters affect Asian option price. However, in order to find which parameters influence the price the most, a relative comparison must be made. In figure 17 a comparison among σ, μ, δ, and λ is made.

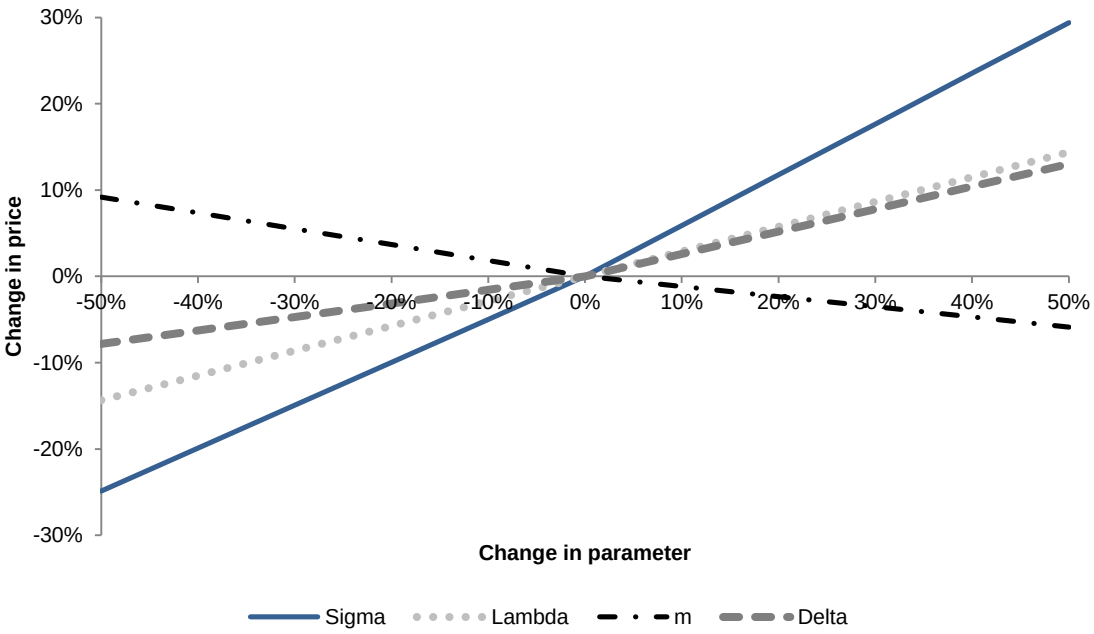


Figure 17. Comparison of Asian option price sensitivity to different parameters. The change in price is calculated against an Asian option with initial values of $\sigma = 0.2$, $\lambda = 0.5$, $\delta = 0.3$, $\mu = -0.25, F_0 = 50$, $X = 50$, $r = q = 0.01$, $t = 1$. A change in a parameter is a relative change from its initial value.

The figure depicts how a relative change in parameters from initial values changes Asian option price. The results show that the diffusion volatility has the largest influence on the price. Also, for all, but log jump size, there is a positive correlation between change in a parameter and change in price. The fact that the log jump size has a negative correlation is also observable when looking at its relative surface in figure 23.

6 Conclusions

The aim of this research is to examine whether information contained in market prices is accurate and useful. To find evidence, I propose two research questions: 1) Do market prices contain valuable information?; and 2) Is information implied in market prices useful for pricing exotic derivatives?

To answer the first research question, I calibrate the Merton jump-diffusion model to oil call options from October 2009 to May 2015. I find that the retrieved parameters capture market events over time accordingly and properly. In particular, the average jump process amplitude escalates (declines) when there is an upward (downward) pressure on oil prices while the diffusion volatility and the standard deviation of jump process amplitude rises with increasing volatility in oil returns. The fluctuations in the average jump process amplitude and the standard deviation of jump process amplitude are more pronounced and these parameters also capture smaller events in the market, while the diffusion volatility varies when there are more fundamental changes in the market. For all parameters, changes are most evident when the parameters are retrieved from options with shorter maturities. In addition, to test predictability I develop a simple trading strategy based on the retrieved diffusion volatility. The results indicate that there is a positive, but marginal gain before transaction costs. Hence, the results do not provide evidence that the future can be predicted from the past market information.

To answer the second research question, I compare Asian option prices estimated from the retrieved parameters with Asian option prices in the market. I find that the relative pricing error fluctuates over time with the highest levels obtained when there is alleviated volatility in the market. The fluctuations are most pronounced for options with shorter maturities, but in all cases the errors are within acceptable levels. In particular, I show that market prices always lie within the range of 95% confidence interval around the estimated prices of ITM, ATM and OTM options. In addition, when Asian options are priced using parameters that are retrieved from call options with similar moneyness, the relative pricing error is substantially improved. Finally, I perform sensitivity analysis and find that in the current market environment an Asian option price is the most sensitive to the diffusion volatility. Therefore, the diffusion volatility must be most carefully examined before the MJD model is applied in practice.

To conclude, the results show that market information implied in oil option prices has strong economic significance. In addition, this information can be used to price simple exotic derivatives, as Asian options. These findings are especially useful for practitioners that price derivatives because the results indicate that exotic derivatives can be priced using market information implied in plain vanilla derivatives. However, before making more general statements, it is important to examine whether the results are similar in nature when other asset classes are considered, other models are used and other derivatives are priced. In addition, it would be valuable for practitioners to see how and if results differ when also other derivatives as put options on the same asset are included in the calibration.

7 References

Aït-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do option markets correctly price the probabilities of movement of the underlying asset?, Journal of Econometrics 102, 67-110.

Andersen, Leif, and Jesper Andreasen, 2000, Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing, Review of Derivatives Research 4, 231-262.

Askari, Hossein, and Noureddine Kirchene, 2008, Oil price dynamics (2002-2006), Energy Economics 30, 2134-2153.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, The Journal of Finance 52, 2003-2049.

Bates, David S., 1991, The crash of'87: Was it expected? The evidence from options markets, The Journal of Finance 46, 1009-1044.

Bates, David S., 1996, Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options, Review of Financial Studies 9, 69-107.

Bates, David S., 2000, Post-'87 crash fears in the S&P 500 futures option market, Journal of Econometrics 94, 181-238.

Bayraktar, Erhan, and Hao Xing, 2011, Pricing Asian options for jump diffusions, Mathematical Finance 21, 117-143.

Björk, Tomas, 2009, Arbitrage theory in continuous time, 3rd edition (Oxford University Press, Oxford).

Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, Journal of Political Economy 81, 637–654.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, The Journal of Finance 62, 1453-1490.

Chernov, Mikhail, and Eric Ghysels, 2000, A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation, Journal of Financial Economics 56, 407-458.

Christoffersen, Peter, Kris Jacobs, and Chayawat Ornthanalai, 2012, Dynamic jump intensities and risk premiums: Evidence from S&P 500 returns and options, Journal of Financial Economics 106, 447-472.

CME Group, 2016a, Energy products, retrieved from http://www.cmegroup.com/trading/energy

CME Group, 2016b, WTI average price options contract specs, retrieved from http://www.cmegroup.com/trading/energy/crude-oil/west-texas-intermediate-wti-crude-oil-calendar-swap-futures_contractSpecs_options.html

Cont, Rama, and Peter Tankov, 2004, Nonparametric calibration of jump-diffusion option pricing models, Journal of Computational Finance 7, 1-49.

Coval, Joshua D., Jakub W. Jurek, and Erik Stafford, 2009, The economics of structured finance, Journal of Economic Perspectives 23, 3-25.

Das, Sanjiv R., and Silverio Foresi, 1996, Exact solutions for bond and option prices with systematic jump risk, Review of Derivatives Research 1, 7-24.

Derman, Emanuel, and Iraj Kani, 1994, Riding on a smile, Risk 7, 139-145.

Dumas, Bernard, Jeff Fleming, and Robert E, Whaley, 1998, Implied volatility functions: Empirical tests, The Journal of Finance 53, 2059-2106.

Dupire, Bruno, 1994, Pricing with a smile, Risk 7, 18-20.

Giot, Pierre, Sébastien Laurent, and Mikael Petitjean, 2010, Trading activity, realized volatility and jumps, Journal of Empirical Finance 17, 168-175.

Glasserman, Paul, 2004, Monte Carlo methods in financial engineering (Springer, New York).

He, Changhong, J. Shannon Kennedy, Thomas F. Coleman, Peter A. Forsyth, Yuying Li, and Kenneth R. Vetzal, 2006, Calibration and hedging under jump diffusion, Review of Derivatives Research 9, 1-35.

Heston, Steven L., 1993, A closed-form solution for options with stochastic volatility applications to bond and currency options, Review of Financial Studies 6, 327-343.

Hull, John, and Alan White, 1987, The pricing of options with stochastic volatilities, The Journal of Finance 42, 281-300.

Kearney, Fearghal, Finbarr Murphy, and Mark Cummins, 2015, An analysis of implied volatility jump dynamics: Novel functional data representation in crude oil markets, The North American Journal of Economics and Finance 33, 199-216.

Kou, Steven G., 2002, A jump diffusion model for option pricing, Management Science 48, 1086-1101.

Matsuda, Kazuhisa, 2005, Parametric regularized calibration of Merton jump-diffusion model with relative entropy: What difference does it make?, working paper, The City University of New York.

McIntyre, Michael L., and David Jackson, 2009, Market moves and the information content of option prices, International Review of Economics & Finance 18, 327-340.

Merton, Robert C., 1976, Option pricing when the underlying stock returns are discontinuous, Journal of Financial Economics 3, 125-144.

Murphy, Finbarr, and Ehud I. Ronn, 2015, The valuation and information content of options on crude-oil futures contracts, Review of Derivatives Research 18, 95-106.

Rappoport, Peter, and Eugene N. White, 1994, Was the crash of 1929 expected?, The American Economic Review 84, 271-281.

Rubinstein, Mark, 1994, Implied binomial trees, The Journal of Finance 49, 771-818.

Tankov, Peter, and Ekaterina Voltchkova, 2009, Jump-diffusion models: A practitioner's guide, Banque et Marchés 99.

Yan, Shu, 2011, Jump risk, stock returns, and slope of implied volatility smile, Journal of Financial Economics 99, 216-233.

8 Appendix: sensitivity surfaces

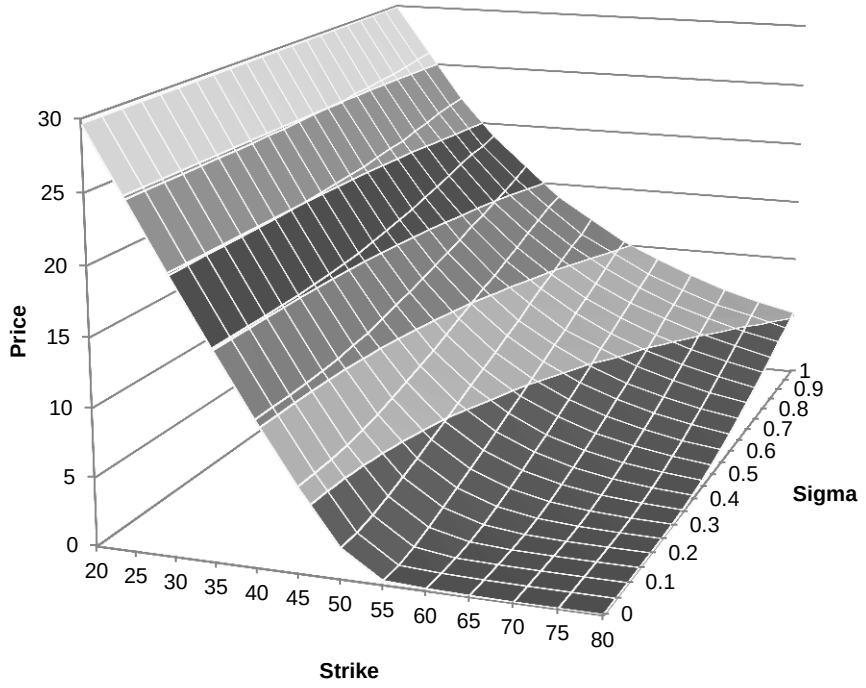


Figure 18. Diffusion volatility surface of an Asian option. A sensitivity surface of σ while keeping other variables constant. $\lambda = 0.5$, $\delta = 0.3$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

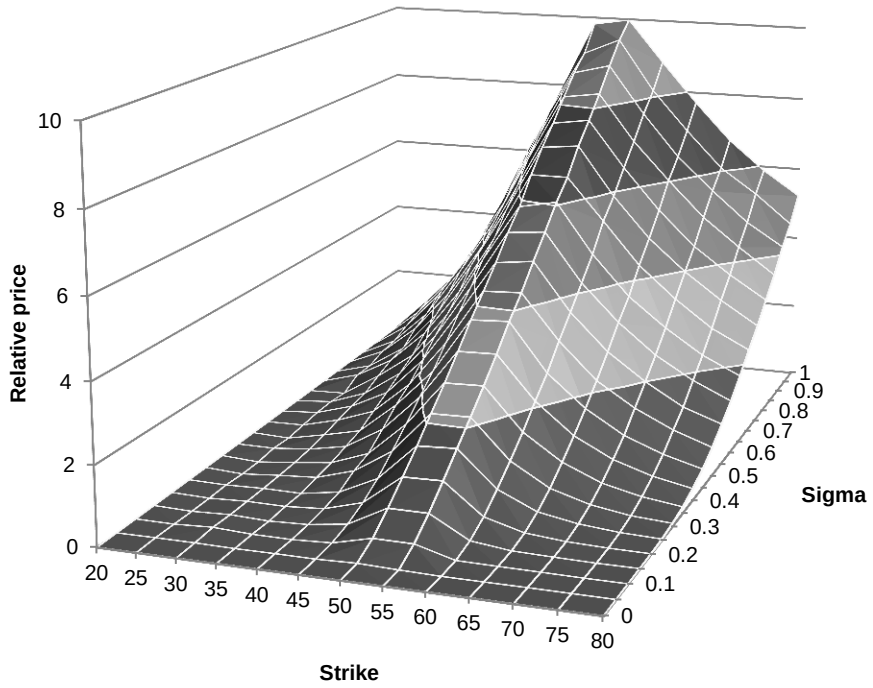


Figure 19. Diffusion volatility relative surface of an Asian option. A relative sensitivity surface of σ. The relative price is measured against the price obtained by the lowest value of σ. Other variables are kept constant. $\lambda = 0.5$, $\delta = 0.3$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

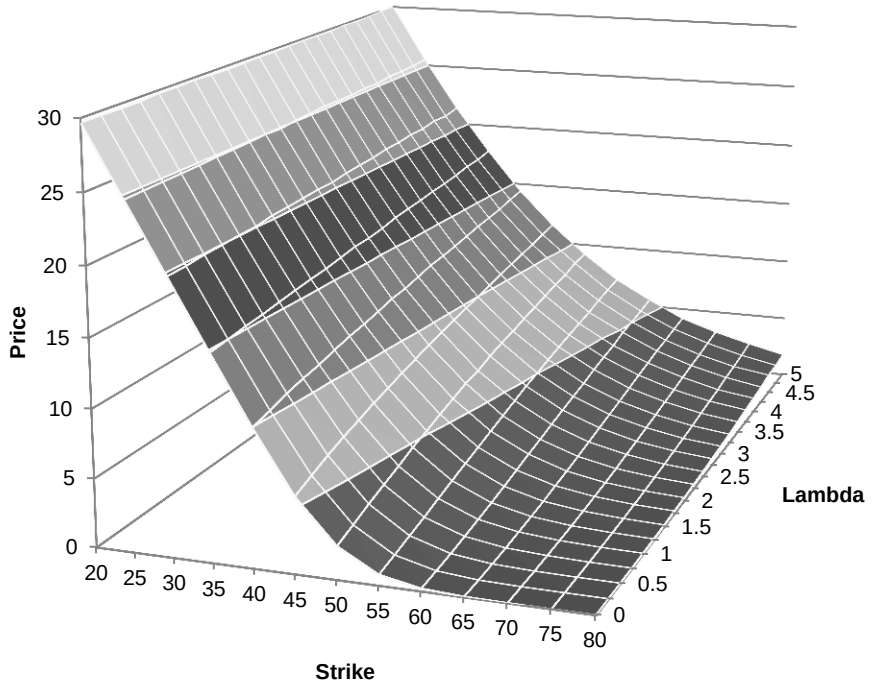


Figure 20. Jump intensity surface of an Asian option. A sensitivity surface of λ while keeping other variables constant. $\sigma = 0.2$, $\delta = 0.3$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

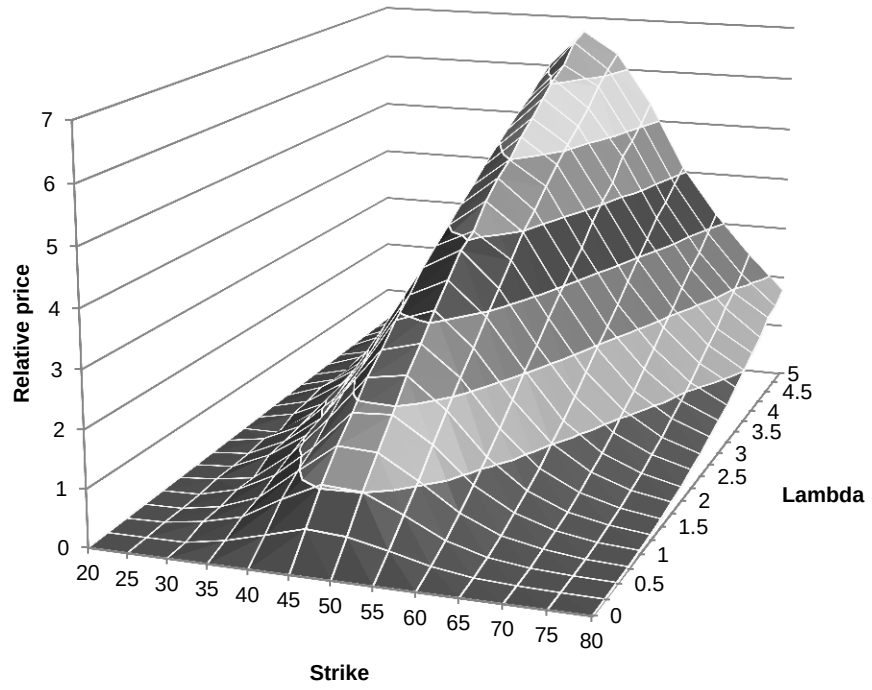


Figure 21. Jump intensity relative surface of an Asian option. A relative sensitivity surface of λ. The relative price is measured against the price obtained by the lowest value of λ. Other variables are kept constant. $\sigma = 0.2$, $\delta = 0.3$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

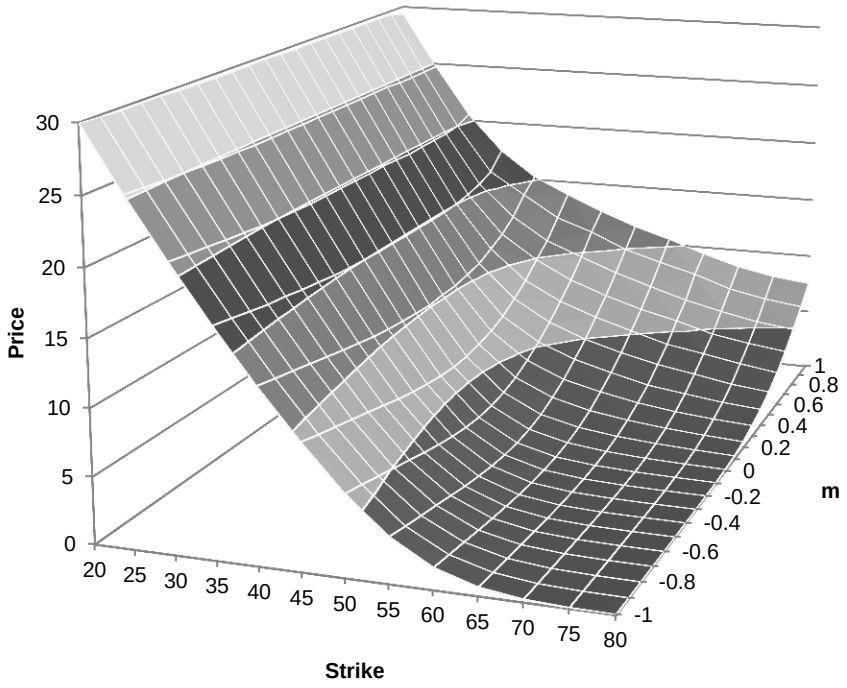


Figure 22. Log jump size surface of an Asian option. A sensitivity surface of μ while keeping other variables constant. $\sigma = 0.2$, $\lambda = 0.5$, $\delta = 0.3$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

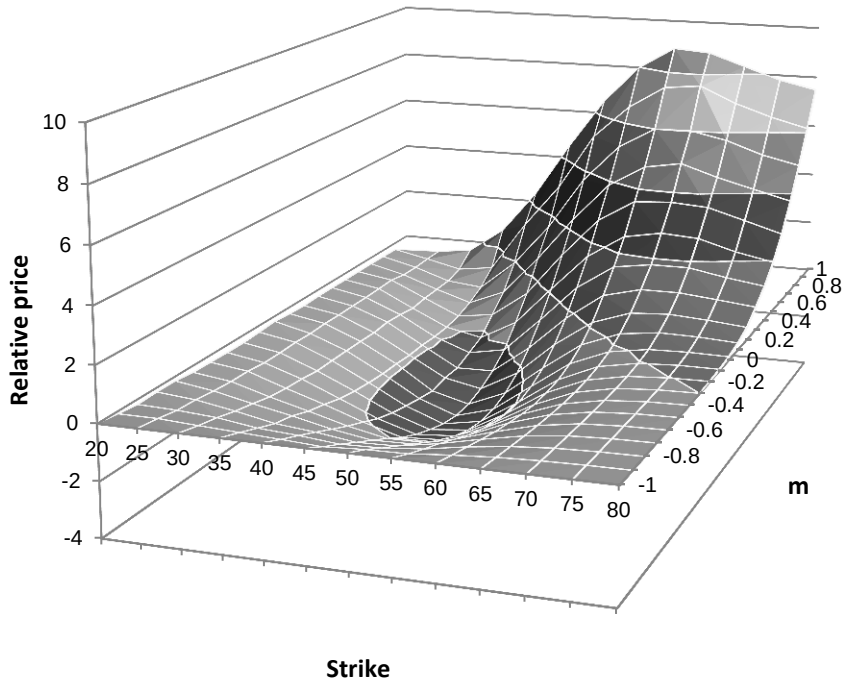


Figure 23. Log jump size relative surface of an Asian option. A relative sensitivity surface of μ. The relative price is measured against the price obtained by the lowest value of μ. Other variables are kept constant. $\sigma = 0.2$, $\lambda = 0.5$, $\delta = 0.3$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

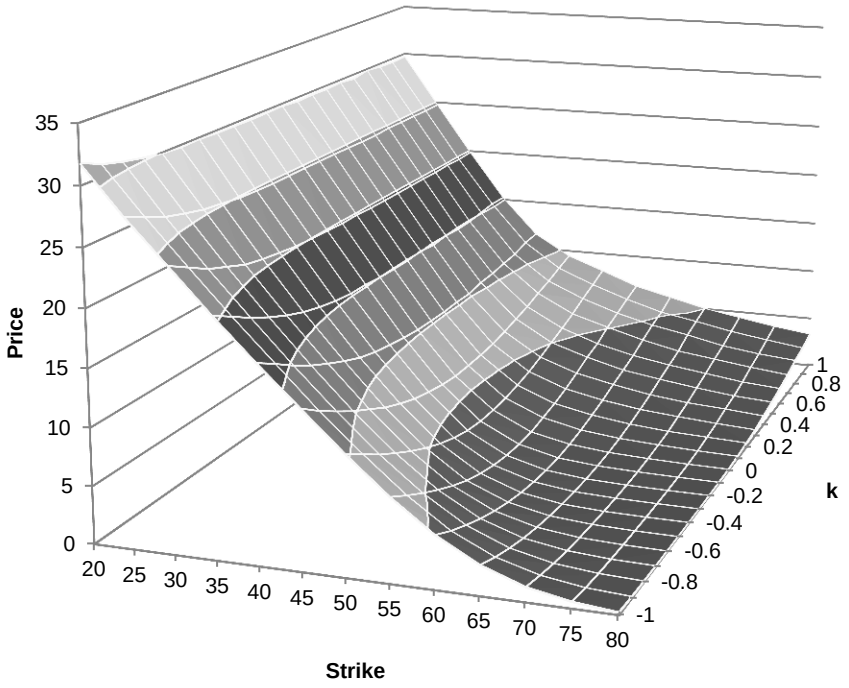


Figure 24. Average jump process amplitude surface of an Asian option. A sensitivity surface of k while keeping other variables constant. $\sigma = 0.2$, $\lambda = 0.5$, $\delta = 0.3$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

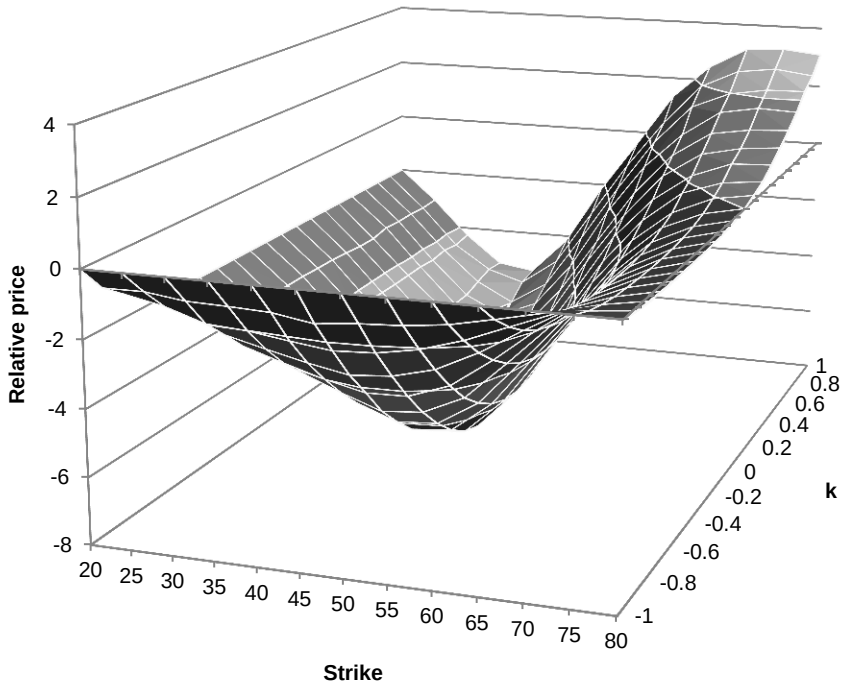


Figure 25. Average jump process amplitude relative surface of an Asian option. A relative sensitivity surface of k. The relative price is measured against the price obtained by the lowest value of k. Other variables are kept constant. $\sigma = 0.2$, $\lambda = 0.5$, $\delta = 0.3$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

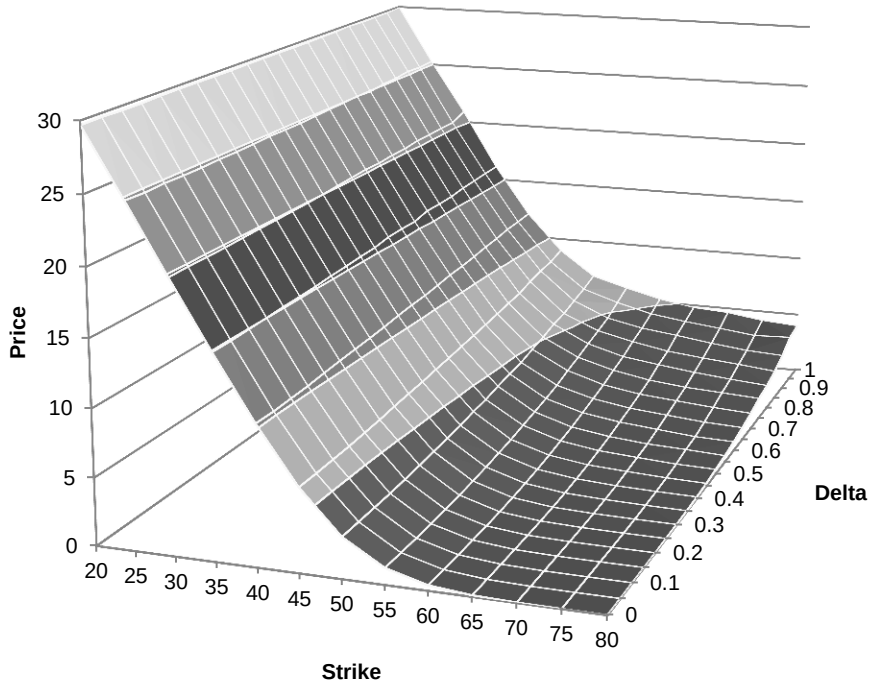


Figure 26. Standard deviation of jump process amplitude surface of an Asian option. A sensitivity surface of δ while keeping other variables constant. $\sigma = 0.2$, $\lambda = 0.5$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.

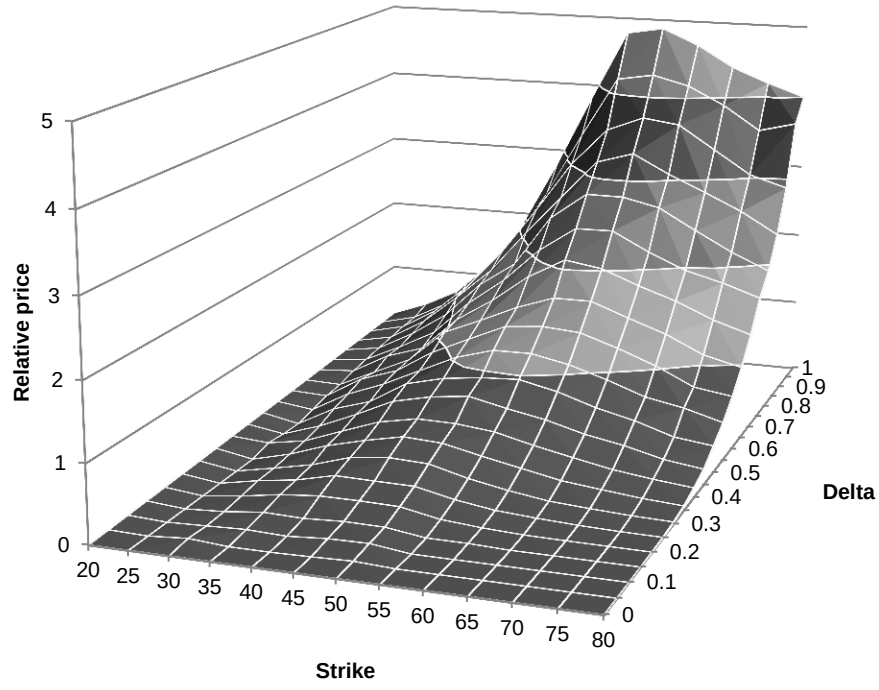


Figure 27. Standard deviation of jump process amplitude relative surface of an Asian option. A relative sensitivity surface of δ. The relative price is measured against the price obtained by the lowest value of δ. Other variables are kept constant. $\sigma = 0.2$, $\lambda = 0.5$, $\mu = -0.25$, $F_0 = 50$, $r = q = 0.01$, $t = 1$.