Norwegian School of Economics
Bergen, Spring 2013

On the Relevance of Jumps for the Pricing of S&P 500 Options

With Particular Emphasis on the Adjustment for Systematic Risk in Jump-Diffusion Models

Børge Langedal and Sindre Sunde

Supervisor: Jørgen Haug

Master Thesis in Financial Economics

NORWEGIAN SCHOOL OF ECONOMICS

Abstract

Jump-diffusions are a class of models that is used to model the price dynamics of assets whose value exhibit jumps. The first part of this thesis discusses the implications of such models for the pricing of derivatives. Particular emphasis is put on explaining the adjustment for systematic risk. Efforts are made to link purely mathematical arguments with economic theory and intuitive explanations.

In the second part, the theoretical framework for derivatives pricing are applied to answer the question whether jumps are relevant for the pricing of European options with the S&P 500 index as the underlying asset. Analysis of the distributional properties of log-returns leads to the suggestion of a specific jump-diffusion model for the dynamics of this index. The model is calibrated to market data on a daily basis for a period of 80 trading days prior to and 80 trading days after what is considered the outbreak of the financial crisis of 2008. Obtained values of the jump-diffusion parameters implicit in option prices establish that jumps are relevant for their value.

Preface

This thesis is written as part of our major in financial economics at the Norwegian School of Economics (NHH). Provided with the opportunity to devote an entire semester to study a topic of interest, the subject of jump-diffusion models within derivatives pricing appeared as a choice of particular appeal.

Through the work on this paper, we have gained valuable knowledge on both the theoretical and practical aspects of jump-diffusion models. As the chosen area of research is one that is characterized by an extensive use of advanced mathematics and statistics, the project has also been challenging. This has required in-depth studies of selected parts of the relevant literature.

We would like to thank our supervisor, Jørgen Haug, for helpful guidance and interesting discussions, along with his accessibility throughout the work of this thesis. While also making the topic more exciting, this has contributed significantly to the learning outcome of the process.

Bergen, 19th of June 2013

Børge Langedal and Sindre Sunde

Contents

1. Introduction

1.1 Motivation

Derivatives are flexible and powerful investment tools. Due to the wide range of payoff functions, they are used to achieve a variety of different goals. These goals range from hedging risk, obtaining exposure to an underlying that is non-tradable (e.g. air temperature), exploiting arbitrage opportunities, speculation, or simply leverage one's exposure to an underlying. Thus, it is of great value to acquire the knowledge required to price derivatives with precision.

In spite of the flexibility these financial contracts exhibits, the framework used to price them are general and applicable to a wide range of valuation problems. Within this framework, a model for the price dynamics of the underlying asset(s) is needed. The choice of model may have important implications for pricing. In some markets, jump-diffusion models seem to better capture the real dynamics than models not accounting for jumps. Hence, to study the implications for derivatives pricing when the underlying asset follows a jump-diffusion is of great interest.

Furthermore, the impression of the authors is that much of the existing literature on jump-diffusions mainly focuses on the mathematical aspects. This may hide the economic principles these models are built upon. Hopefully, by putting emphasis on explaining the underlying economic arguments, this thesis can add value to the discussion of jump-diffusion models.

Besides examining the theoretical aspects of jump-diffusion models, it is also desirable to get acquainted with the practical application. Due to the widespread attention directed towards the U.S. stock market, along with the liquidity in the market for derivatives, options written on the broad S&P 500 index is considered exciting to study in the context of a jump-diffusion model.

A final motivating factor is related to the information implicit in option prices. Option pricing models are for instance used to extract investor expectations regarding future volatility in the underlying asset. While this is often summarized in a single parameter called *implied volatility*, the use of jump-diffusion models may provide additional information regarding the assessment of risk.

1.2 Topic

The topic of this paper is twofold. First, focus is directed towards the development of a framework for pricing financial contracts derived from assets whose dynamics could be represented by a jump-diffusion. Particular emphasis is put on explaining the adjustment for systematic risk within such models. The objective is to highlight the underlying economic arguments that the pricing model is based upon.

Second, the thesis aims to determine whether price jumps in the underlying asset are relevant for the pricing of S&P 500 index options of European type. To investigate this, a pricing model is developed and then calibrated to historical option prices. Focusing on a time interval comprising a period prior to and a period after the outbreak of the financial crisis of 2008 is considered appropriate. Studying this particular period allows for answering two related research questions of interest. One is the question if there was an assessed risk of a market crash in September 2008. The other is to what extent the market's perception of risk changed after the financial crisis hit.

1.3 Refinements

Certain refinements are considered appropriate in order to ensure a narrow focus. These are justified by the fact that including them complicates the analysis without adding significant value. It is assumed that the reader is familiar with basic option pricing theory.

A *derivative security*, or a *contingent claim*, can be defined as an instrument whose value depends on the price of another asset, often referred to as the *underlying* (Hull, 2012). There exists a variety of such contracts distinguished by different payoff functions. In this thesis, however, only one type is considered. That is, call options of European type. To be clear, this is a contract that gives the holder the right to buy the underlying security at a predefined price at a specific date.

The price of derivatives depends upon parameters whose value is uncertain. For simplifying purposes, some of these are assumed to be constant. In particular, this assumption is applied to the *risk-free rate of return* and the *dividend yield* of the underlying security.

In the financial literature, *stochastic volatility models* have gained widespread popularity. However, such models are not considered here.

1.4 Thesis Structure

The following structure is a reflection of what is considered the best approach to achieve the stated objectives of the thesis.

Chapter 2 presents the theoretical foundation for jump-diffusion models. Intuitive reasons to why modeling jumps may add value is provided through relating model components to the occurrence of rare and normal events. Then, in chapter 3, it is turned to the issue of adjusting for systematic risk within such models. Due to this subjects' fundamental importance for pricing and hedging, attempts are made to uncover the underlying economic principles.

Chapter 4 marks the start of the applied part of the thesis. In this chapter, a pricing model for European options with the S&P 500 index as the underlying asset is suggested. The choice of model is supported by the results of statistical tests of the distributional properties of S&P 500 returns. In order to extract investor expectations from market data, the model is calibrated to historical option prices. The choice of calibration approach and its implementation is described in chapter 5, while the chapter 6 is devoted to the presentation and discussion of the results.

Finally, conclusions are summarized in chapter 7. Here, a discussion of limitations is also provided, along with suggestions for further research. A flowchart of the thesis' structure is presented in Exhibit 1.1.

Exhibit 1.1 - Thesis Flowchart

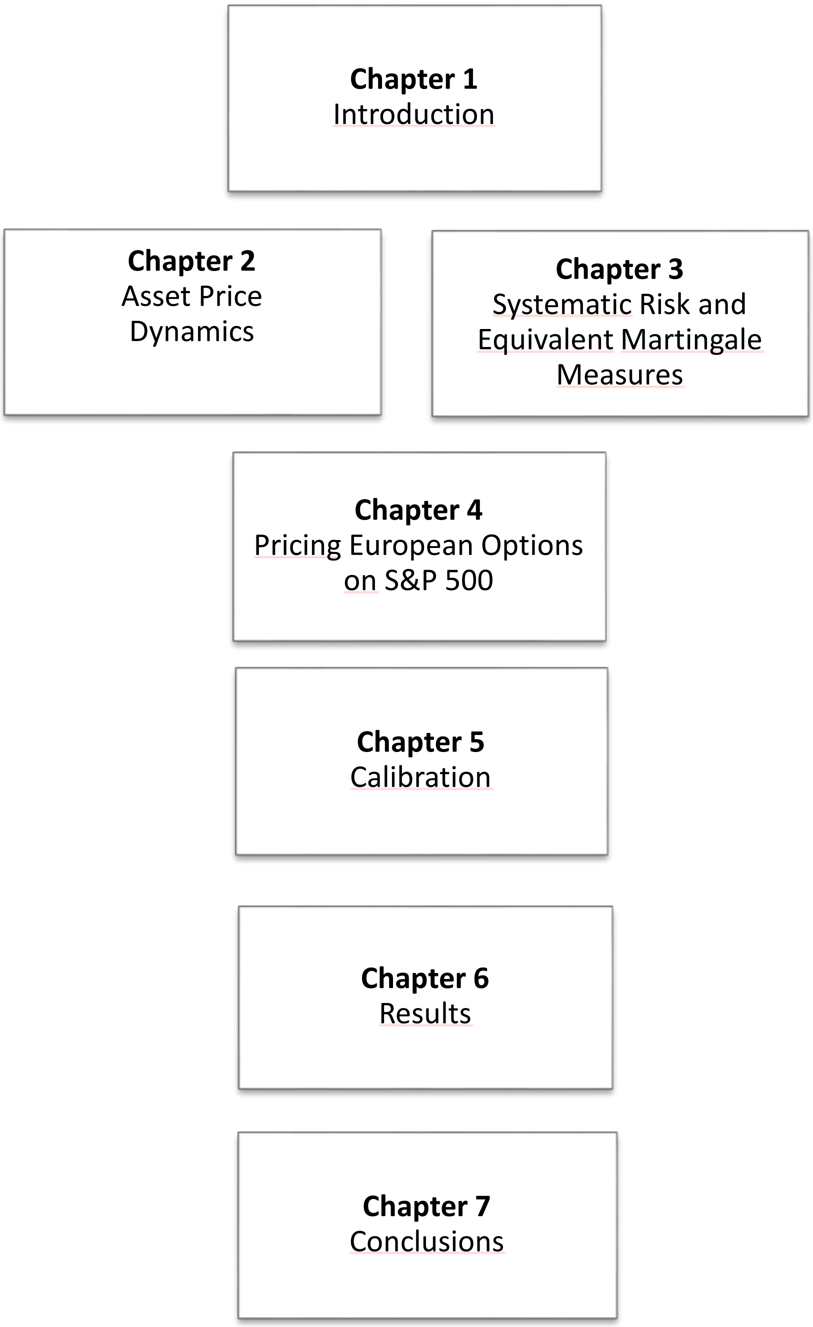

2. Modeling Asset Price Dynamics

There is a variety of ways to model jumps in the price of an asset. In this chapter, a general jump-diffusion is developed. The approach is based on Wiener and Poisson processes.

2.1 The General SDE

The change in the value of an asset over time is in general uncertain, i.e. it is stochastic (Hull, 2012). Stochastic differential equations (SDEs) provide a framework for modeling asset price dynamics. It can be shown that, under some mild assumptions[1], the behavior of a continuous-time stochastic process S_t can be approximated by the general SDE

$$dS_t = \mu(S_t, t)dt \ + \ b(S_t, t)dX_t \, , \tag{1}$$

where

> $\mu(S_t, t)$ is the *drift coefficient*,
> $b(S_t, t)$ is the *diffusion coefficient*,
> dX_t is an *innovation term*.

The first term on the right hand side of (1) represents the expected change in the security price over the infinitesimal time interval dt (Neftci, 1996). $\mu(S_t, t)$ is then the instantaneous absolute expected return on t he asset. Since investors are risk-averse, they demand compensation for taking on non-diversifiable risk. Hence, the drift rate will deviate from the risk-free rate of return if there are systematic risks inherent in S_t.

Unpredictable changes in the price of S_t in the given time interval is represented by the second term. This is sometimes referred to as the *dispersion term*. It consists of the diffusion coefficient $b(S_t, t)$ and the innovation term, dX_t. The latter incorporates the uncertainty in the price process, and has expectation equal to zero - i.e. it is a martingale[2] (Neftci, 1996).

---

[1] See Neftci (1996) ch.7 p.136-137 for a discussion of these assumptions.

[2] The concept of martingales is thoroughly explained in chapter 3.

Note that both the drift and diffusion term, $\mu(S_t, t)$ and $b(S_t, t)$, are $I_t - adapted$[3] (Neftci, 1996). That is, their values are known given the information set I_t.

2.2 Introducing Jump-Diffusion Models

The dynamic behavior of the underlying asset(s) is given by the general SDE (1). In order to use this for pricing purposes one has to specify its distributional properties. Critical to this is the distinction between rare and normal events (Neftci, 1996).

2.2.1 Rare and Normal Events

According to Neftci (1996), a rare event is defined as something that has a "large" size and occurs infrequently. These differ from normal events, which occur in a routine fashion with smaller magnitude.

Consider an observation interval, dt. The formal distinction between rare events and normal events is the way their size and their probability of occurrence vary with this interval (Neftci, 1996). As the interval gets smaller, the size of normal events also gets smaller. However, because they are ordinary, their probability of occurrence is not zero. That is, in short time intervals, it will always be a non-zero probability that some normal event occurs. For rare events, this is not the case. As $dt \to 0$, the probability of occurrence also goes to zero. However, in contrast to normal events, the size of the event may not shrink. Hence, it represents a discontinuous jump in the price of the asset in question.

The intuition behind rare and normal events can be explained by the nature of price sensitive news. Normal events can be seen as small price changes due to the flow of "non-noticeable" news. For the stock market, such news may include small changes to investor expectations about future corporate earnings due to FED statements. The following normal event is the marginal change in prices caused by this. Or for a market like oil, a normal event can be an

[3] A variable a_t is said to be $I_t - adapted$ if its value is included in the information set, I_t. That is, a_t is known given I_t. The information set contains all relevant information at time t, and may include historical prices, trading volumes, market volatility, etc. For a more detailed explanation, see Neftci (1996).

unexpected, but marginal increase in demand due to changing weather forecasts. Such news is the cause for the majority of price changes.

On the other hand, rare events can be interpreted as consequences of "big" news. The stock market crash in 1987 is a good example of a rare event. A recent oil related example of such an event, is the political and social unrest in Libya, which caused a large jump in the price of oil. Another example is the large change in the stock price of a pharmaceutical company receiving FDA denial for a promising drug. During a short time interval, the probability of such events approaches zero. Still, when they occur, their size may not be very different whether one looks at large or small time intervals (Neftci, 1996).

2.2.2 A General Jump-Diffusion Model

As stated, dX_t in the general SDE (1) has to be modeled. It is clear from the above discussion that this innovation term should account for both continuous and discontinuous price changes (Neftci, 1996). Two basic building blocks for doing this are the Wiener process for normal events, and the Poisson process for rare events.

The standard Wiener process, denoted W_t, is a natural choice for modeling normal events (Neftci, 1996). This process has normally distributed increments, dW_t, with expectation and variance equal to zero and dt, respectively[4]. The normal distribution has tails that extend to infinity. However, since the variance is time-dependent, the tails will disappear as the time interval dt approaches zero. Hence, the distribution will be concentrated on zero. That is, for small time intervals, the Wiener process is only suitable for modeling small price changes. This is consistent with the discussion above. Exhibit 2.1 illustrates the evolution of a standard Wiener process over time, along with its corresponding increments. It is noted that the scales of the y-axes are different.

[4] Formally, the standard Wiener process has the following important properties:

1) $W_0 = 0$
2) $W_t - W_s \sim \mathrm{N}(0, t - s) \quad for \quad t \geq s$
3) $W_{t_0},\ W_{t_1} - W_{t_0}, W_{t_2} - W_{t_1}, \dots, W_{t_n} - W_{t_{n-1}} \qquad t_0 \leq t_1 \leq \dots \leq t_n$
   is independent (increments) for any integer n
4) W has continuous sample paths, i.e. $W_t(\omega) = W(t, \omega)$ is continuous in t for a given ω (state).

Exhibit 2.1 – Left Panel: Illustration of a Standard Wiener Process
Right Panel: The Increments of a Standard Wiener Process

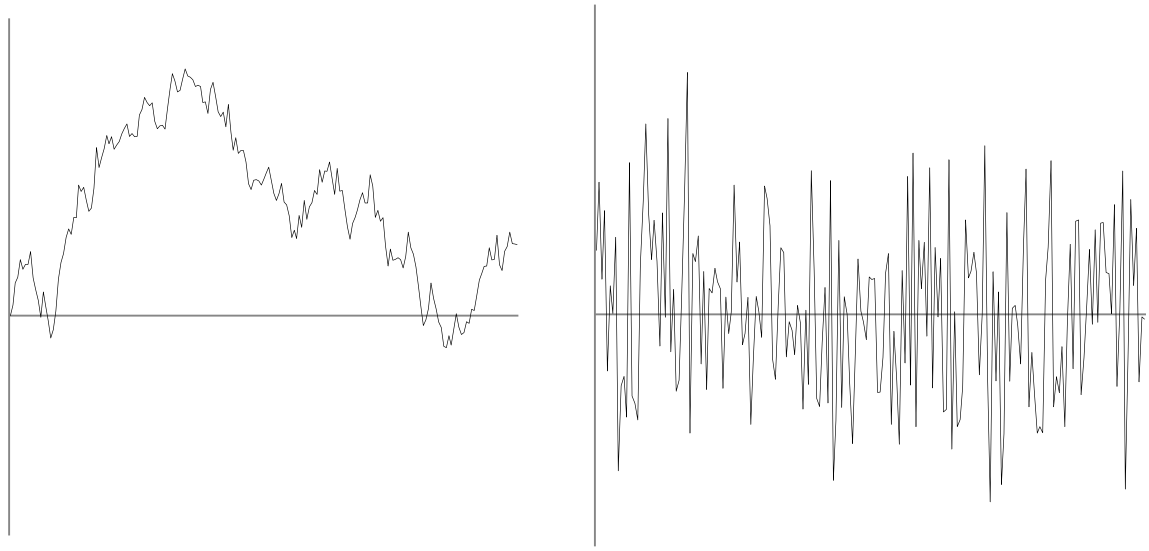

However, the Wiener process is not appropriate for modeling rare events. Instead, a process that is capable of generating large price changes in very small time increments is needed. In other words, the actual process must exhibit discontinuous jumps, i.e. the process must have outcomes that are independent of dt. This can be modeled in different ways. Frequently suggested in the literature on derivatives pricing are Poisson processes[5]. Such processes will be used to model jumps in this thesis.

A particular type of Poisson process that is suitable for modeling jumps in financial markets are *the Poisson counting process*, N_t. This process represents the total number of changes that occur until time t, and is Poisson distributed. Hence, its expectation and variance are identical. The increments in N_t, denoted dN_t, can take on two possible values. Either they are zero, meaning no change, or they are equal to one, representing change[6] (Neftci, 1996).

[5] See for example Neftci (1996), Wilmott (2007) or Hull (2012).

[6] The increments has the following probability distribution:

$$dN_t = \begin{cases} 1 \; with\; probability\; \lambda dt \\ 0 \; with\; probability\; 1-\lambda dt \end{cases}$$

Also important is it that the number of changes occurring in non-overlapping intervals is independent.

Exhibit 2.2 - Illustration of a Poisson Counting Process[7]

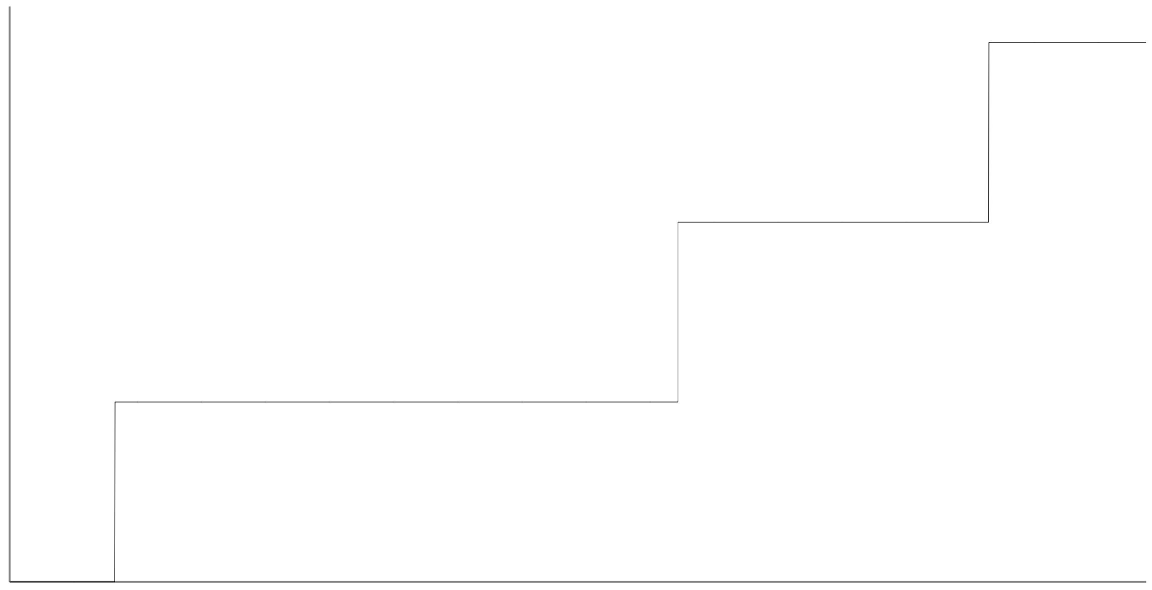

As explained above, N_t has constant jump sizes equal to one. When modeling financial markets this seems to be unrealistic. However, it is fairly easy to allow for random jump sizes. Let N_t be a *Poisson counting process* with intensity λ. Further, let Y_t be a stochastic process with a predetermined distribution $f(y)$ and mean $\kappa = E[Y_i]$. If Y_t also is independent of the Poisson process N_t, one can define the *compound Poisson process* (Shreve, 2004)

$$Q_t = \int_0^t Y_t dN_t = \sum_{i=1}^{N_t} Y_i \,. \tag{2}$$

This can be seen as an extension of the pure *counting process*. The jumps in Q_t arrive at the same times as before. That is, the jumps occur when N_t equals 1. However, whereas the jumps in N_t are constantly 1, the size of the jumps in Q_t is random (determined by the distribution of Y_t). The compound Poisson process has mean $\kappa\lambda t$ (Shreve, 2004).

[7] A minor flaw in the exhibit is that the jumps, i.e. the vertical lines, are not perfectly vertical. Still, the figure is considered suitable for illustrational purposes.

Exhibit 2.3 - Illustration of a Compound Poisson Process

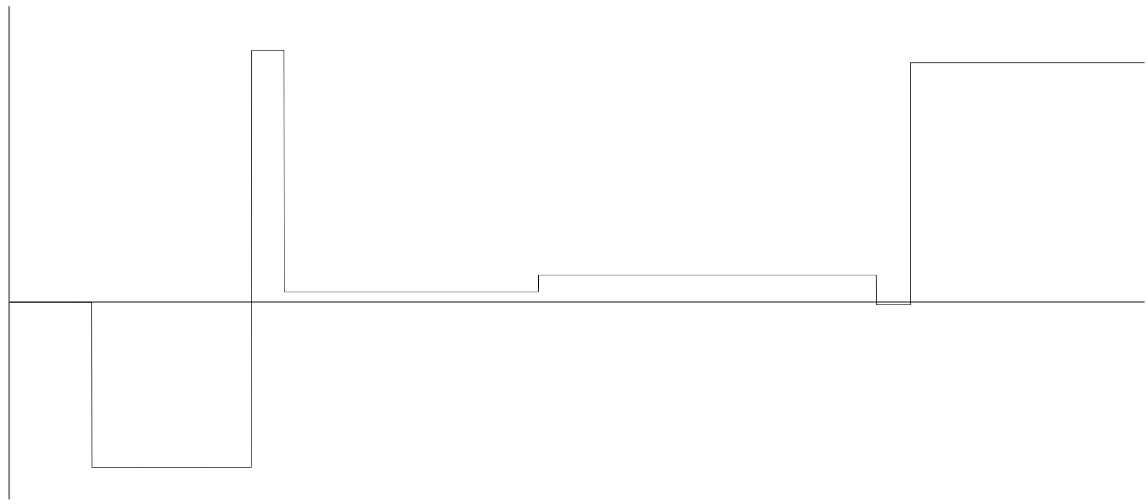

From now on, the compound Poisson process will be used for modeling jumps. In order to be consistent with the general SDE (1), it has to be compensated. That is, Q_t has to be adjusted by subtracting its mean. Thus, one can define

$$J_t = N_t - \kappa\lambda t\,, \tag{3}$$

where J_t is a *compensated compound Poisson process*, i.e. $E[J_t] = 0$.

By splitting up the innovation term into a *standard Wiener process* for normal events and a *compensated compound Poisson process* for jumps, one ends up with a general model for asset price dynamics given by (Nefcti, 1996)[8]

$$dS_t = \mu(S_t, t)dt + \sigma_1(S_t, t)dW_t + \ dJ_t\,. \tag{4}$$

Here,

$\mu(S_t, t)$ is the expected change in S_t,

$\sigma_1(S_t, t)$ is the diffusion coefficient conditional on no jump,

W_t is a standard Wiener process,

J_t is a compensated compound Poisson, $J_t = \sum_{i=1}^{N_t} \sigma_2(S_t, i) \ - \kappa\lambda t$, with intensity λ, and random jump size $\sigma_2(S_t, t)$; $\kappa \equiv E[\sigma_2(S_t, t)]$.

[8] The notation in Neftci (1996) is slightly different.

As before, the dt-term represents the expected change in S_t during dt. In contrast to (1), the dispersion is now modeled by two separate terms, i.e. a diffusion term and a jump-term. Hence, this model is slightly less general than (1), and is often referred to as a jump-diffusion (Neftci, 1996). Furthermore, the parameters $\mu(S_t, t)$, $\sigma_1(S_t, t)$ and $\sigma_2(S_t, t)$ are still $I_t - adapted.$ A random sample of a jump-diffusion model is illustrated in Exhibit 2.4.

Exhibit 2.4 - Illustration of a Jump-Diffusion

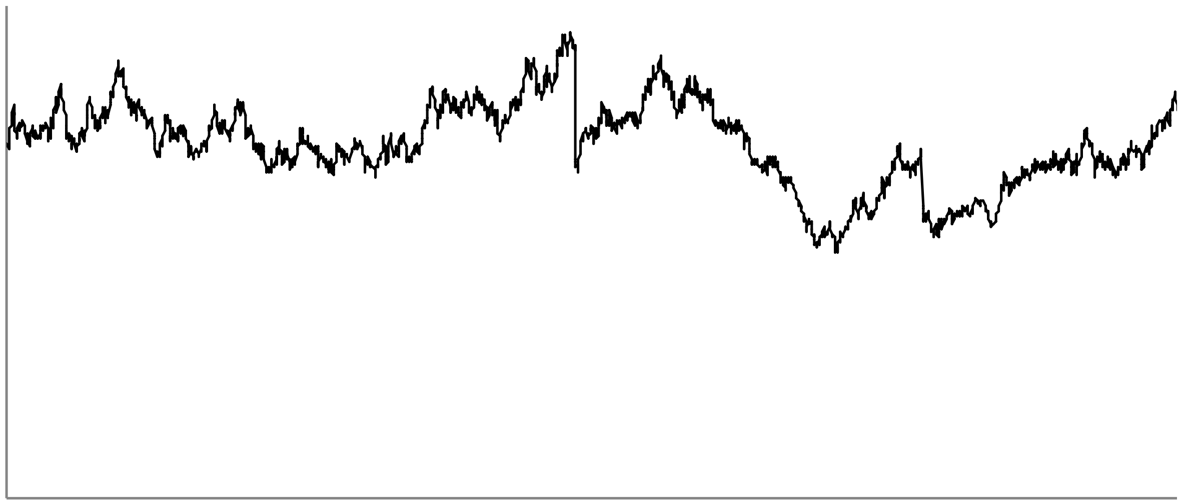

It is noted that the dispersion terms in (4) have to be martingales to be consistent with (1). In addition, they have to be independent, i.e. the Wiener process and the Poisson process have to be independent at every instant t (Neftci, 1996).

Note that (4) also can be written as[9]

$$dS_t = [\mu(S_t, t) - \kappa\lambda]dt + \sigma_1(S_t, t)dW_t + \sigma_2(S_t, t)dN_t. \tag{5}$$

This notation makes the construction of the jump term more clear. A special case of the jump-diffusion is when the jump sizes, $\sigma_2(S_t, t)$, are constantly equal to zero. The resulting classes of models are called diffusions.

[9] As noted, $Q_t = \int_0^t \sigma_2(S_t, t)dN_t$. In differential form, $dQ_t = \sigma_2(S_t, t)dN_t$. Thus $dJ_t = \sigma_2(S_t, t)dN_t - \kappa\lambda dt$.

3. Systematic Risk and Equivalent Martingale Measures

Systematic risk is one of the most discussed topics in financial theory. Asset prices depend critically upon the size of this parameter. This is the focus of the following section.

3.1 Definition

Before elaborating on *systematic risk*, it seems appropriate with a definition. According to Hull (2012), it is risk that is related to the return from the market as a whole and cannot be diversified away. On the other hand, *non-systematic risk*, also referred to as *idiosyncratic risk*, is risk that is unique to the asset and can be diversified away (Hull, 2012).

3.2 The General Pricing Problem

Asset prices are determined by *the law of one price*[10]. This implies that security prices must preclude arbitrage opportunities[11]. It is, however, in many cases not a straightforward task to obtain a fair price for an asset. To illustrate the general pricing problem, it can be useful to consider a one-period setting. The expected rate of return of an asset S_t at time t is given by

$$\mu_t = \frac{E[S_t]}{S_0} - 1\,, \tag{6}$$

where S_0 is the initial value of the asset and μ_t is the expected rate of return. Note that μ_t here is defined as a percentage, while $\mu(S_t, t)$ in (1) and (4) is expressed as an absolute value. In other words, $\frac{\mu(S_t,t)}{S_0}$ is equivalent to μ_t . Rearranging terms in (6) gives an equation for the present value,

$$S_0 = \frac{E[S_t]}{1+\mu_t}. \tag{7}$$

---

[10] The Law of One Price states that identical cash flows must have the same price. If this is not the case, an arbitrage exists, which is not consistent with economic theory. For more on this, see for example Berk, J. and P. DeMarzo (2010).

[11] In its simplest form, an arbitrage opportunity is said to exist if it is possible to achieve a riskless profit greater than the riskless rate of return by taking simultaneous positions in different assets (Neftci, 1996).

Equation (7) illustrates that uncertainty are present both in the numerator and denominator. It states that the value today is equal to the expected value at time t discounted with the appropriate discount rate, which is the asset's expected rate of return. Since investors are assumed to be risk-averse, they will demand a premium for taking on non-diversifiable risk. For a t ypical risky asset, μ_t has to be greater than the risk-free rate of return, r_t. If not, investors will only invest in risk-free assets. Note that μ_t sometimes are smaller than r_t. This is the case when the asset in question provides insurance, i.e. it has a negative covariance with the market portfolio. In conclusion, μ_t is determined by the degree of systematic risk (Neftci, 1996).

Academics have identified numerous challenges with estimating this parameter. These are heavily discussed in modern textbooks in finance, and are not considered here[12]. In fact, the expected rate of return is almost impossible to accurately estimate ex ante (Hull, 2012). As a consequence, pricing methods that maneuver around this problem have been developed.

3.3 Change of Measure

3.3.1 The EMM Approach

A method that does not require an estimate of μ_t is that of *Equivalent Martingale Measures* (EMM) (Miltersen, 2005). These are alternative probability measures that are used for pricing purposes. Formally, $\mathbb{Q}$ is an equivalent martingale measure relative to $\mathbb{P}$ if

$$\mathbb{Q}(\mathrm{E}) = 0 \Leftrightarrow \mathbb{P}(E) = 0 \tag{8}$$

and

$$\frac{C_t}{B_t} = E^{\mathbb{Q}}\left[\frac{C_T}{B_T}|I_t\right], \tag{9}$$

[12] See, e.g., Berk, J. and P. DeMarzo (2010) or Brealey, Myers and Marcus (2009).

for any T, $0 \leq t \leq T$, and for any price process, C_t, in the economy. B_t is referred to as the *numeraire*, i.e. the security price used as the discount factor. The choice of B_t is a matter of convenience and should be chosen to best simplify the calculations[13].

In short, the EMM approach is a method that makes all discounted price processes martingales. A benefit from this is that it allows use of the bank account, A_t, as a deflator – i.e. the risk-free rate of return:

$$S_0 = \frac{E[S_t]}{1+\mu_t} = \frac{E^{\mathbb{Q}}[S_t]}{1+r_t}. \tag{10}$$

Throughout this thesis, only the bank account is used as numeraire. The corresponding probability measure will be referred to as *the risk-neutral measure*.

The existence of an equivalent martingale measure is closely related to the absence of arbitrage. Actually, it c an be shown that if there are no a rbitrage opportunities in the economy then an equivalent martingale measure exists. This is known as *the fundamental theorem of asset pricing*[14] (Cont and Tankov, 2004a).

Another important factor with respect to changing probability measure is *market completeness*. Cont and Tankov (2004a) state that markets are complete if the economy contains enough assets such that all contingent claims can be *replicated*. That is, one can create a portfolio that has the exact same properties as the contingent claim[15]. This implies the existence of a unique equivalent martingale measure.

---

[13] Formally, the numeraire can be any non-dividend paying asset, with price process B_t such that $B_t \geq 0$, for all t, i.e. strictly positive prices (Miltersen, 2005).

[14] The fundamental theorem of asset pricing states that "the market model defined by $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ and asset prices $(S_t)_{t\in[0,T]}$ is arbitrage-free if and only if there exist a probability measure $\mathbb{Q} \sim \mathbb{P}$ such that the discounted assets $(\hat{S}_t)_{t\in[0,T]}$ are martingales with respect to $\mathbb{Q}$". (Cont and Tankov, 2004a)

[15] A perfect hedge (or equivalently perfect replication) is an investment strategy that exactly offset any gains or losses for an existing investment. In a B&S-economy with an underlying asset S_t and savings account A_t, a perfect hedge is defined as a self-financing strategy (ϕ, θ) for a contingent claim H if

$$H = \phi_0 S_0 + \theta_0 A_0 + \int_0^t \phi_u dS_u + \int_0^t \theta_u dA_u \,.$$

Here, ϕ and θ represents the number of the underlying asset and the investment in the savings account, respectively. (Cont & Tankov, 2004a, and Miltersen 2005)

It should be emphasized that while most pricing models are arbitrage-free, not all are complete (Cont & Tankov, 2004a). In such cases, there are multiple equivalent martingale measures. Consequently, some contingent claims cannot be perfectly replicated. As will become evident later, this is the case when allowing for jumps. Cont and Tankov (2004a) argue that in the real world, markets are in general incomplete.

3.3.2 Intuition

There are two ways of changing measure (Neftci, 1996). First, the original shape of the distribution can be changed. Second, one can change the *mean* of the distribution, while leaving the variance unchanged. The latter is particularly used in pricing models for contingent claims. The intuition is that the original probability measure includes a premium for systematic risk. When changing measure from $\mathbb{P}$ to the new measure $\mathbb{Q}$, this premium is removed. Note that the sample paths of the stochastic processes are unchanged - it is only the probability weights that are changed in the transformation (Miltersen, 2005). The change of mean is illustrated in Exhibit 3.1.

Exhibit 3.1 - Change of Measure

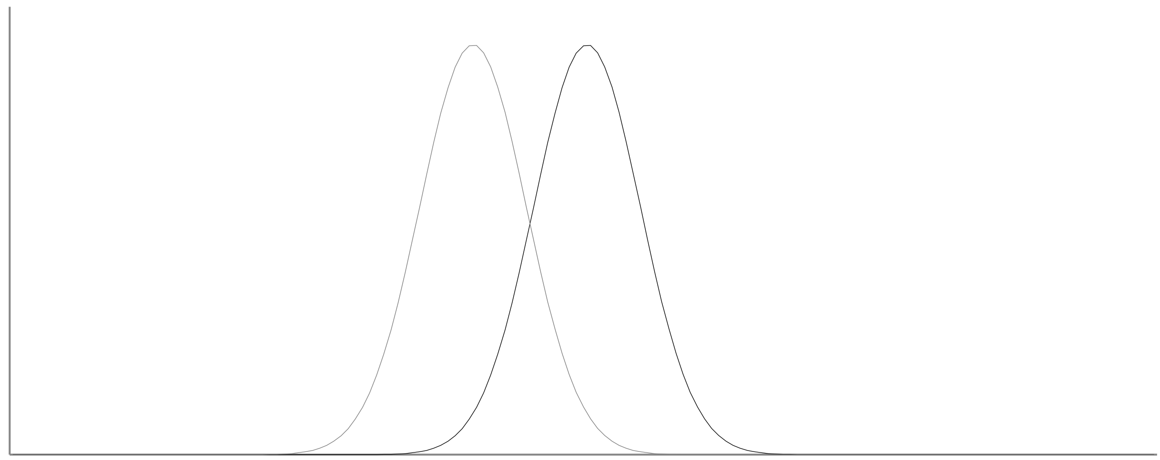

To be clear, changing probability measure is just a method of adjusting for systematic risk. The way this is done is dependent on the model in question. This will become evident in the subsequent sections.

Keep in mind that the probability measure $\mathbb{Q}$ is fictitious. On the other hand, the original probability measure $\mathbb{P}$ is the *real* or the *subjective* probability measure. That is, it reflects the market's belief about the future (Miltersen, 2005).

3.4 Adjusting for Systematic Risk in Jump-Diffusions

Consider now the general jump-diffusion model proposed in equation (4),

$$dS_t = \mu(S_t, t)dt + \sigma_1(S_t, t)dW_t + \; dJ_t \,. \tag{4}$$

Keep in mind that J_t is a compensated compound Poisson process. That is a compound Poisson process, Q_t, adjusted by subtracting its mean such that it is a martingale,

$$J_t = Q_t - \kappa\lambda t \,. \tag{3}$$

The parameters λ and κ represents the *jump intensity* and the *expected jump size*, respectively. It is assumed that Q_t, and hence J_t, has random jump sizes $\sigma_2(S_t, t)$. The jump-diffusion model can then be written as[16]

$$dS_t = [\mu(S_t, t) - \kappa\lambda]dt + \sigma_1(S_t, t)dW_t + \; \sigma_2(S_t, t)dN_t \,. \tag{5}$$

In order to use this for pricing purposes, the dynamics under $\mathbb{Q}$ must be derived.

3.4.1 A Special Case – Diffusions

A special case is when $\sigma_2(S_t, t) = 0$. As previously mentioned, these models are called diffusions. Here, the price dynamics is only dependent on one stochastic process, the Wiener process:

$$dS_t = \mu(S_t, t)dt + \sigma_1(S_t, t)dW_t \,. \tag{11}$$

Consequently, the price evolves continuously over time and replicating arguments can be used. Hence, the market is complete. In accordance with the above discussion, this implies the existence of a u nique equivalent martingale measure. The discounted price process, $S_t^* = \frac{S_t}{A_t}$, is now a martingale under $\mathbb{Q}$.

By using Itô's lemma[17] and Girsanov's[18] theorem, the risk-adjusted dynamics of S_t is obtained:

---

[16] See section 2.2.2.

[17] For an explanation of Itô's lemma, see for example Hull (2012), ch. 13.

$$dS_t = r_t S_t dt + \sigma_1(S_t, t)d\widetilde{W}_t \,. \tag{12}$$

Here, $\widetilde{W}_t$ is a Wiener process under the new probability measure. The relation between this and the Wiener process under $\mathbb{P}$, is given by

$$\widetilde{W}_t = W_t + \theta t, \tag{13}$$

where θ is the *market price of risk*. In order for S_t^* to be a martingale under the risk-neutral measure, the following equation must hold:

$$\theta = \frac{\mu(S_t,t) - r_t S_t}{\sigma_1(S_t,t)} \,. \tag{14}$$

The intuition is consistent with previous explanations. That is, systematic risk has to be adjusted for when changing from $\mathbb{P}$ to the risk-neutral measure.

As a result, the expected return under $\mathbb{Q}$ is equal to the risk-free rate of return. Notice that the volatility remains unchanged.

3.4.2 Jump-Diffusions

Now, consider the case when $\sigma_2(S_t, t)$ is stochastic[19] given S_t, i.e. a jump-diffusion. As noted, the model then has price jumps that may include systematic risk. Compared to the model above, this poses additional challenges for risk adjustment.

When allowing for discontinuous jumps, markets are no longer complete. As a consequence, there is no longer a unique equivalent martingale measure. Hence, in mathematical terms, there exist many theoretical prices. In order to choose the right price, economic arguments must be used[20] (Cont & Tankov, 2004).

Formally, there is defined a Wiener process (W_t) and a (compensated) compound Poisson process (J_t) on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ (Shreve, 2004). Assuming there is a

---

[18] Girsanov's theorem is discussed in Neftci (1996), ch. 14.

[19] $\sigma_2(S_t, t)$ can also be a constant.

[20] This is further discussed in section 3.5.

single filtration $\mathcal{F}(t), t \geq 0$ for both processes, they must be independent. Thus, the processes can be considered separately when changing measure (Shreve, 2004).

To obtain the risk-adjusted dynamics for S_t, the same principles as before are applied. That is, the discounted price process must be a martingale under $\mathbb{Q}$.

Now, consider the dynamics under the original probability measure, $\mathbb{P}$,

$$dS_t = \mu(S_t,t)dt + \sigma_1(S_t,t)dW_t + \, dJ_t \,. \tag{4}$$

By using the *Itô-Doeblin formula for jump processes*[21], the dynamics of the discounted price process are obtained:

$$dS_t^* = \left[-\frac{S_t r_t}{A_t} + \frac{\mu(S_t,t)}{A_t} - \frac{1}{2}\sigma_1(S_t,t)^2 0\right]dt + \frac{\sigma_1(S_t,t)}{A_t}dW_t + \frac{1}{A_t}dJ_t \,. \tag{15}$$

When simplifying terms,

$$dS_t^* = [\mu(S_t,t) - \, S_t r_t]\frac{1}{A_t}dt + \sigma_1(S_t,t)\frac{1}{A_t}dW_t + \frac{1}{A_t}dJ_t \,. \tag{16}$$

The next step is to find the risk-adjusted dynamics for the underlying by changing measure from $\mathbb{P}$ to $\mathbb{Q}$. As before, Girsanov's theorem is used to change measure for the Wiener process. For the (compound) Poisson process, this is conducted in a similar way[22]. Given that there is systematic risk inherent in the price jumps, the change of measure affect both the jump intensity and the jump size. That is, under risk-neutral measure one obtain a new intensity $\tilde{\lambda}$ such that

$$\tilde{\lambda} = \lambda + \lambda^c \,, \tag{17}$$

where λ is the intensity under $\mathbb{P}$, and λ^c is the change in intensity. Further, the distribution of the jump size will change from $\sigma_2(S_t,t)$ to a new distribution $\tilde{\sigma}_2(S_t,t)$. The expected jump size under $\mathbb{Q}$ is then equal to $\tilde{\kappa}$ such that

$$\tilde{\kappa} = \kappa + \kappa^c \,. \tag{18}$$

---

[21] See e.g. Cont and Tankov (2004a), section 8.3.2, for a presentation of the Itô-Doeblin formula for jump processes.

[22] For a discussion of how to change measure for a compound Poisson process, see e.g. Shreve (2004).

Here, κ is the expected jump size under $\mathbb{P}$, while κ^c represents the change. Despite the above changes, J_t is a compound Poisson process under $\mathbb{Q}$ (Shreve, 2004).

When going from $\mathbb{P}$ to $\mathbb{Q}$, J_t is no longer a martingale. However, an adjustment of the mean can solve this. That is, one can define

$$\tilde{J}_t = J_t + \left(\lambda\kappa - \tilde{\lambda}\tilde{\kappa}\right)t \tag{19}$$

such that $\tilde{J}_t$ is a martingale under $\mathbb{Q}$. The risk-neutral dynamics for S_t^* is then given by

$$dS_t^* = [\mu(S_t, t) - S_t r_t]\frac{1}{A_t}dt + \sigma_1(S_t, t)\frac{1}{A_t}d(\widetilde{W}_t - \theta t) + \frac{1}{A_t}d\left(\tilde{J}_t - \left(\lambda\kappa - \tilde{\lambda}\tilde{\kappa}\right)t\right)$$

$$= \left[\mu(S_t, t) - S_t r_t - \sigma_1(S_t, t)\theta - \lambda\kappa + \tilde{\lambda}\tilde{\kappa}\right]\frac{1}{A_t}dt + \sigma_1(S_t, t)\frac{1}{A_t}d\widetilde{W}_t + d\tilde{J}_t, \tag{20}$$

where $\widetilde{W}_t$ and $\tilde{J}_t$ are independent of each other (Shreve, 2004). As stated, S_t^* must be a martingale under the risk-neutral measure. Hence, the dt-term in (20) must equal 0. That is,

$$\mu(S_t, t) - S_t r_t - \sigma_1(S_t, t)\theta - \lambda\kappa + \tilde{\lambda}\tilde{\kappa} = 0\,. \tag{21}$$

Rearranging terms, the market price of risk equation is given by

$$\mu(S_t, t) - S_t r_t = \sigma_1(S_t\ t)\theta + (\lambda\kappa - \tilde{\lambda}\tilde{\kappa}). \tag{22}$$

This equation includes important economic insight. The left hand side represents the total risk premium for holding the underlying asset. Thus, the right hand side illustrates the decomposition of this premium. Here, the first term can be interpreted as the total compensation for diffusion risk, i.e. risk related to the Wiener process. Then θ yields the diffusion premium per unit diffusive risk, $\sigma_1(S_t\ t)$.

The second term represents the market price of jump risk. In general, this contains premiums for both the jump size and the rate of occurrence. When jump size is not priced, $\kappa = \tilde{\kappa}$ and the market price of jump risk reduces to a premium for the intensity. On the other hand, if jump size is priced but timing is not, $\lambda = \tilde{\lambda}$ and only a compensation for size is included.

There are now three unknowns, θ, $\tilde{\lambda}$ and, $\tilde{\kappa}$, and one equation. This implies a multitude of possible prices, and thus the need for economic reasoning to obtain a unique price. In practice, market calibration[23] is commonly used to estimate the parameters (Shreve, 2004).

In accordance to the above discussion, the risk-neutral dynamics of S_t is then given by

$$dS_t = S_t r_t dt + \sigma_1(S_t, t)d\widetilde{W}_t + d\tilde{J}_t \tag{23}$$

or equivalently

$$dS_t = \left[r_t S_t - \tilde{\kappa}\tilde{\lambda}\right]dt + \sigma_1(S_t, t)d\widetilde{W}_t + \tilde{\sigma}_2(S_t, t)d\widetilde{N}_t \,. \tag{24}$$

Here,

> r_t is the risk-free *rate of return*,
> $\sigma_1(S_t, t)$ is as before,
> $\widetilde{W}_t$ is a *standard Wiener process*,
> $\tilde{\sigma}_2(S_t, t)$ is the risk-adjusted jump-size, where $E[\tilde{\sigma}_2(S_t, t)] \equiv \tilde{\kappa}$,
> $\widetilde{N}_t$ is a Poisson process with intensity $\tilde{\lambda}$.

Note that the price dynamics under $\mathbb{Q}$ has an expected rate of return equal to the risk-free rate. This is consistent with risk-neutral pricing arguments. Observe also that if jump risk is idiosyncratic, the jump intensities and the jump sizes are equal in both probability measures. Then, the market price of risk equation is the same as in the situation with no jumps.

A special case of (4) is when the jump-sizes $\sigma_2(S_t, t)$ are deterministic. Then $\kappa = \tilde{\kappa} = \sigma_2(S_t, t)$, and the market price of risk equation reduces to

$$\theta = \frac{\mu(S_t, t) - S_t r_t - (\lambda - \tilde{\lambda})\sigma_2(S_t\ t)}{\sigma_1(S_t\ t)} \,. \tag{25}$$

---

[23] This is further discussed in the following.

3.5 Hedging and Pricing under Jump-diffusions

Section 3.3.1 established that there is a one-to-one correspondence between arbitrage-free (payoff-replication) pricing and equivalent martingale measures. When markets are complete, as in the diffusion-case, there exist a unique pricing measure[24]. In other words, it is only one arbitrage-free way to price a contingent claim; the value equals the cost to replicate it (Cont & Tankov, 2004a).

When markets are incomplete, as for jump-diffusions, there are in general infinitely many pricing measures (Cont & Tankov, 2004a). Essentially this means that there is no hedging strategy that perfectly replicates the contingent claim in question. Hence, when setting up a hedging portfolio there is risk that cannot be eliminated. From an economic point of view, the value of the claim should then equal the cost of the hedge, plus a premium for the unhedgeable risk. However, since there are different ways of measuring risk, there are also different ways of hedging. This implies that there exist a multitude of possible prices, dependent on the risk aversion of the investors.

As noted in Xu (2005) and Cont & Tankov (2004a), there are two major approaches for hedging and pricing contingent claims when markets are incomplete. First, one can use so-called *utility-based methods*. Here, one incorporates the investors' attitude to risk via utility functions, using the underlying asset(s) to construct a hedging portfolio. Since it is difficult to determine investors' preferences for risk, this approach is difficult to use in practice (Cont & Tankov, 2004a).

The second approach is called (implied) *risk-neutral modeling*, and will be the focus in this thesis. Here, one obtains the risk-neutral dynamics for the underlying asset directly by choosing an equivalent measure $\mathbb{Q}$ that represents qualitative properties of the asset's price (Cont & Tankov, 2004a). More specific, one assumes that the underlying follows a given risk-neutral model with certain parameters. These parameters are extracted from market prices[25] for liquid contingent claims (e.g. plain vanilla options), and the model is then used for hedging and pricing of other exotic or illiquid derivatives. The intuition is that the market

[24] See section 3.3.1.

[25] This is also known as model calibration, and is further discussed in chapter 5.

chooses the right pricing measure, reflecting the investors' risk aversion[26]. It is important to note that this approach may include options (and other contingent claims) to construct the hedging portfolio. This requires well-functioning markets for such contracts. If this is not the case, reasonable hedging portfolios may not exist (Xu, 2005).

One can relate the above discussion to the general jump-diffusion in section 3.4. W hen markets are incomplete (as is the case for jump-diffusions), the risk-neutral dynamics has no direct relation to the real pricing measure $\mathbb{P}$. More specific, the jump intensity ($\tilde{\lambda}$) and expected jump size ($\tilde{\kappa}$) under $\mathbb{Q}$ are unknown. These 'free parameters' can then be estimated using risk-neutral modeling, as described above. This is further discussed in chapter 4 and 5.

3.6 Summary

This chapter and the previous one laid out the theoretical groundwork for jump-diffusion models. By using a *Wiener process* for normal events and a (compound) *Poisson process* for extreme events, the model can represent all types of disturbances that may affect financial markets (Nefcti, 1996).

Much emphasis has been attributed to the discussion of *systematic risk* and *equivalent martingale measures*. In general, since investors are assumed to be risk-averse, asset prices depend crucially on their level of systematic risk. That is, investors demand a compensation for taking on risk that is correlated to the risk in the return from the *market portfolio*. In our model, this includes both diffusive risk and jump risk.

Since it is difficult to identify the investors' risk-preferences, the job of pricing assets may seem impossible. However, it turns out that there exists a pricing method that maneuvers around the problem, without the concern of risk-aversion. In the financial literature this is known as the *method of equivalent martingale measures.* Essentially, this is a pricing method that allows for treating a risky asset as if it was risk-free.

[26] More precisely, this reflects the risk aversion of the average investor. The individual investors will in general have different attitudes to risk.

The existences of martingale measures are nearly related to the absence of arbitrage. According to *the fundamental theorem of asset pricing* such a measure exist if the market is arbitrage-free. Further, if the market is *complete* there exist a unique pricing measure.

On the other hand, if the market is said to be *incomplete*, there exist a multitude of possible prices corresponding to different pricing measures. This stems from the fact that it is impossible to perfectly replicate the asset in question. In these situations one must use economic arguments, such as those presented in section 3.5, to choose the right measure.

4. Pricing European Options on S&P 500

With the theoretical aspects of the framework for derivatives pricing in place, the remainder of this thesis is devoted to its application. In particular, the focus will be directed towards the implications of jumps for the pricing of European options with S&P 500 as the underlying asset. S&P 500 is a diverse stock market index comprising 500 large companies in leading industries in the U.S., and is by many regarded as the best representation of the market (Standard & Poor's, 2013). For the purposes of this thesis, it is important to note that it is a price index, i.e. it does not account for dividends.

This chapter begins with a discussion of the traditional assumption that stock returns follow a normal distribution, followed by an empirical examination of the distributional properties of historical S&P 500 returns. This leads to the proposition of a jump-diffusion model for the dynamics of this index. The chapter concludes by presenting a closed form solution for the price of a European call on S&P500.

4.1 Traditional Assumptions about Stock Returns

The distributional properties of asset prices, and thus their rate of return, have important implications for investment decisions. Numerous models and applications of finance rely upon the assumption that stock prices follow a lognormal distribution (Hull, 2012). Consequently, the distribution of logarithmic stock returns is assumed to follow a normal distribution. This is also one of the underlying assumptions of the Black-Scholes (B&S) option-pricing model[27] (Black and Scholes, 1973).

The B&S model is a widely used model for pricing European options. It is based on the assumption that stock prices evolve according to a *geometric Brownian motion*[28]. That is, the underlying stock is assumed to exhibit the dynamics given by the SDE

$$dS_t = \mu S_t dt + \sigma_1 S_t dW_t \,. \tag{26}$$

[27] The model is not derived here, as it assumed that it is familiar for the reader. See Black and Scholes (1973).

[28] Note that a Brownian motion is equivalent to a Wiener process.

In this specific diffusion model, μ and σ_1 are assumed to be constants, representing the expected rate of return and the standard deviation of the return, respectively. This is consistent with previous explanations. The solution to the SDE is given by

$$S_t = S_0 e^{\left(\mu - \frac{1}{2}\sigma_1{}^2\right)t + \sigma_1 dW_t} \quad . \tag{27}$$

Since S_t is a lognormal process, it follows that the distribution of logarithmic stock returns are normal. That is,

$$R_t = \ln\left(\frac{S_t}{S_0}\right) \sim N\left[\left(\mu - \frac{1}{2}\sigma_1{}^2\right)t, \sigma_1{}^2 t\right], \tag{28}$$

where R_t denotes the logarithmic return (Hull, 2012). From here on, this will be referred to as the *log-return*.

4.2 Distributional Properties of S&P 500 log-returns

To determine whether the empirical distribution is consistent with the aforementioned assumptions, the distributional properties of log-returns from the S&P 500 index are investigated. Obviously, it is essential that the assumptions leading to a pricing model is consistent with the real dynamics.

Two approaches are applied to assess whether the log-returns of the S&P 500 are normal. First, descriptive statistics and graphical representation of the time series are used. Keller (2009) advocates the use of such methods. Since this approach requires subjective judgment of the distributional properties, formal statistical tests are also conducted in order to determine normality. For this, the commonly used *Anderson-Darling* test is applied to the data sample. All statistical tests are conducted in the software *MiniTab 16*.

4.2.1 Data

Data was downloaded from Yahoo! Finance on May 21st. Closing prices for the S&P 500 index (ticker: ^GSPC) is collected for the period starting in May 1990 and ending in May 2013 (finance.yahoo.com, 2013a). The frequency of the obtained data is daily, weekly and monthly, respectively. From this, log-returns are computed for all frequencies. Only results based on 5795 daily observations are presented in the text. However, identical analyses are conducted to data with weekly and monthly frequencies. The results of these are presented in Appendix A.

4.2.2 Graphical Analysis

A presentation of the empirical distribution of daily log-returns, along with its descriptive statistics is provided in Exhibit 4.1. F or illustrational purposes, the theoretical normal distribution is plotted in the same diagram.

Exhibit 4 1 - Distribution of Daily log-returns for S&P 500, 1990-2013

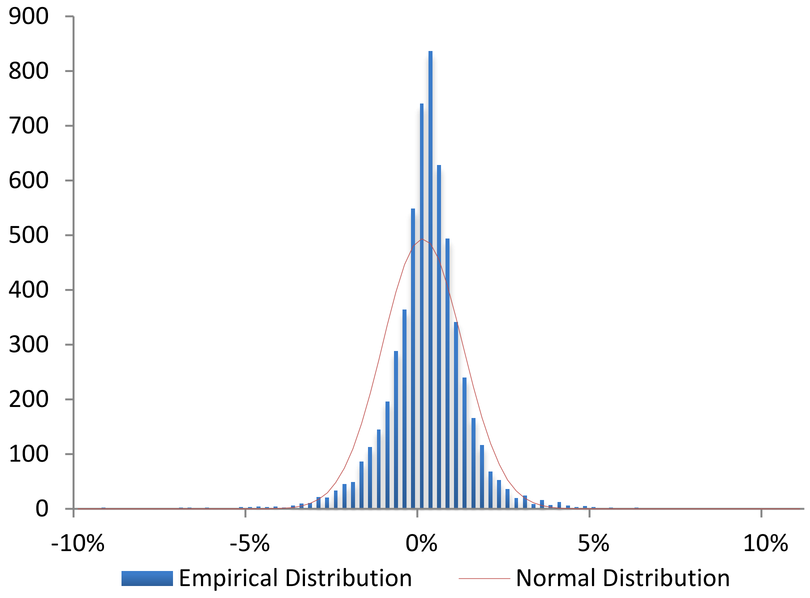

Descriptive statistics	
Mean	0,03 %
Std.dev	1,17 %
Skewness	-0,2333
Kurtosis	8,518
Median	0,05 %
Max	10,96 %
Min	-9,47 %
Q1	-0,48 %
Q3	0,58 %
Observations	5795

By visual inspection of the empirical distribution of daily log-returns, it is evident that it does not perfectly resemble the normal distribution. In particular, it displays *leptokurtic* features. That is, the peak is higher and the tails are heavier than that of the normal distribution. The last point is evident from the frequency of extreme observations. These are seen by looking closely at the tails of the histogram. It is noted that the data was checked for obvious measurement errors, for which none was identified.

The descriptive statistics suggest that the distribution of log-returns is close to symmetric. A *skewness* of -0.23 indicates that it is slightly skewed to the left. Further, *kurtosis* equal to 8.518 confirms the views posted above. This implies an *excess kurtosis* of 5.518. The latter measure provides a direct comparison with the normal distribution, which has a kurtosis equal to 3.

4.2.3 A Formal Normality Test

A powerful and frequently used method to detect if a given data sample departs from normality is the *Anderson-Darling* test. It is considered beyond the scope of this paper to explain how this test is conducted. Instead, a brief description of this method for assessing the distributional properties of S&P500 log-returns is provided.

The test statistic, denoted *AD*, is the squared difference between the empirical and fitted *cumulative distribution functions* (CDFs) (D'Agostino & Stephens, 1986). Here, each observation is weighted in such a way that the tails of the distribution is accentuated. In order to determine if the sample follow a particular distribution, the test statistic is compared to critical values of the fitted distribution. Smaller values of AD thus indicate a better fit to the theoretical probability distribution. The null hypothesis is that the data is normally distributed. If the computed p-value is less than the chosen significance level, this is rejected.

It should be noted that p-values and critical values for the Anderson-Darling test in most cases only is approximated since it does not have a usable distribution (D'Agostino & Stephens, 1986). The results of the test are displayed in Exhibit 4.2.

Exhibit 4.2 - Anderson-Darling Test for Daily log-returns

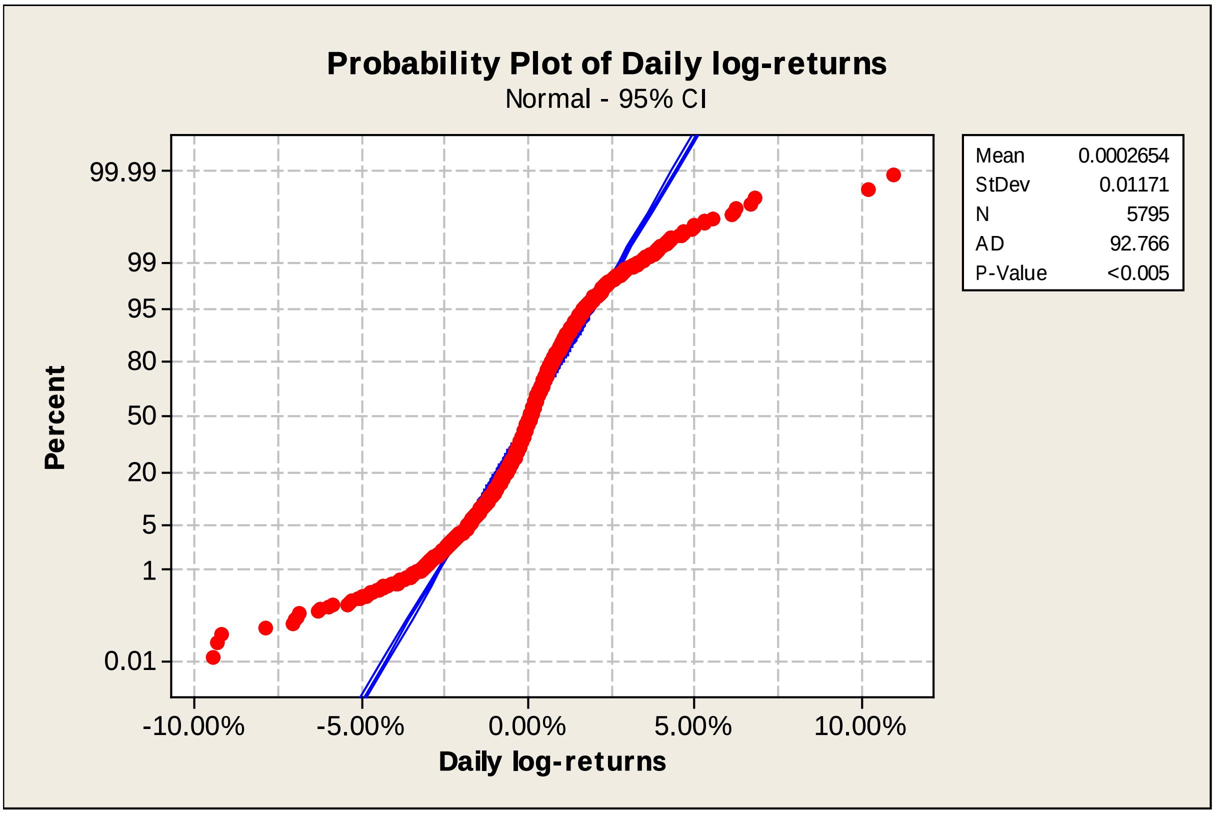

In the probability plot, the empirical CDF and the CDF of the normal distribution are compared. These are illustrated by the red and blue lines, respectively. Note that the y-axis is scaled in such a way that the latter is a straight line. Non-normal data are thus indicated by large deviations from the straight blue line. Both ends of the red line exhibit a flatter slope than the normal CDF. This implies that the empirical CDF exhibit fatter tails. In other words, there are more extreme observations of daily log-returns from S&P 500 than a fitted normal distribution would produce.

MiniTab returns a value of the AD test statistic equal to 92.766. The corresponding p-value is less than 0.005. Hence, the assumption of normality is rejected for daily log-returns from S&P 500 at a significance level of less than 0.5%.

In order to verify the conclusion produced by the Anderson-Darling test, two other normality tests are applied to the data sample. These are the *Kolmogorov-Smirnov* test and the *Ryan-Joiner* test[29]. These methods yield the same conclusion as above. Results from these tests are presented and discussed in Appendix A.

4.2.4 Conclusion on Distributional Properties

Both graphical and formal analysis of the distributional properties of empirical data suggests that daily S&P 500 log-returns do not follow a normal distribution. Rather, it is evident from the analyses that that the empirical distribution exhibit leptokurtic features, i.e. it has fatter tails and a sharper peak. Consequently, it is concluded that daily log-returns are non-normal.

This conclusion is in direct conflict with traditional assumptions concerning the dynamics of stock prices. In particular, geometric Brownian motion fails to capture the distributional properties of daily log-returns from S&P 500. Hence, to model these features, other stochastic processes are needed (Benth, 2004).

Note that this conclusion is in accordance with the arguments of Benth (2004). He states that log-returns are far from normal for short time horizons. Given that daily log-returns are iid, however, the central limit theorem implies that weekly and monthly log-returns will move towards a normal distribution (Benth, 2004). This is due to the additive property of log-

[29] Note that the Ryan-Joiner test is also called the *Shapiro-Wilk* test.

returns. Still, normality tests for weekly and monthly log-returns presented in Appendix A suggest otherwise. The question whether daily log-returns are independent can be discussed, as empirical evidence also suggests long-range dependency of returns. However, the issue of autocorrelation is ignored as it is of greater importance to models not considered in this thesis[30].

4.3 A Suggested Pricing Model for S&P 500

Based on t he above discussion, alternatives to geometric Brownian motion should be considered for the S&P 500 index. There exists an extensive literature on such models (Hull, 2012). Due to the limited scope of this thesis only one model is applied. As in Bates (1991), it is suggested that the price dynamics of the S&P 500 index is given by

$$\frac{dS_t}{S_t} = [\mu - \kappa\lambda - d_t]dt + \sigma dW_t + (Y_t - 1)dN_t. \tag{29}$$

Here,

μ is the cum-dividend *expected rate of return* on the asset,
d_t is the dividend yield,
σ is the diffusion coefficient conditional on no jump,
W_t is a standard Wiener process,
$(Y_t - 1)$ is the percentage jump given a Poisson event, where $\ln Y_t$ is normally distributed: $\ln Y_t \sim N\,(\alpha, \delta^2)$,
$E[Y_t - 1] \equiv \kappa = e^{\alpha + \frac{\delta^2}{2}} - 1$,
$Var[Y_t - 1] \equiv \upsilon^2 = (e^{\delta^2} - 1)e^{2\alpha + \delta^2}$,
N_t is a Poisson process with intensity λ .

The model is often referred to as *Merton's jump-diffusion model*. Most of the time, it is identical to the B&S-model (i.e. geometric Brownian motion). However, at an average of λ times per year, S_t jumps discretely by $(Y_t - 1)$ percent. It should be noted when $\lambda = 0$, the process has no jumps. This is also the case when both α and δ equals 0.

[30] The time-dependency of returns is of more interest to models that account for stochastic and/or mean-reverting volatility.

Bates (1991) shows that for constant $\mu - d_t$, the variance, skewness and kurtosis for $\ln(\frac{S_{t+T}}{S_t})$ are given by

$$Variance = v^2 T = \{\sigma^2 + \lambda[\alpha^2 + \delta^2]\}T \tag{30}$$

$$Skewness = \lambda\alpha[\alpha^2 + 3\delta^2]\frac{T^{\frac{-1}{2}}}{v^3} \tag{31}$$

$$Kurtosis = 3 + \lambda[\alpha^4 + 6\alpha^2\delta^2 + 3\delta^4]\frac{T^{-1}}{v^4} \tag{32}$$

As the holding period (T) increases, the distribution will converge towards the normal distribution (Bates, 1991).

Consistent with the notation in previous sections, Merton's model can also be written as

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t + J_t\,, \tag{33}$$

where J_t is a compensated compound Poisson process, $J_t = \sum_{i=1}^{N_t}(Y_i - 1) - \kappa\lambda t$.

4.3.1 The Risk-Neutral Measure

Merton's original paper from 1976 assumes that jump risk is unsystematic. Hence, neither the jump size nor the intensity is priced, and the jump distribution is the same under the real and the risk-neutral measure.

When considering a large stock index, such as S&P 500, Merton's assumptions regarding jump risk, seems to be too simple (Bates, 1991)[31]. Because of this, it is allowed for systematic jumps. Hence, in this framework both jump-size and intensity may change when going from $\mathbb{P}$ to $\mathbb{Q}$. The theory of changing measure for jump-diffusions was presented in section 3.4. This will now be applied to our suggested price model.

Consider the general jump-diffusion,

[31] When combining many stocks in a large portfolio, unsystematic risk will be diversified away. Hence, large indexes such as S&P 500 will contain little unsystematic risk. (Berk and DeMarzo, 2011)

$$dS_t = \mu(S_t, t)dt + \sigma_1(S_t, t)dW_t + dJ_t \,. \tag{4}$$

To be consistent with Merton's notation, this is alternatively written as

$$dS_t = [\mu(S_t, t) - \kappa\lambda]dt + \sigma_1(S_t, t)dW_t + \sigma_2(S_t, t)dN_t \,. \tag{5}$$

Given that there is systematic risk inherent in the jumps, the change of measure affect both the jump-size and the intensity. As stated earlier, the risk-neutral dynamics is then given by

$$dS_t = \left[r_t S_t - \tilde{\kappa}\tilde{\lambda}\right]dt + \sigma_1(S_t, t)d\widetilde{W}_t + \tilde{\sigma}_2(S_t, t)d\widetilde{N}_t \,. \tag{24}$$

Now, take a second look at Merton's jump-diffusion model (29). This can be considered as a special case of the more general jump-diffusion (5). Specifically, by comparing terms, one gets that

$$\mu(S_t, t) = \mu S_t$$

$$\sigma_1(S_t, t) = \sigma S_t$$

$$\sigma_2(S_t, t) = (Y_t - 1)S_t \,.$$

The risk-adjustment is then straightforward. One should, however, be aware of one particular point. When going from $\mathbb{P}$ to $\mathbb{Q}$, the density function of the jump-sizes $\sigma_2(S_t, t)$ may change. This change is dependent on the original distribution.

In our model, the jump-sizes under $\mathbb{P}$ are denoted $(Y_t - 1)$. Here $\ln(Y_t)$ is normally distributed with mean α and variance δ^2. Then, as shown in Gerber and Shiu (1994), the distribution under $\mathbb{Q}$ also is normal but with different parameters[32]. More specific, the result is changed mean and unchanged variance[33] (Gerber and Shiu, 1994). That is, when changing measure, $(Y_t - 1)$ will simply turn to a new variable $(\tilde{Y}_t - 1)$, where $\ln(\tilde{Y}_t)$ is normally distributed with new mean $\tilde{\alpha}$ and unchanged variance δ^2. This is consistent with the discussion in section 3.3.2. As before, the economic intuition is that the variable contains a

[32] Here, a so-called *Esscher transform* is used. Essentially this method takes a density function $f(x)$ and transform it to a new probability function $f(x; h)$ with parameter h. (Gerber and Shiu, 1994)

[33] When changing measure, a normally distributed variable $X \sim N(\alpha, \delta^2)$ will turn to a new normally distributed variable $\tilde{X} \sim N(\alpha + h\delta^2, \delta^2)$. The parameter h must be determined such that the new measure is an equivalent martingale measure (Gerber and Shiu, 1994).

premium for systematic risk. When going from $\mathbb{P}$ to $\mathbb{Q}$ this premium is removed, leaving the variance unchanged.

From the above discussion it is clear that the risk-neutral dynamics for our suggested model (29) is given by (Bates, 1991)

$$\frac{dS_t}{S_t} = \left[r_t - \tilde{\kappa}\tilde{\lambda} - d_t\right]dt + \sigma d\widetilde{W}_t + \left(\tilde{Y}_t - 1\right)d\widetilde{N}_t. \tag{34}$$

Here,

r_t is risk-free rate of return,
d_t, σ, and δ are as before,
$\widetilde{W}_t$ is a standard Wiener process,
$(\tilde{Y}_t - 1)$ is the percentage jump given a Poisson event, where $\ln \tilde{Y}_t$ is normally distributed: $\ln \tilde{Y}_t \sim N(\tilde{\alpha}, \delta^2)$,
$E[\tilde{Y}_t - 1] \equiv \tilde{\kappa} = e^{\tilde{\alpha}+\frac{\delta^2}{2}} - 1$,
$Var[\tilde{Y}_t - 1] \equiv \tilde{v}^2 = (e^{\delta^2} - 1)e^{2\tilde{\alpha}+\delta^2}$,
$\widetilde{N}_t$ is a Poisson process with intensity λ.

It should be noted that when all jump-risk is unsystematic, $\lambda = \tilde{\lambda}$ and $\alpha = \tilde{\alpha}$. Then the jumps are equal under both measures, and the model coincides with Merton's jump-diffusion. However, if both jump intensity and jump-size contains systematic risk, $\lambda \neq \tilde{\lambda}$ and $\alpha \neq \tilde{\alpha}$.

4.3.2 Closed-Form Solution

When jump risk is idiosyncratic, Merton (1976) provides a closed-form formula for the price of a European call option on an underlying following (29). It is straightforward to extend this formula to our model[34].

Consider a European call (C) with time to maturity T and strike K. The risk-free rate of return (r) is assumed to be constant. If its underlying (S) follows (29) with systematic jump risk, the price of C is given by:

[34] This is simply done by replacing λ and α in Merton's formula, with $\tilde{\lambda}$ and $\tilde{\alpha}$.

$$C(S,T,K) = \sum_{n=0}^{\infty} \frac{e^{-\hat{\lambda}}\left(\hat{\lambda}T\right)^n}{n!} C_{BS}(S,T,K,\sigma_n^2, r_n, d_t)\,, \tag{35}$$

where

$$\sigma_n^2 = \sigma^2 + \frac{n}{T}\delta^2\,,$$

$$\hat{\lambda} = \tilde{\lambda}\left(1+\tilde{k}\right),$$

$$r_n = r - \tilde{\lambda}\tilde{k} + \frac{n}{T}\ln(1+\tilde{\kappa})\,.$$

C_{BS} represents the well-known B&S-formula (Black and Scholes, 1973)

$$C_{BS}(S,T,K,\sigma_n^2, r_n, d_t) = S_t\Phi(d_1) - Ke^{-rT}\Phi(d_2)\,, \tag{36}$$

where

$$\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{\frac{s^2}{2}} ds\,,$$

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r_n - d_t + \frac{\sigma_n^2}{2}\right)T}{\sigma_n\sqrt{T}}\,,$$

$$d_2 = d_1 - \sigma_n\sqrt{T}\,.$$

4.3.3 The Hedging Portfolio

From a financial point of view, pricing of options using (35) and (36) makes no sense without a replicating strategy. Since the model exhibits discontinuous jumps the market is incomplete. This means that there is no such thing as a perfect hedge, and there exist a multitude of pricing measures. Hence, in order to price and hedge under these circumstances one must use techniques such as those presented in section 3.4.

It is considered beyond the scope of this thesis to present a replicating strategy for (29). Note, however, that there exists a wide literature on hedging of such models. This includes Rebonato (2004), Cheang and Chiarella (2011), etc. As a general result, the hedging portfolio includes other contingent claims.

5. Calibration

In order to answer whether jumps are relevant for the pricing of European options on S&P 500, the suggested model is calibrated to historical option prices. This chapter subsequently describes the applied calibration algorithm, the data used in the process, and discusses the implementation of the procedure.

5.1 Method

Before presenting the applied calibration approach, a general discussion about market calibration is provided.

5.1.1 General Remarks

Calibration is a common way of determining unknown parameters in financial models (Hull, 2012). In general, it involves identifying the model parameters that leads to the best fit to market data. Several calibration methods are proposed in the literature on option pricing models, and the best approach is dependent on the problem at hand[35] (Gilli and Schumann, 2010). The success of the calibration is critically dependent on an appropriate goodness-of-fit measure, along with a robust method to solve the resulting optimization problem.

The choice of goodness-of-fit measure, or objective function, should ensure that the calibration process determines the value of the desired parameters in a reasonable fashion. Considerations to make are whether the objective function should comprise of absolute or relative price differences, whether to use the absolute value or square these differences, and if a weighting scheme should be applied (Gilli and Schumann, 2010). Standard methods to solve the resulting optimization problem are based on the derivatives of the objective function, such as the *Newton-Raphson* method. In the case of an *ill-posed* problem, i.e. a problem with no unique or stable solution, however, such methods may fail[36] (Cont and Tankov, 2004b).

[35] For discussions of different calibration routines, see, e.g., Cont and Tankov (2004b), Hull (2012) or Gilli and Schumann (2010).

[36] See Cont and Tankov (2004b) for a calibration approach that mitigates this problem.

It should be noted that the existence of a unique solution to the calibration problem is unlikely if model and market prices are restricted to be equal (Cont and Tankov, 2004b). This is due to the fact that market prices are quoted as bid-ask intervals with predefined tick sizes, and are thus not necessarily exact. Hence, the problem is often reduced to a minimization problem.

5.1.2 Choice of Calibration Approach

Here, the objective of the calibration procedure is to obtain precise estimates of the diffusion volatility (σ), the risk-adjusted jump intensity ($\tilde{\lambda}$), the risk-adjusted mean jump-size ($\tilde{\kappa}$) and the volatility of the risk-adjusted jump-size ($\tilde{\upsilon}$). Since $\tilde{\kappa}$ and $\tilde{\upsilon}$ follows directly from $\tilde{\alpha}$ and δ, it does not matter which two of these parameters that are estimated directly through calibration, and which two that are subsequently determined indirectly. In this case, the procedure is easier to implement when estimates of $\tilde{\kappa}$ and δ are derived, while $\tilde{\upsilon}$ and $\tilde{\alpha}$ are indirectly computed. Estimates of implied dividend yield (d) are not obtained here, as it is, in general, uncommon to calibrate parameters for which good estimates are easily observable.

In line with the above discussions, the applied calibration strategy is given by:

$$\min_{\sigma,\tilde{\lambda},\tilde{\kappa},\delta} f(\sigma,\tilde{\lambda},\tilde{\kappa},\delta) = \min_{\sigma,\tilde{\lambda},\tilde{\kappa},\delta} \sum_{i=1}^{n} \left(C_i^{Obs} - C_i^{*}\right)^2, \tag{37}$$

where

n is the number of *calibrating instruments*,
C_i^{Obs} is the observed market price of calibrating instrument i,
C_i^{*} is the price obtained by the suggested model for instrument i.

(37) states that the model is fitted to market data by minimizing the sum of squared absolute differences between observed market prices and model prices with respect to the unknown parameters $\sigma, \tilde{\lambda}, \tilde{\kappa}, \delta$. Since only one daily estimate of these parameters are obtained by

choosing calibrating instruments that satisfies a set of required characteristics, it is appropriate to use absolute price differences in the objective function[37].

Obviously, the calibrating instruments are traded European type call options with the S&P500 index as the underlying asset. Some remarks regarding the choice of these should be emphasized. First of all, the number of calibrating instruments cannot be smaller than the number of unknown parameters in order for the strategy (37) to work. In addition, it is of critical importance to the results that the chosen calibrating instruments and the instruments being valued are identical, i.e. that they have the same exercise price and time to maturity, respectively (Hull, 2012). The characteristics of the chosen options are thoroughly explained in section 5.2.

Despite its absence in (37), weighting schemes could have been included (Cont and Tankov, 2004b) . In particular, it was considered to construct relative weights that accounts for the liquidity of the given option. It is generally accepted that the more liquid the asset, the more confidence is attributed to its price. A possible choice of weights reflecting this is the inverse of the bid-ask spread of the specific option, assuming that the bid-ask spread is inversely related to liquidity. However, the bid-ask spread is not always possible to obtain. An alternative is to exclusively choose options with relatively high trading volumes. This is the chosen strategy for the calibration procedure in this thesis.

5.2 Data

Historical end-of-day prices for SPX options were obtained from Historical Option Data (historicaloptiondata.com, 2013). These are options of European type with the S&P 500 index as the underlying (cboe.com, 2013a). Exhibit 5.1 presents the product specifications.

When determining the relevance of jumps, it seems reasonable to choose a sample period that includes abrupt price changes. An obvious choice is the global financial crisis of 2008. From September 2008 until the end of the year, a total of 32 daily log-returns exceeded 3

[37] It can in some cases be more suitable to use relative price differences in the objective function. A rationale for choosing relative price differences rather than absolute price differences is to achieve a more equal weighting of in-the-money and out-of-the-money options. This is based on an assumption that relative price differences exhibit less variation than absolute price differences in this case. Due to the choice of calibrating instruments, this is not relevant here.

Exhibit 5.1 - Fact Sheet for S&P 500 Index Options

Symbol: SPX

Underlying: The Standard & Poor's 500 Index is a capitalization-weighted index of 500 stocks from a broad range of industries. The component stocks are weighted according to the total market value of their outstanding shares. The impact of a component's price change is proportional to the issue's total market value, which is the share price times the number of shares outstanding. These are summed for all 500 stocks and divided by a predetermined base value. The base value for the S&P 500 Index is adjusted to reflect changes in capitalization resulting from mergers, acquisitions, stock rights, substitutions, etc.

Multiplier: $100.

Premium Quote: Stated in decimals. One point equals $100. Minimum tick for options trading below 3.00 is 0.05 ($5.00) and for all other series, 0.10 ($10.00).

Strike Prices: In-,at- and out-of-the-money strike prices are initially listed. New series are generally added when the underlying trades through the highest or lowest strike price available.

Strike Price Intervals: Five points. 25-point intervals for far months.

Expiration Months: Up to twelve (12) near-term months. In addition, the Exchange may list up to ten (10) SPX LEAPS® expiration months that expire from 12 to 60 months from the date of issuance.

Expiration Date: Saturday following the third Friday of the expiration month.

Exercise Style: European - SPX options generally may be exercised only on the last business day before expiration.

Last Trading Day: Trading in SPX options will ordinarily cease on the business day (usually a Thursday) preceding the day on which the exercise-settlement value is calculated.

Settlement Value: Exercise will result in delivery of cash on the business day following expiration. The exercise-settlement value, SET, is calculated using the opening sales price in the primary market of each component security on the last business day (usually a Friday) before the expiration date. The exercise-settlement amount is equal to the difference between the exercise-settlement value and the exercise price of the option, multiplied by $100.

Position and Exercise Limits: No position and exercise limits are in effect. Each member (other than a market-maker) or member organization that maintains an end of day aggregate position in excess of 100,000 contracts in SPX and Mini-SPX (10 Mini-SPX options equal 1 SPX full value contract) for its proprietary account or for the account of a customer, shall report certain information to the Department of Market Regulation. The member must report information as to whether such position is hedged and, if so, a description of the hedge employed e.g. stock portfolio current market value, other stock index option positions, stock index futures positions, options on stock index futures; and for customer accounts, provide the account name, account number and tax ID or social security number. A report must be filed when an account initially meets the aforementioned applicable threshold. Thereafter, a report must be filed for each incremental increase of 25,000 contracts. Reductions in an options position do not need to be reported. However, any significant change to the hedge must be reported.

Margin: Purchases of puts or calls with 9 months or less until expiration must be paid for in full. Writers of uncovered puts or calls must deposit / maintain 100% of the option proceeds* plus 15% of the aggregate contract value (current index level x $100) minus the amount by which the option is out-of-the-money, if any, subject to a minimum for calls of option proceeds* plus 10% of the aggregate contract value and a minimum for puts of option proceeds* plus 10% of the aggregate exercise price amount. (*For calculating maintenance margin, use option current market value instead of option proceeds.) Additional margin may be required pursuant to Exchange Rule 12.10.

Cusip Number: 648815

Trading Hours: 8:30 a.m. - 3:15 p.m. Central Time (Chicago time).

Position and Exercise limits are subject to change.

Source: cboe.com

standard deviations[38]. The largest single day percentage drop in the S&P 500 index occurred on the 15th of October, and equaled 9.47%.

The first sharp price drop in this period is identified on the 9th of September, when the S&P 500 index declined by 3.47%. By considering a time frame that is centered on this date, the observation interval comprises two distinct periods. One period is characterized by normal market conditions and another period is dominated by a high frequency of extreme events. This allows for analyzing the impact of price jumps on investor expectations. Perhaps more interesting, it may provide answers to whether, and to what extent, the significant financial turmoil of 2008 was expected. Consequently, the chosen sample period is 80 trading days prior to, and 80 trading days after September 9th, i.e. May 12th 2008 to January 5th 2009. Exhibit 5.2 shows daily closing prices for the S&P 500 index from January 2007 until May 2013. The dashed black lines illustrate the chosen time frame, while the red line is at September 9th 2008.

Exhibit 5.2 - S&P 500 Index, January 2007 - May 2013

[38] The standard deviation of daily log-returns used here is based on the period from 1990 until 2013. See Exhibit 4.2 for a recap of the descriptive statistics corresponding to this sample period.

To ensure the robustness of the estimates, high liquidity levels are desired for the options applied in the calibration procedure. However, due to the extreme uncertainty surrounding financial markets in the chosen timeframe, both financial prices and liquidity exhibited large fluctuations. This is illustrated by a varying and occasionally large bid-ask spread in the dataset of SPX option prices. As the liquidity of out-of-the-money options for a large fraction of the trading days is insufficient for calibration purposes, only in-the-money options are considered.

In addition, the *volatility smile*, i.e. the tendency of implied volatility to be higher for deep out-of-the-money and deep in-the-money options than for near-the-money options, should be accounted for (Hull, 2012). Hence, for the implicit parameters to be comparable within the sample period, it is required that they are obtained from a set of options with similar *moneyness*[39]. Since the liquidity is the highest and most stable for options with moneyness between 1.05 and 1.10 in the dataset, this appears as the best choice of moneyness range.

Furthermore, *the term structure of volatility* represents how the implied volatility varies for different maturities (Wilmott, 2007). In order to prevent this phenomenon from disturbing the estimates, the time to maturity for the calibrating instruments is consistently as close as two months as possible. That is, the time to maturity ranges from 45 to 75 days. This is also among the most traded maturities. Observe that plotting the implied volatility as a function of strike and time to maturity yields the *implied volatility surface.*

In addition to the calibrating instruments, estimates of the risk-free rate of return are needed. For reasons not elaborated on he re, this rate of return can only be approximated[40]. A frequently used approach is to derive the rate of return from U.S. Treasury securities with an appropriate maturity and use this as an approximation of the risk-free rate of return (Berk and DeMarzo, 2011). This approach is thus applied here.

[39] Moneyness is a measure of how deep an option is in the money, and is calculated as the spot price of the underlying as a fraction of the strike price. The argument is based on an assumption that implied volatility correlates better with moneyness than exercise price.

[40] There are several issues related to determining the risk-free rate of return. For instance, when pricing a given cash flow, its maturity should ideally be equal to the maturity of the asset from which the risk-free rate is derived. Moreover, it is difficult to identify any completely risk-free assets in the real world. Thus, one can at best obtain an approximation of the risk-free rate of return. For more on the practical estimation of the risk-free rate of return, see, e.g., Koller, Goedhart and Wessels (2010) or Berk and DeMarzo (2011).

In particular, the risk-free rate of return is approximated by 3-month U.S. Treasury Bills. Each option is matched with the spot rate corresponding to the particular quote date. It is noted that, ideally, the maturity of the options and the Treasury Bills should be the same. Here, the maturity of the options is 15 to 45 days shorter than that of the 3-month Treasury Bills. The choice is justified by the fact that interest rates in the sample period was very low and exhibited small deviations between maturities. Additionally, rates from 3-month Treasury Bills are more easily obtainable than for shorter maturities. Historical rates for the sample period were retrieved from Yahoo! Finance (ticker: ^IRX) (finance.yahoo.com, 2013b).

The last variable that needs to be determined is the dividend yield. This parameter is for simplifying purposes assumed constant, and is calculated as the 5-year historical average dividend yield for the component stocks of the S&P500 index. The applied value of this parameter is 1.77%, based on da ta downloaded from Standard & Poor's (standardandpoors.com, 2013).

5.3 Implementation

The calibration procedure is implemented in MS Excel. An Excel function for the price of a European Call is created in the VBA developer, given that the underlying's dynamics is described by the SDE (29). The absolute difference between the obtained market prices and corresponding model prices are then squared and aggregated. By adjusting the values of the unknown parameters, the resulting sum is minimized by applying Excel's Solver add-in. The VBA code applied to create the function for the call price is presented in Appendix B.

The chosen method in Solver is *non-linear Generalized Reduced Gradient* (GRG), which is a proven and reliable approach to solving non-linear problems (Harmon, 2011). Each time the GRG algorithm is run in Solver, the starting point is slightly changed until a solution is obtained. One should be aware that if the objective function or any constraints is non-convex, the method might arrive at a locally optimal solution (Solver.com, 2013). To increase the probability of obtaining the optimal solution, the procedure should be repeated for a range of initial values and realistic constraints should be imposed. Optimally, the initial guess reflects good knowledge on the specific problem. Appendix C provides a test of the non-linear GRG method's robustness.

According to Cont and Tankov (2004b), gradient-based methods may fail for optimization problems such as (37). This is also indicated by the partial derivatives of the theoretical call price with respect to the unknown parameters. Graphical representation of these suggests that the optimization problem is non-convex. Exhibit 5.3 illustrates the sensitivity of the call price to $\tilde{\kappa}$ and δ, while Exhibit 5.4 shows the sensitivity of the prices for a range of parameter values of σ and $\tilde{\lambda}$.

As a consequence of the ill-posedness of the optimization problem (37), the chosen calibration method is carefully applied. In particular, for each set of calibration instruments, the GRG algorithm is run for a range of initial guesses. Additionally, appropriate constraints are imposed for the unknown parameters. The optimal solution is then the set of parameter values that produces the lowest aggregate pricing error, which is given by the objective function (37).

Exhibit 5.3 - Theoretical Call Prices for Different Parameter Values of $\tilde{\kappa}$ and δ

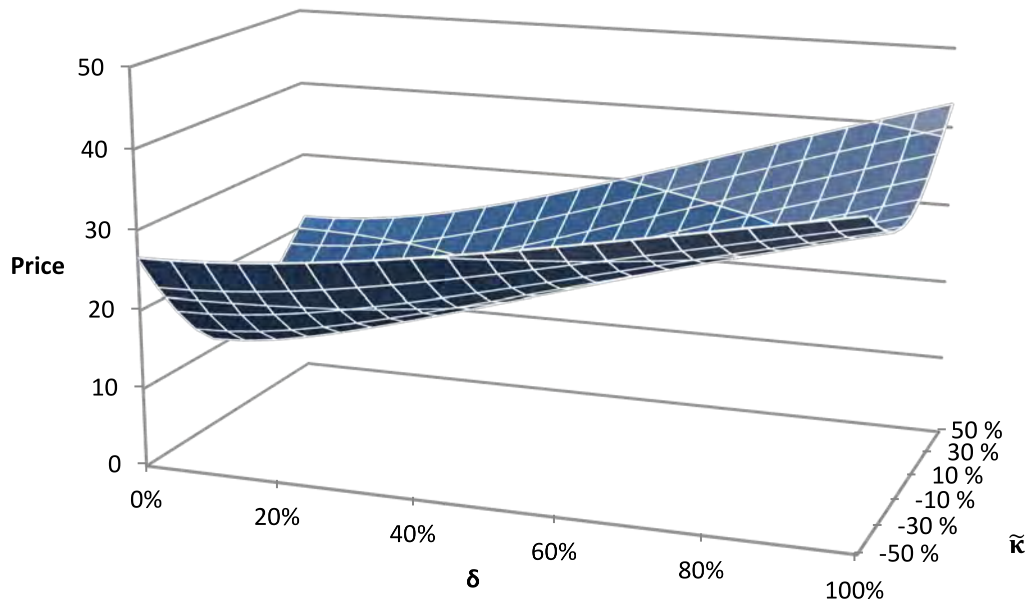

Fixed values of other parameters: $S_0 = 100, K = 100, r_f = 5\%, T = 1, q = 0, \sigma = 20\%$ and $\tilde{\lambda} = 1$.

Exhibit 5.4 - Theoretical Call Prices for Different Parameter Values of σ and $\tilde{\lambda}$

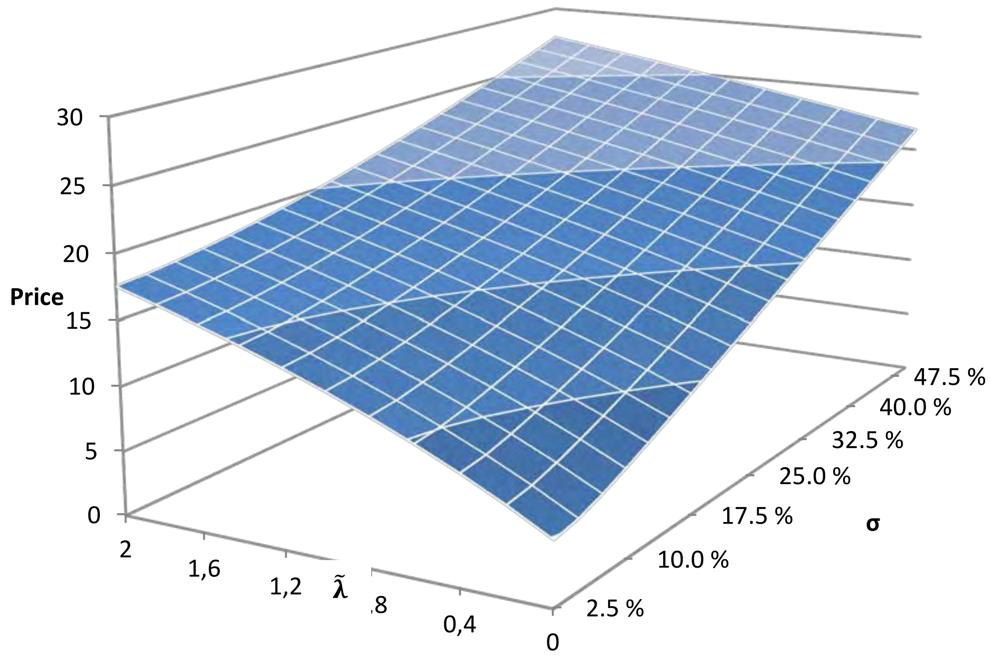

Fixed values of other parameters: $S_0 = 100, K = 100, r_f = 5\%, T = 1, q = 0, \tilde{\kappa} = 0\ and\ \delta = 30\%.$

6. Results

This section presents the results of the calibration. The obtained time-series of the implied parameters are first discussed individually, before they are evaluated in conjunction with regards to total implied volatility, given by the square root of (30). Statistical tests are performed on relevant parameters in order to determine the significance of the findings.

Two points should be emphasized regarding the results presented below. First, inferences are based upon the risk-neutral parameters. As explained above, only in the case of σ and δ are these identical to those of the real world. Bates (1991) argues, however, that the risk-neutral parameters do not differ significantly from the corresponding real world parameters, given reasonable assumptions[41].

Second, the parameter values implicit in option prices are allowed to change. This is in contrast to the assumptions that the jump-diffusion model is built upon. That is, it is inconsistent with the assumption about constant or slow-changing parameters. However, doing so allows for the generation of a time-series of the implied parameters, which allows for investigating the change in sentiment following jump events (Bates, 1991).

6.1 Implied Diffusion Coefficient

The implied diffusion coefficient (σ) averaged at 9.62% in the sample period as a whole, with an associated standard deviation of 5.06%. The largest estimate exceeds 34% is obtained for September 15th, a day when S&P 500 dropped 4.71%. On average, both the diffusion coefficient and the variability of the coefficient estimate are greater in the period after September 9th than the period before this date. This is evident from a 0.94 percentage point greater average and 2.81 percentage point greater standard deviation in this period. From Exhibit 6.1, where the daily estimates are plotted, this is easily observable.

[41] These include plausible assumptions regarding relative risk aversion and about the extent to which jumps in the S&P 500 index affects total wealth (Bates, 1991).

Exhibit 6.1 - Implied Daily Diffusion Coefficient (σ)

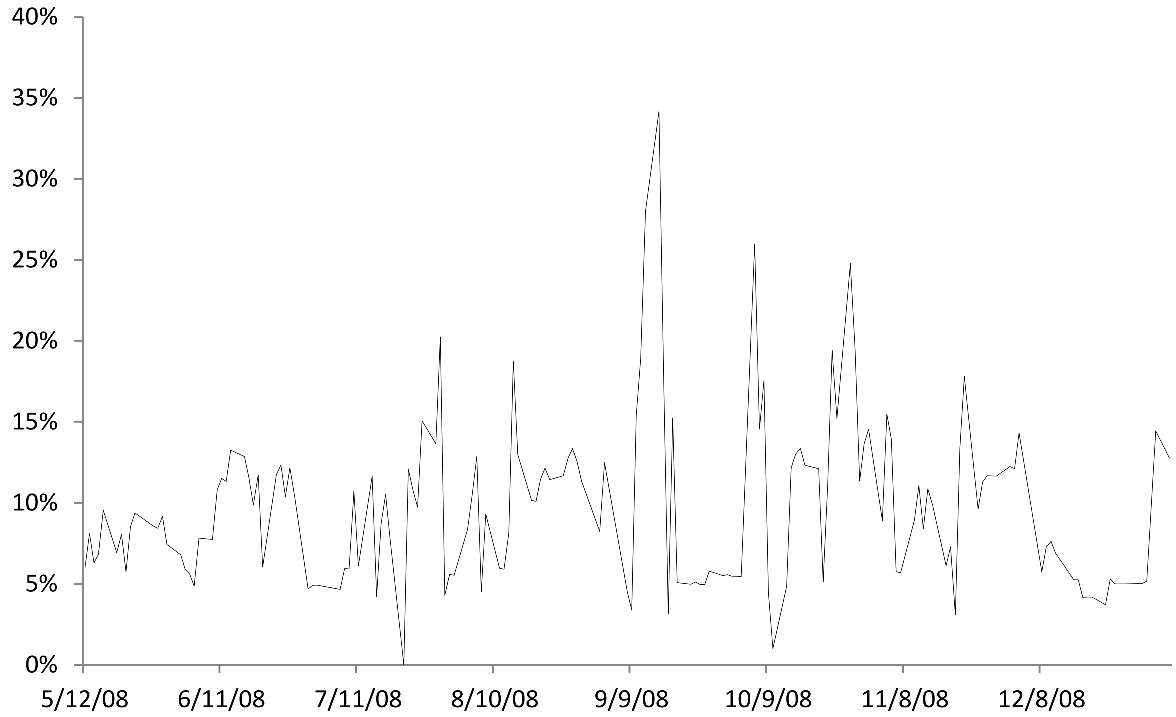

6.2 Implied Jump Intensity

Daily estimates of implied jump intensity ($\tilde{\lambda}$) varies from a minimum of 0.3 to a maximum of 13.81. The average for the whole period of 160 trading days is 5.73 jumps per year and the standard deviation is 2.18. The estimated parameter values suggest that the market on average expects a higher annual frequency of jumps in the last 80 days of the sample. Compared to the first period, the average is larger by a magnitude of 1.57, while the standard deviation is only slightly increased. Exhibit 6.2 illustrates the evolution of the implied jump intensity.

It is difficult to infer from the results that the negative shock on September 9th was expected. Still, it is noted that the implied number of jumps was relatively high before this date. This suggests an assessed risk of jumps occurring, however not to a very large extent.

It seems reasonable that investors revise their expectations regarding jumps after their occurrence. From Exhibit 6.2, it is though unclear to what direction they are adjusted. One might expect that shortly after a jump has occurred, investors expect fewer jumps in the near future. The sharp drop in the implied jump frequency on September 9th may indicate this.

Exhibit 6.2 - Daily Implied Jump Intensity ($\tilde{\lambda}$)

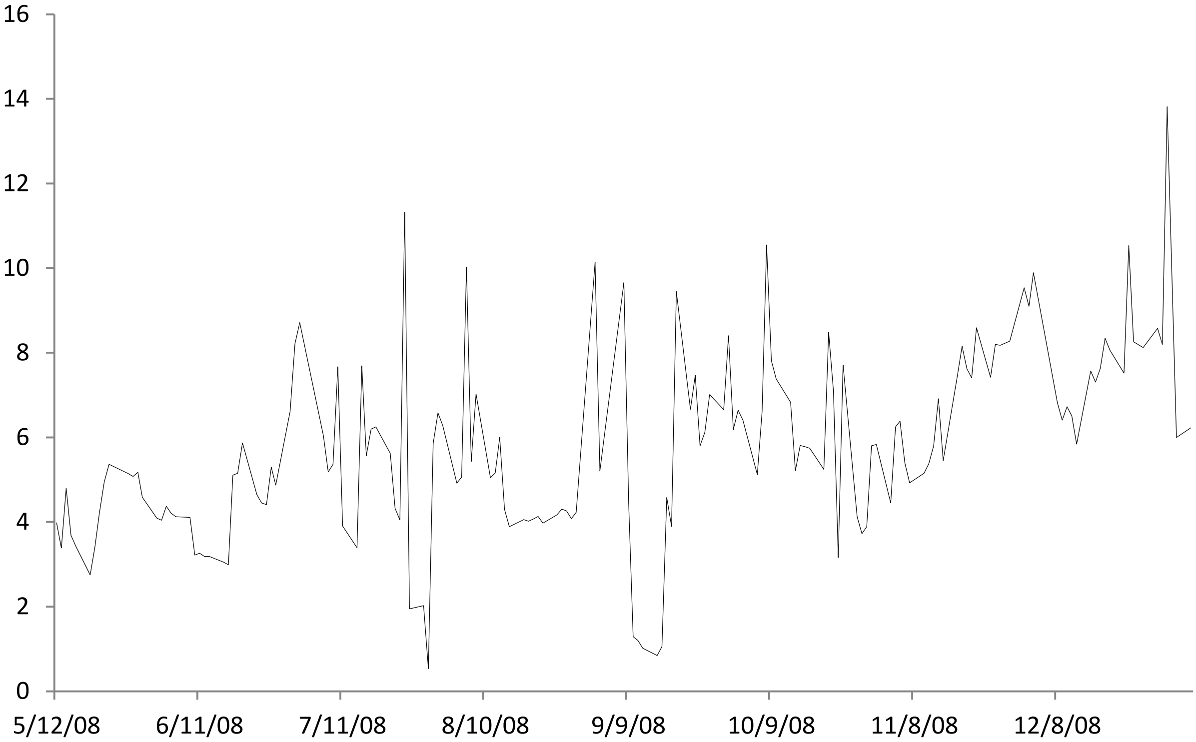

On the other hand, this is somewhat inconsistent with the observed increase in the estimated parameter value in the days after September 16th and that it then remains at relatively high levels throughout the sample period. A plausible explanation is related to the extreme market conditions characterizing the months following 15th of September. Suddenly, daily returns of magnitudes that are perceived as jumps occurred very frequently, while they before were few and far between[42]. As the uncertainty surrounding financial markets became extreme, it is intuitive that the assessed risk of jumps increased.

6.3 Implied Mean Jump Size

Obtained estimates of implied mean percentage jump size conditional on a jump occurring ($\tilde{\kappa}$) range from 0% to 7.29%. The associated average and standard deviation is 0.67% and 1.33%, respectively. The results are presented in Exhibit 6.3.

[42] As mentioned in section 5.2, a total of 32 obs ervations of daily log-returns exceeded 3 standard deviations from September 2008 until the end of the year. In comparison, from January 2007 until the start of September 2008, only a total of 3 daily log-returns of this magnitude are observed.

Exhibit 6.3 - Daily Implied Mean Jump Size ($\tilde{\kappa}$)

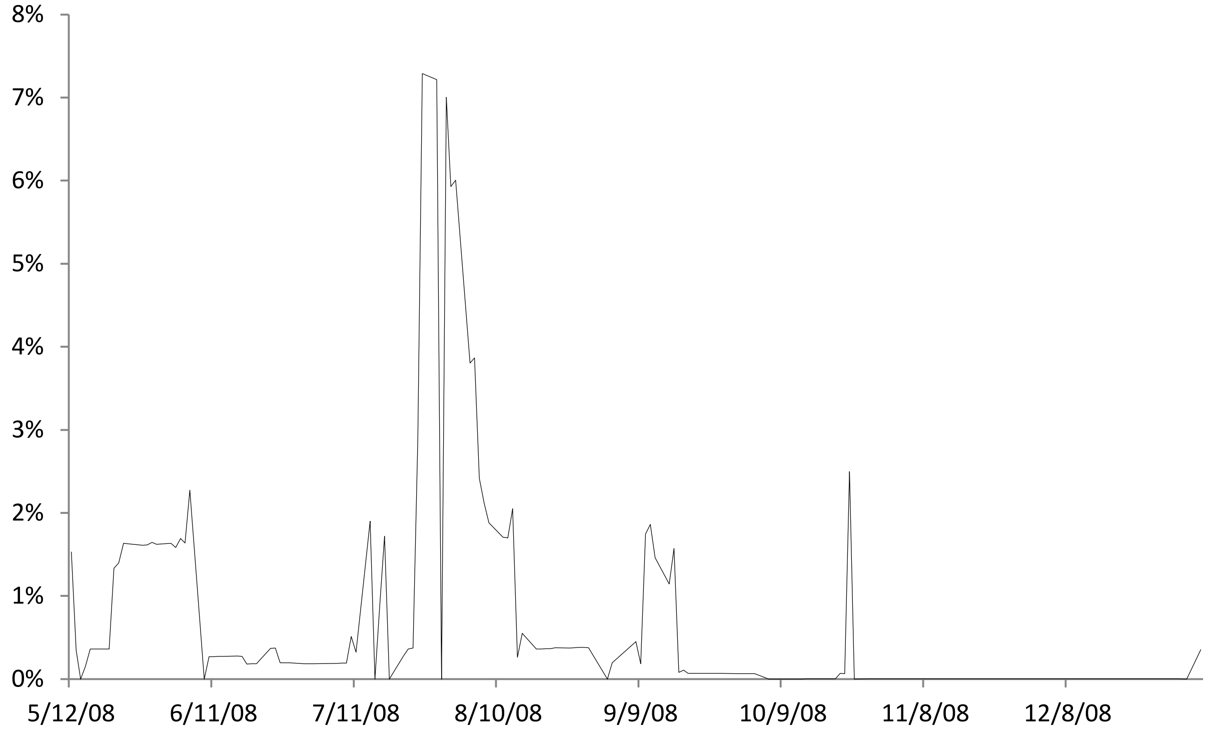

There is little or no evidence of an expected negative price jump in the derived parameter values for the mean jump size. In fact, no daily estimate yielded a negative value of the mean jump size. However, results indicate that the average of implied mean jump sizes is larger and more volatile in the first half of the sample, i.e. before the first jump event. While positive values are frequently observed in this period, the by far most common estimate in the second period is 0%. This is confirmed by a median observation of 0% for the period following September 9, reflecting the uncertainty that characterized financial markets.

6.4 Implied Standard Deviation of Percentage Jump Sizes

For the whole sample period, the average implied standard deviation of percentage jump sizes ($\tilde{\upsilon}$) is 16.69%. The estimates vary between 2.8% and 43.16% and have a standard deviation of 7.52%. Not surprisingly, there is a marked change in the implied volatility of jump sizes after September 9th. This is clear from Exhibit 6.4, which provides an illustration of the obtained estimates of this parameter.

Exhibit 6.4 - Daily Implied Standard Deviation of Percentage Jump Sizes ($\tilde{\upsilon}$)

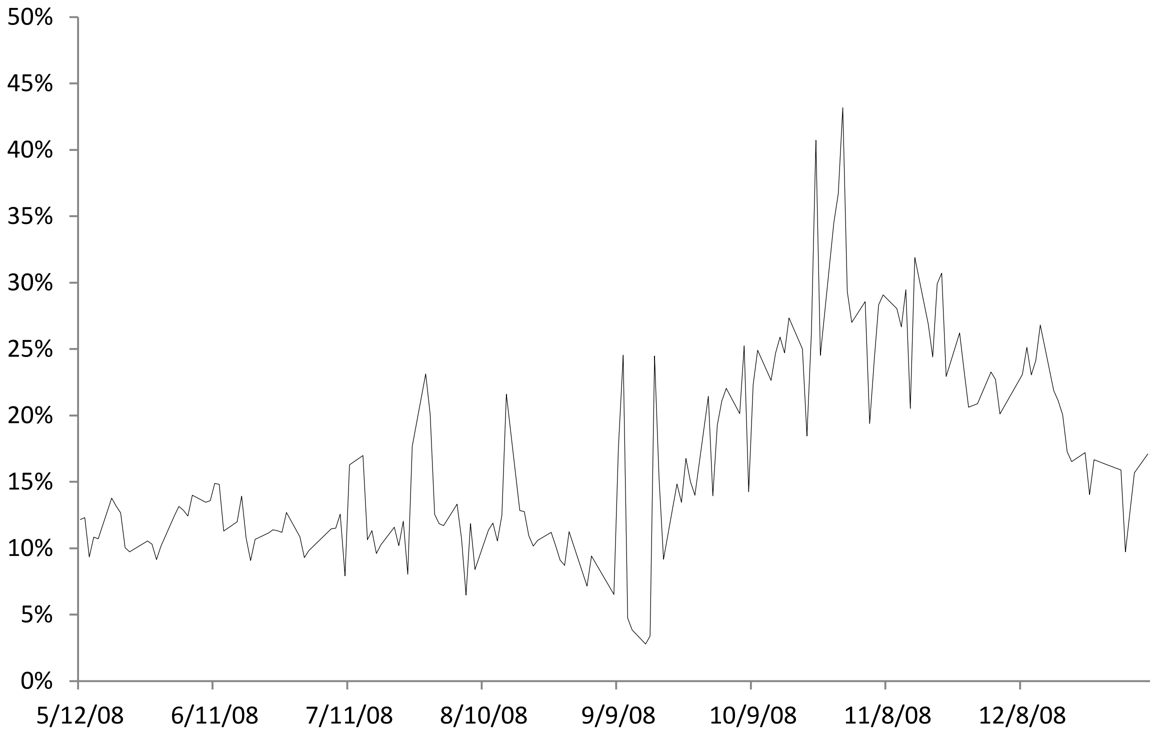

With the exception of a couple of observations in August, this parameter evolved in a relatively stable fashion until the occurrence of the first jump event. The average implied standard deviation of jumps is estimated to 11.74% for the period before this event. Evaluated in conjunction with positive estimates of implied mean jump sizes in the same period, investors did not seem to attach a high probability to the possibility of a market crash. On September 8, the implied standard deviation of percentage jump sizes is 6.53%, and the associated implied mean jump size is estimated to 0.45%.

After this date, the parameter estimates exhibit an increasing trend until the midst of November. This is a direct reflection of the significant financial turmoil associated with this period. In contrast to the first half of the sample, the corresponding average implied standard deviation of jump sizes for the second half is 21.64%.

6.5 Implied Total Volatility

By putting together the estimates of the individual parameters, the total volatility implicit in option prices in the given time frame is derived. From May 12 to September 8, the average implied volatility from the jump-diffusion model is 26.65%. Compared to historical values obtained from the VIX index[43], this is a reasonably high estimate. However, it is less than half the size of the implied volatility of the subsequent period, for which the average is 55.88%. The highest estimated value is obtained for October 29, and equals 84%. Exhibit 6.5 depicts the implied volatility from the jump-diffusion model along with the VIX index for the same period.

Exhibit 6.5 - Daily Implied Volatility From Jump-Diffusion Model and VIX Index

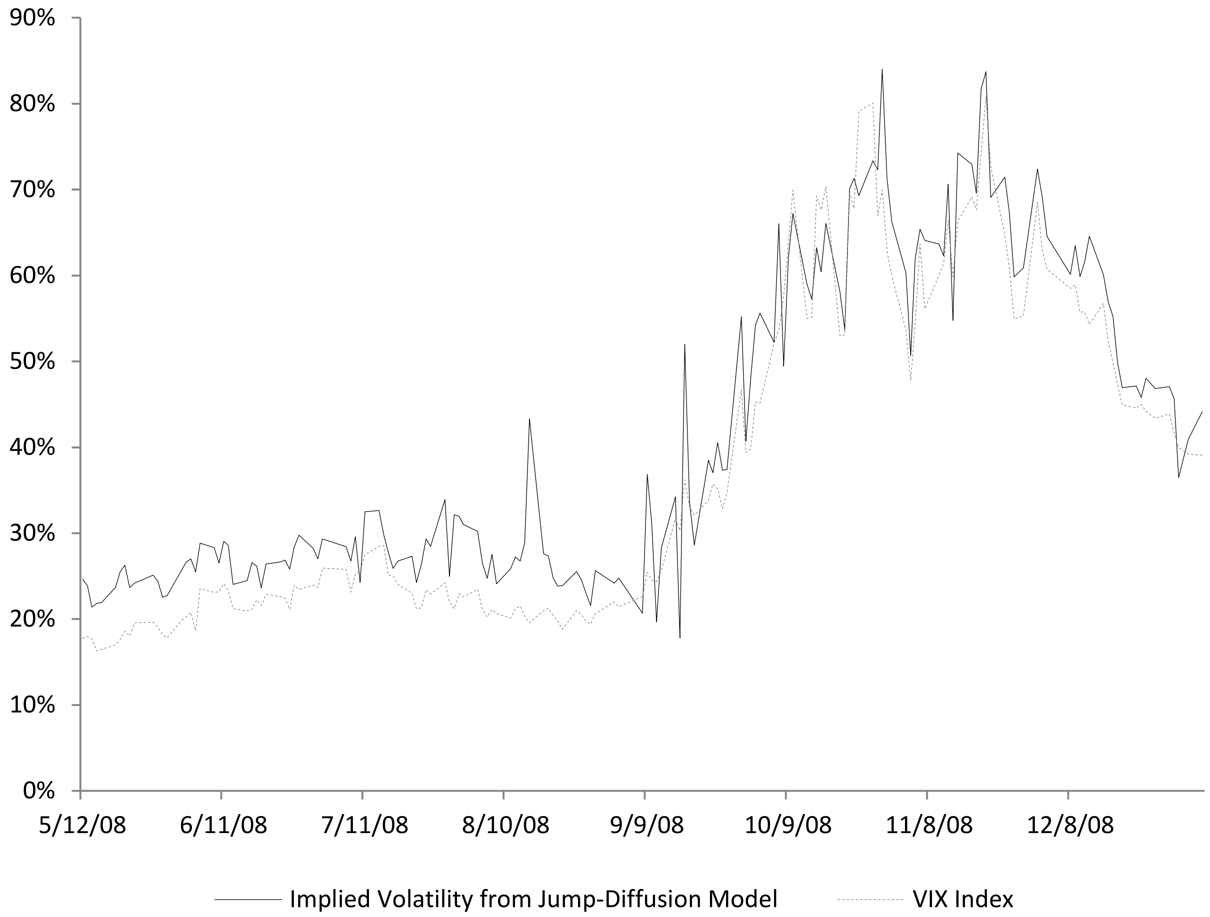

The VIX index is included in the exhibit in order to verify the obtained estimates. It is clear that the graphs track each other fairly well, even though the resemblance is not perfect. Some deviations should however be expected, as the VIX index value is calculated from a many near-term options for a range of exercise prices (cboe.com, 2013b). Additionally, some of the jump-diffusion parameters are risk-neutral.

[43] The CBOE Volatility Index (VIX) measures the implied volatility in the S&P 500 index conveyed by near-term S&P 500 stock option prices (cboe.com, 2013b).

6.6 Pricing Error

For almost all days in the sample, the jump-diffusion model fits the data well. This is evident from the consistent achievement of a low pricing error, which is measured by the objective function presented in section 5.1.2. T he most notable exception is the days following September 9. Due to low liquidity, as indicated by very large bid-ask spreads and low trading volumes, option prices were inaccurate and exhibited large degrees of variation for these days. Hence, the fitting of the parameters inevitably produced some undesirably large pricing errors. Exhibit 6.6 presents the aggregate squared pricing errors.

Exhibit 6.6 - Aggregate Squared Pricing Errors

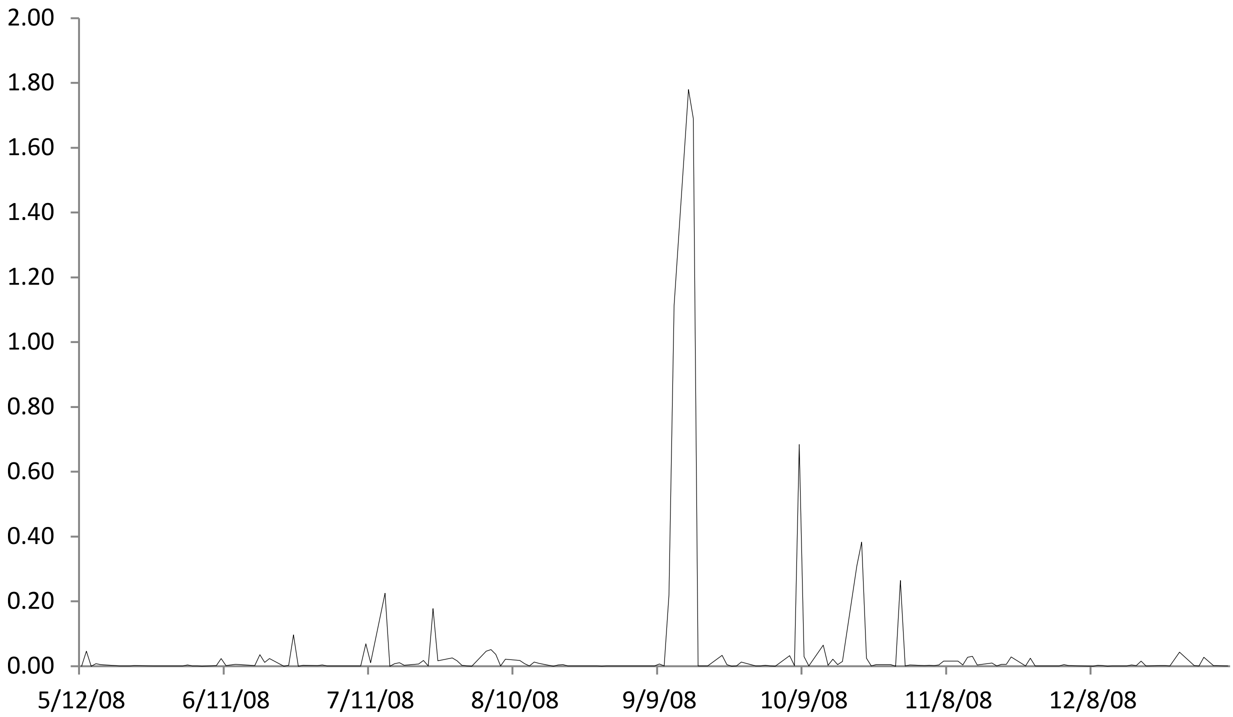

6.7 The Relevance of Jumps

The above results suggest that jumps are relevant for the pricing of European options written on the S&P 500 index. This is firstly indicated by the positive values obtained for the jump intensity. In addition to a positive probability for the occurrence of jumps, it is also required that there is a positive probability that their size is greater than 0. This is found to be the case, as the estimated values of the volatility of the jump size are positive. Consequently, it does not matter that the expected mean jump size is frequently estimated at or around 0. However, in order to establish if jumps are relevant, formal statistical tests are conducted.

For the purposes of this paper, it is sufficient to test the significance of the obtained estimates of the implied jump intensity and the implied standard deviation of the percentage jump size. It is considered unimportant to formally test the values of the implied mean jump size. There are two reasons for this. First, this will not affect conclusions on the relevance of jumps. Second, it seems obvious that this parameter is not significantly different from 0 for reasonable confidence levels.

To enable the use of tests that require normally distributed variables, the 160 daily estimates of the implied parameters subject to testing is transformed into a total of 32 non-overlapping averages comprising of 5 daily estimates. These averages are split into two samples. The first is the 16 averages obtained from the sub-period from May 12 until September 8, while the second comprises the averages estimated from September 9 until January 5 2009. According to the central limit theorem, these averages are normally distributed given a sufficiently large sample size and that the averages are iid (Keller, 2009).

Neither graphical analysis nor the Ryan-Joiner[44] test for normality rejected the hypothesis that the averages from the two sample periods follow a normal distribution. Hence t-tests are applied in all cases.

The hypothesis of no jumps occurring, i.e. $\tilde{\lambda} = 0$, is rejected at a 1% significance level, which is confirmed by a t-statistic of 18.45 and 14.70 for the first and second sample period, respectively. In the same manner is the hypothesis of deterministic jump sizes[45], i.e. $\tilde{\upsilon} = 0$, rejected at 1% significance levels. The corresponding t-statistics is calculated to 25.83 and 14.04. Hence, it is formally established that jumps in the S&P 500 index are relevant for the pricing of European options with this index as the underlying.

Lastly, it is tested whether investor expectations regarding these parameters significantly differs before and after the 9th of September. A two-sample t-test comparing the means of the

[44] The Ryan-Joiner test is preferred instead of the Anderson-Darling test because of the relatively small sample size.

[45] No variation in the mean jump size is implies deterministic jump sizes. Note that if $\tilde{\kappa}$ is not significantly different from 0, this is equivalent to testing whether the size of the jumps significantly differs from 0.

implied parameters for the two sample periods confirms on a 99% confidence level that both the implied jump intensity and the implied standard deviation of percentage jump sizes are greater in the latest sample.

Appendix E presents descriptive statistics for the 5-day averages for all parameters, also those not tested, results of the Ryan-Joiner tests for normality, and printouts from the performed t-tests.

6.8 Summary

Summing up, the estimated parameter values implicit in option prices provide three main insights. First, it is established that jumps in the S&P 500 index are relevant for the pricing of European options written on this index. A standard t-test confirms that both the implied jump intensity and the implied standard deviation of percentage jump sizes significantly differ from 0. Consequently, there was an assessed risk of jumps in the S&P 500 index for the period considered.

Second, the results do not provide evidence that investors in any way expected a market crash in September 2008. That is, there is no indication of a stronger than normal perception of downside risk. This conclusion is supported by a non-negative implied mean jump size. Implied values of the other jump parameters also indicate that the significant financial turmoil starting at the 9th of September was unexpected.

Third, the frequent occurrence of jumps in the months following September 9th of 2008 is found to significantly alter investor expectations regarding jump events. A t-test applied to the samples before and after this date confirms that a higher frequency of jumps was expected in the period after the financial crisis struck. In addition, there was a perceived risk of jumps of greater magnitude. That is, the implied standard deviation of jump sizes is found to be significantly larger in the second sample period. These findings are not surprising, given the extreme market conditions associated with this period.

7. Conclusions

In the following, the findings of this thesis are summarized. Also provided is a discussion of limitations related to the analysis, before suggestions for further research is presented in the end.

7.1 Concluding Remarks

Empirical observations suggest that most financial markets occasionally are characterized by abrupt and large price changes. This highlights the need for accurate pricing models that allows for discontinuous jumps. In the literature, such models are known as jump-diffusions. Compared to traditional pricing models, these pose additional challenges related to the adjustment for systematic risk. In particular, when allowing for jumps, markets are in general no longer complete. This precludes the construction of a perfectly replicating portfolio.

In this thesis, the theoretical framework for jump-diffusions is established. It is shown how to overcome the aforementioned challenges in order to use such models for pricing purposes. Especially, the adjustment for systematic risk is thoroughly covered with particular focus on economic intuition.

With the theoretical aspects in place, a specific jump-diffusion model is suggested for the S&P 500 index. Calibration to market data from May 2008 to January 2009 shows that jumps were relevant for the pricing of options written on the S&P 500 index. That is, an assessed risk of jumps in the index is established for this period. Furthermore, no indication of crash fears is found prior to the outbreak of the financial crisis. Lastly, the analysis shows that more frequent jumps were expected after the crisis struck.

7.2 Limitations

As noted, there exist a range of models that may explain the dynamics of S&P 500. It is not given that the chosen jump-diffusion model provides the best fit. For instance, jump-diffusion models with stochastic volatility may better describe the development in the S&P 500 index. However, it is a trade-off between efficiency and the value added by more complex models. If making the model more complex provides little extra value, there are probably good arguments for not doing so.

Ideally, several independent estimates of the implicit parameters should be derived for each trading day. Calibrating the model to an array of options with different strike prices and different maturities might improve the reliability of the results. In addition, considering puts in the calibration procedure would possibly increase the validity of the estimates. The main reason that this is not done here, is the issue of liquidity. Only call options with moneyness in the range between 1.05 and 1.10, along with a maturity between 45 and 75 days, had sufficiently high and stable trading volumes in order to be suitable for calibration. That is, options without these characteristics had, at least occasionally, bid-ask spreads so large that fairly precise estimates of the implicit parameters were very difficult to obtain.

Also, a longer time-series of the jump-diffusion parameters could have been obtained, as this probably would reduce the estimation error. On the other hand, it is undesirable to include outdated observations that are of little relevance. Hence, there is a trade-off between statistical significance and reliability. Furthermore, calibration is costly in terms of computational time.

Lastly, it is noted that there are several sophisticated optimization methods that are more robust than non-linear GRG for ill-posed calibration problems, see, e.g., Cont and Tankov (2004b). Still, this method is applied in this thesis for two reasons. First, it produces reliable and satisfactory results. This is evident by the small pricing errors that are consistently achieved. Second, the method is relatively easy to implement compared to other approaches. Hence, the need for more complex methods is considered non-existing.

7.3 Suggestions for Further Research

The framework applied in this paper is applicable to a variety of other research areas. For instance, the jump-diffusion model may be applied to a different time period, to other derivatives, or to other markets.

First, it seems interesting to reproduce the analysis by using current market prices for S&P 500 index options. Important insights about investors' assessment of risk may then be revealed. This is of particular interest since the S&P 500 index recently have been fluctuating around all-time-high-levels. Is there, for instance, a perceived risk of an impending negative price jump?

Furthermore, the same analysis can be conducted by using American rather than European options. Estimating implicit parameters from American options would require an expansion related to deriving the price of the given option. A simulation-based pricing approach or a closed-form approximation may be used. This will probably lead to fewer problems related to liquidity, as American options in general are more actively traded than their European counterparts. Also, this enables similar analysis of assets for which only American options is available.

Last but not least, the analysis can be extended to also consider the *importance* of jumps for the accuracy of pricing. That is, examining whether accounting for jumps is important for getting the price right, out of sample.

Appendices

Appendix A: Normality Testing of S&P500 log-returns

Daily Returns

The distributional properties of daily log-returns were also investigated by the Ryan-Joiner and Kolmogorov-Smirnov tests. The latter is similar to the Anderson-Darling test, with the exception that it attaches less weight to the tails of the distribution. The former uses another test statistic, which is more cumbersome to compute, and is better when sample sizes are relatively small (D'Agostino & Stephens, 1986). Results produced by these tests are displayed in Exhibits A1 and A2, respectively.

Exhibit A1 - Ryan-Joiner Test for Normality of Daily log-returns

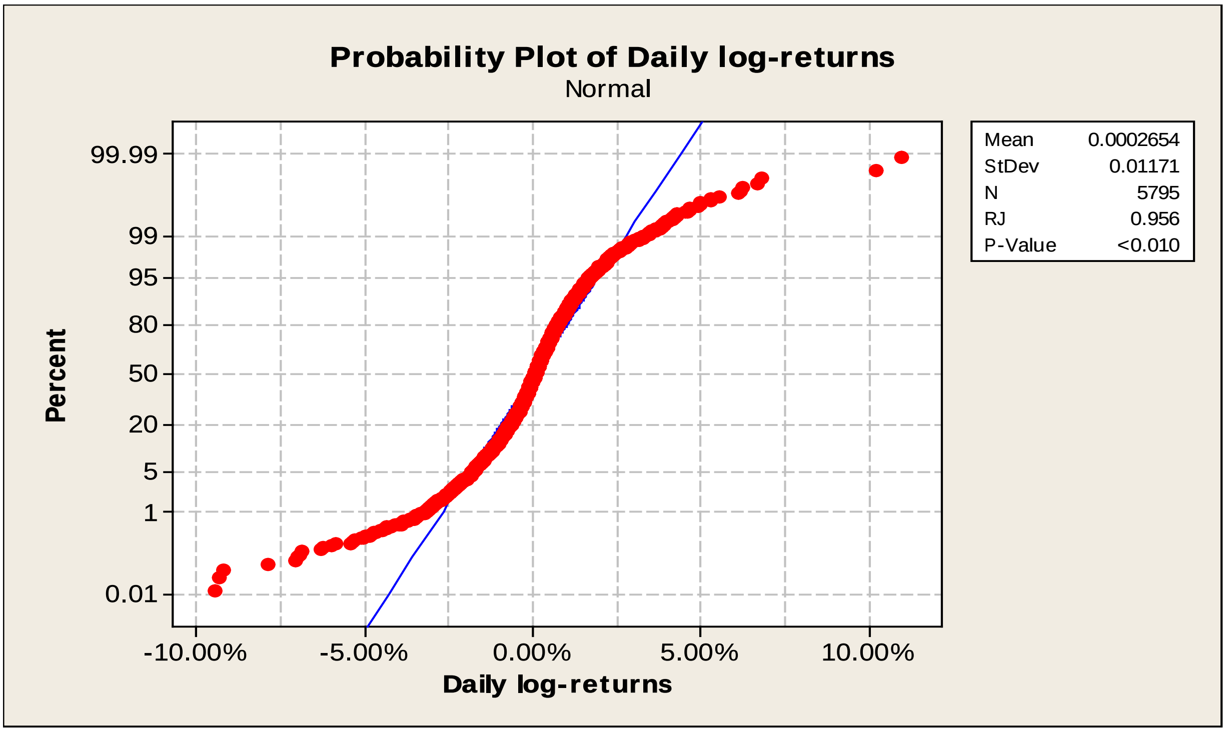

For the Ryan-Joiner test, *RJ* is the test statistic. MiniTab estimates this to 0.956. This corresponds to a p-value of less than 0.01. Consequently, the null hypothesis that the data sample is from a normal distribution is rejected for significance levels of 1% and greater.

Exhibit A2 - Kolmogorov-Smirnov Test for Normality of Daily log-returns

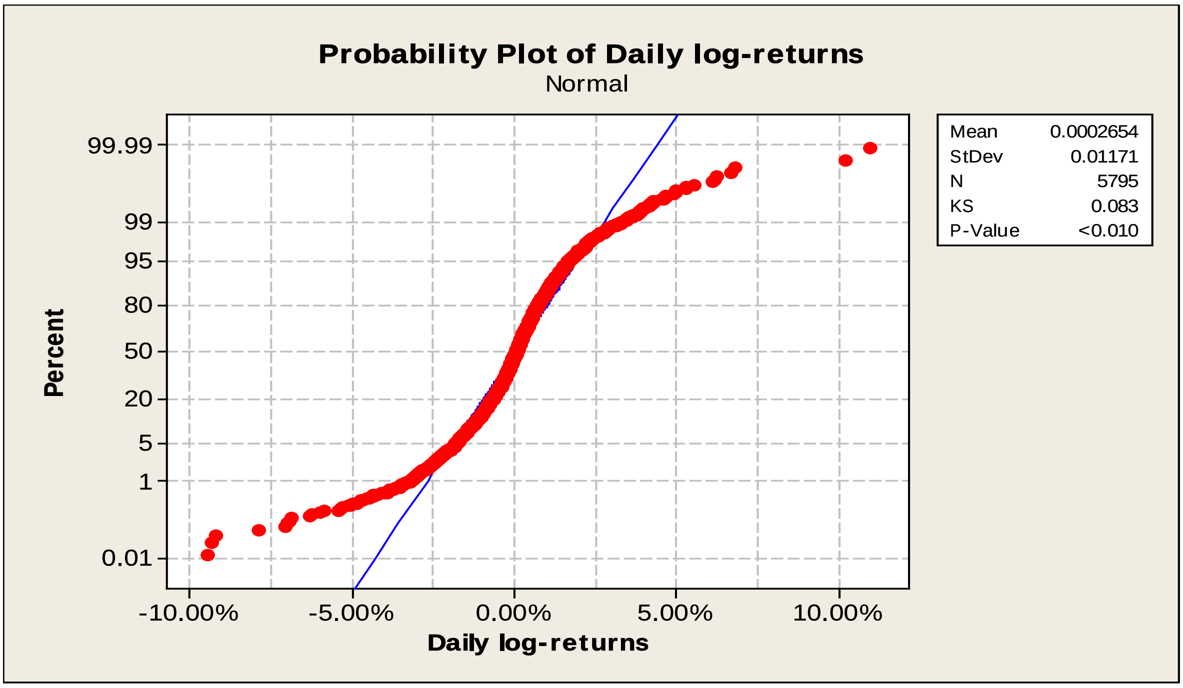

Exhibit A2 shows that The Kolmogorov-Smirnov test yield the same result as the Ryan-Joiner test. That is, the p-value that corresponds to the observed test value, *KS,* is less than 0.01. It follows that the null hypothesis is rejected for significance levels greater than or equal to 1%.

Weekly Returns

Exhibit A3 presents the empirical distribution of 1199 weekly log-returns from S&P500 for the given sample period. Also presented are descriptive statistics and the density function of the normal distribution.

Weekly log-returns seem to exhibit the same features identified for daily log-returns. Visual inspection and the calculated kurtosis suggest that the empirical density is leptokurtic, however to a lesser extent than daily log-returns. Furthermore, a greater magnitude of negative skewness is identified. As noted, this is not consistent with the normal distribution.

Exhibit A3 - Distribution of Weekly log-returns for S&P500, 1990-2013

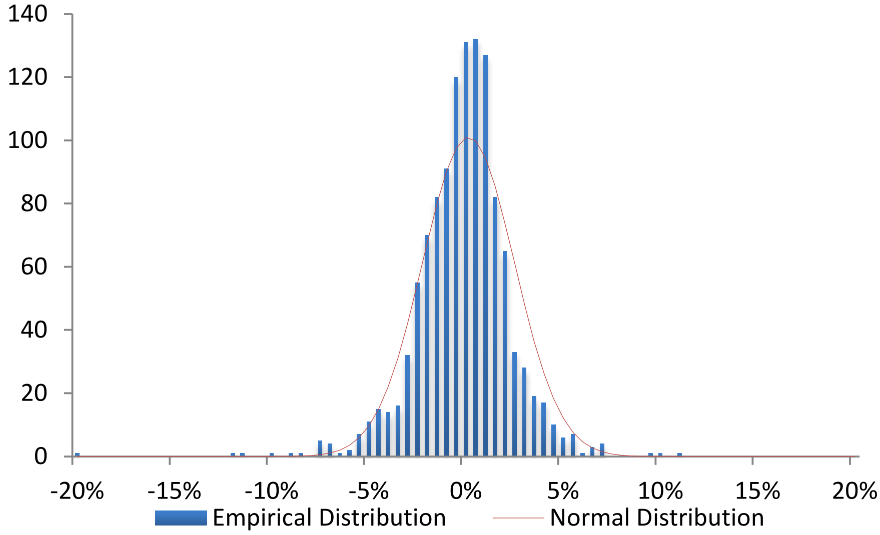

Descriptive statistics	
Mean	0.13%
Std.dev	2.37%
Skewness	-0.7226
Kurtosis	6.6618
Median	0.24%
Max	11.36%
Min	-20.08%
Q1	-1.15%
Q3	1.40%
Observations	1199

The output of the Anderson-Darling test applied to weekly log-returns is provided in Exhibit A4. Clear deviations from the straight line are observed, as the empirical CDF displays a "S-shape". AD is estimated to 10.073, for which the p-value is below 0.005. Hence, normality is rejected for weekly log-returns for 0.5% significance or more. Ryan-Joiner and Kolmogorov-Smirnov approximated p-values were both below 0.01, and yielded thus similar conclusions.

Exhibit A4 - Anderson-Darling Test for Weekly log-returns

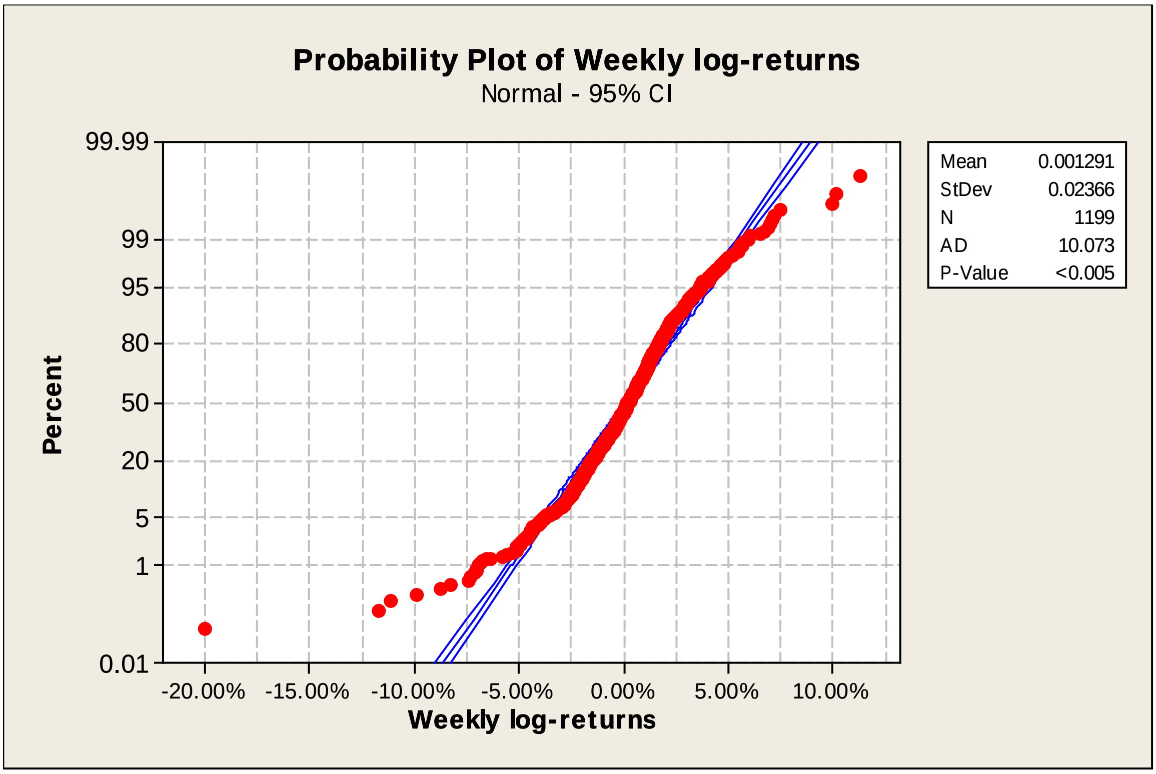

Monthly Returns

The empirical density of monthly log-returns along with descriptive statistics and the normal distribution curve is illustrated in Exhibit A5. It does not display leptokurtic features and the peak is clearly less sharp than for higher frequencies. This is supported by kurtosis of less than 3. Monthly log-returns exhibit the most negatively skewed distribution. Consequently, the figures indicate that the normal distribution is poorly resembled. This may be related to the sample size of 276 observations.

Exhibit A5 - Distribution of Monthly log-returns for S&P500, 1990-2013

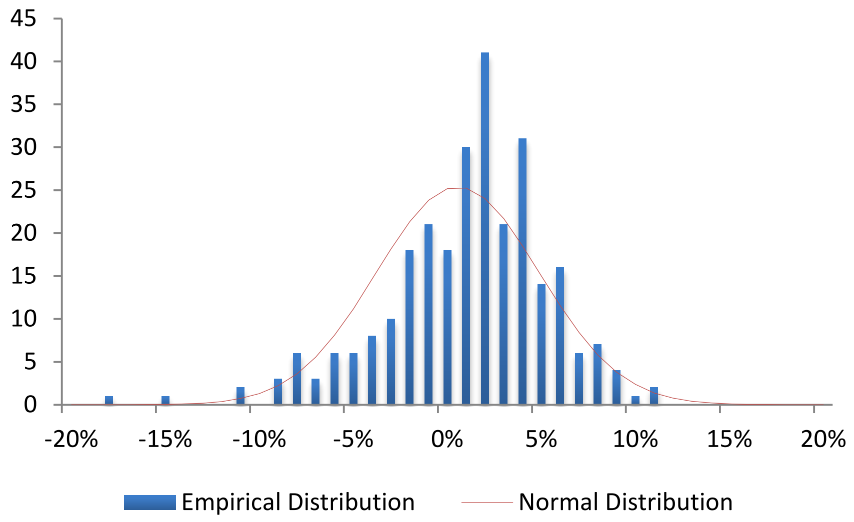

Descriptive statistics	
Mean	0.55%
Std.dev	4.34%
Skewness	-0.8241
Kurtosis	1.6998
Median	1.08%
Max	10.58%
Min	-18.56%
Q1	-1.86%
Q3	3.42%
Observations	276

The results of the Anderson-Darling test for monthly log-returns are shown in Exhibit A6. The “S-shape” observed for higher frequencies are no longer obviously present. However, the red line is far from straight and its shape indicates more outliers than the normal distribution. The test-statistic and p-value provided by MiniTab is 2.277 and less than 0.005, respectively. Hence, the null hypothesis is rejected with the same confidence as for daily and weekly data. Also here, the Ryan-Joiner and Kolmogorov-Smirnov approximated p-values were both short of 0.01. Conclusions are thus similar.

Exhibit A6 - Anderson-Darling Test for Monthly log-returns

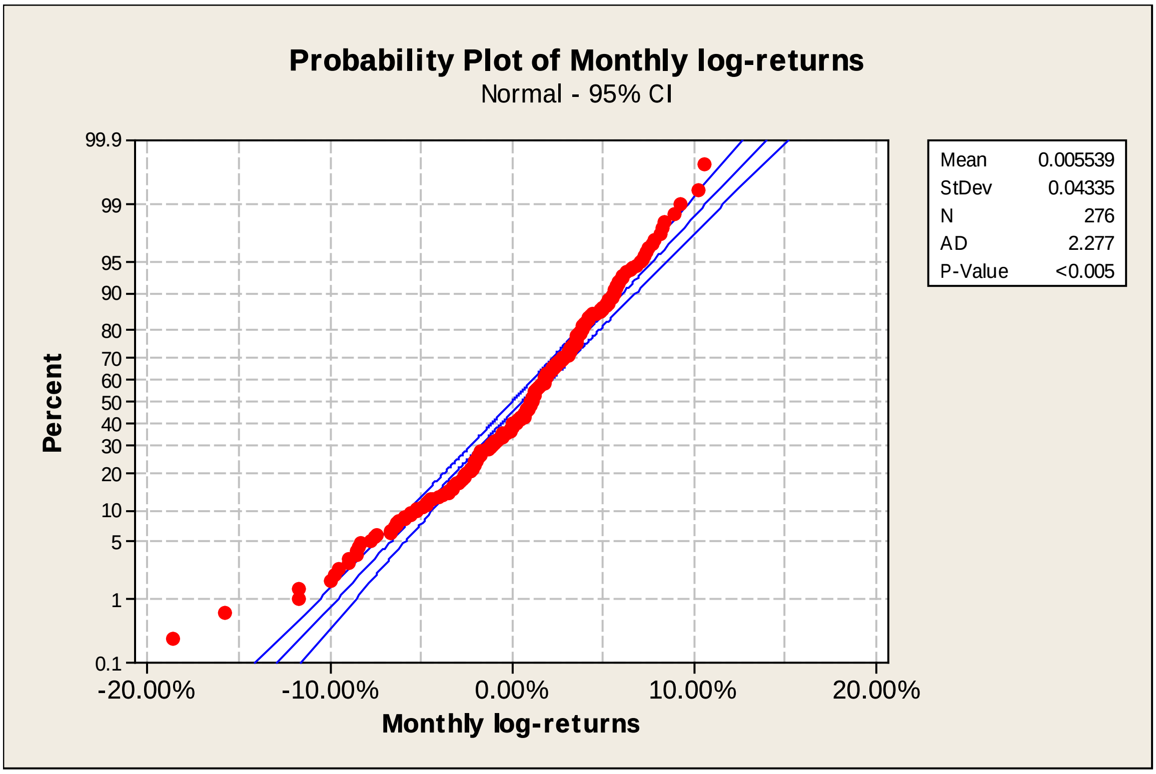

Appendix B: VBA Code Used to Calculate Model Price in Excel

The MS Excel function used to obtain the price of a European Call when the underlying's dynamics is represented by the jump-diffusion suggested in (#), is created by the VBA code presented in Exhibits B1 and B2.

Exhibit B1 – VBA Code for the Price of a European Call

```
Function EuropeanOption(CallOrPut, S, K, v, r, T, q)

Dim d1 As Double, d2 As Double, nd1 As Double, nd2 As Double
Dim nnd1 As Double, nnd2 As Double

d1 = (Log(S / K) + (r - q + 0.5 * v ^ 2) * T) / (v * Sqr(T))
d2 = (Log(S / K) + (r - q - 0.5 * v ^ 2) * T) / (v * Sqr(T))
nd1 = Application.NormSDist(d1)
nd2 = Application.NormSDist(d2)
nnd1 = Application.NormSDist(-d1)
nnd2 = Application.NormSDist(-d2)

If CallOrPut = "Call" Then

 EuropeanOption = S * Exp(-q * T) * nd1 - K * Exp(-r * T) * nd2

Else

 EuropeanOption = -S * Exp(-q * T) * nnd1 + K * Exp(-r * T) * nnd2

End If

End Function
```

Exhibit B2 – VBA Code for the Price of a European Call When Underlying is Modeled by a Jump-diffusion

```
Public Function JumpDiffusionCall(S, K, v, r, T, q, Kappa, Delta, Lambda)

Dim LAMBDAn As Double, SIGMAn As Double
Dim Rn As Double, Sum As Double
Dim i As Integer

LAMBDAn = Lambda * (1 + Kappa)

Sum = 0

For i = 0 To 100

SIGMAn = Sqr(v ^ 2 + (i * Delta ^ 2) / T)

Rn = r - (Lambda * Kappa) + i * (WorksheetFunction.Ln(1 + Kappa)) / T

Sum = Sum + (Exp(-LAMBDAn * T) * (LAMBDAn * T) ^ i / Application.Fact(i)) *
EuropeanOption("Call", S, K, SIGMAn, Rn, T, q)

Next

JumpDiffusionCall = Sum

End Function
```

Appendix C: Robustness Test of Optimization Method in Solver

To ensure that the chosen optimization method is appropriate for the given purposes, its robustness is tested. The test is conducted by confirming that Solver actually arrives at the optimal solution to the optimization problem when it is possible. That is, the method is applied to a problem of the same type as (37), defined in such a way that there exists a unique solution. This is the case if market prices are substituted with prices generated by the model. The initial guess for the value of the unknown parameters should deviate from those used to obtain the model prices. Instead of minimizing the sum in (37), it could now be set equal to zero. If Solver yields a solution by using the non-linear GRG method, it is considered robust for the application in this thesis. Exhibit C1 presents a set of randomly chosen input parameters and the corresponding model generated prices of four European call options.

Exhibit C1 - Model Generated Prices for European Calls

Parameters			***Call 1***	***Call 2***	***Call 3***	***Call 4***
Known						
	Asset Price	S	100	100	100	100
	Strike Price	K	90	100	110	120
	Risk-Free Rate	r	1.00%	1.00%	1.00%	1.00%
	Time to Maturity	T	1	1	1	1
	Dividend Yield	d	2.00%	2.00%	2.00%	2.00%
Unknown						
	Diffusion coefficient	σ	15.00%	15.00%	15.00%	15.00%
	Kappa	$\tilde{\kappa}$	0.00%	0.00%	0.00%	0.00%
	Delta	δ	10.00%	10.00%	10.00%	10.00%
	Lambda	$\tilde{\lambda}$	5	5	5	5
Model Output						
	Call Price	C_i^*	15.00	9.94	6.36	3.98

For relatively good initial guesses, the method consistently succeeds in identifying the optimal solution. However, the success seems to be somewhat sensitive to the precision of the starting values of the unknown parameters. The method still generates parameter values that produce small aggregate pricing errors even for very imprecise starting values. Nonetheless, this test highlights the need for repeating the algorithm for a range of initial guesses. When this is done, the method is considered to be sufficiently reliable for the purposes of this thesis.

Appendix D: Test of Pricing Model

The suggested pricing model is tested for errors in order to be able to put confidence in the results obtained. Bates (1991) presents a table of theoretical prices on European call options when the underlying follows the same dynamics as suggested for S&P 500 in this thesis. By applying the same input parameters, the suggested pricing model should exactly reproduce the prices obtained by Bates (1991). The results of the test are presented in Exhibit D1.

Exhibit D1 - Theoretical Option Prices Compared to Bates (1991)

					Suggested Model		**Bates (1991)**
	Jump-Diffusion Parameters		*Exercise Price K*		*European Call Option c(S,T;K)*		*European Call Option c(S,T;K)*
1)	σ = 0,1414		220		29.49		29.49
	λ = 0		235		16.39		16.39
	γ = 0		250		6.88		6.88
	δ = 0		265		2.04		2.04
			280		0.42		0.42
2)	σ = 0,10		220		29.45		29.45
	λ = 10		235		16.25		16.25
	γ = 0.01		250		6.81		6.81
	δ = 0.03		265		2.17		2.17
			280		0.56		0.56
3)	σ = 0,10		220		29.58		29.58
	λ = 10		235		16.49		16.49
	γ = -0.01		250		6.79		6.79
	δ = 0.03		265		1.88		1.88
			280		0.35		0.35

Fixed values of other parameters: $S_0 = 250,\ r_f = 10\%,\ T = 0{,}25$

No differences are identified between the models. Hence, it is concluded that the pricing model is correct and free of errors. Note that in the above exhibit, the notation of Bates (1991) is used. In terms of the notation used in this thesis, $\tilde{\lambda} = \lambda$ and $\gamma = \tilde{\alpha} + \frac{\delta^2}{2}$, or equivalently $e^{\gamma} = \tilde{\kappa}$.

Appendix E: T-Test for The Relevance of Jumps

Exhibit E1 provides descriptive statistics for the 5-day independent averages of all obtained parameters for the sample period as a whole, and for the two sub-periods.

Exhibit E1 - Descriptive Statistics, 5-day Independent Averages

12.05-05.01		Average	Std.dev	Max	Min	Median	Q1	Q3	Skewness	Kurtosis
	σ	9,62 %	3,22 %	19,96 %	4,64 %	9,21 %	7,99 %	11,83 %	1,02	2,13
	$\tilde{\lambda}$	5,73	1,65	8,65	1,76	5,63	4,64	6,67	-0,08	-0,18
	$\tilde{\kappa}$	0,67 %	1,16 %	4,88 %	0,00 %	0,26 %	0,00 %	0,59 %	2,76	7,71
	$\tilde{\upsilon}$	16,69 %	6,73 %	32,96 %	8,62 %	13,79 %	11,04 %	21,95 %	0,85	-0,48
	$\tilde{\nu}$	41,27 %	17,62 %	75,11 %	22,73 %	30,61 %	26,48 %	60,81 %	0,60	-1,31
	$\tilde{\alpha}$	-0,95 %	1,97 %	3,75 %	-5,28 %	-0,64 %	-2,40 %	-0,08 %	0,21	0,53
	δ	16,42 %	6,56 %	32,03 %	8,57 %	13,65 %	10,96 %	21,67 %	0,83	-0,57
12.05-08.09										
	σ	9,15 %	2,34 %	12,88 %	5,27 %	8,42 %	7,63 %	10,63 %	0,26	-0,97
	$\tilde{\lambda}$	4,95	1,07	6,70	3,18	4,78	4,16	5,67	0,38	-0,81
	$\tilde{\kappa}$	1,19 %	1,45 %	4,88 %	0,19 %	0,51 %	0,28 %	1,48 %	1,94	2,98
	$\tilde{\upsilon}$	11,74 %	1,82 %	16,30 %	8,62 %	11,10 %	10,79 %	12,76 %	0,98	1,60
	$\tilde{\nu}$	26,65 %	2,46 %	31,18 %	22,73 %	26,42 %	24,54 %	28,39 %	0,13	-0,84
	$\tilde{\alpha}$	0,46 %	1,34 %	3,75 %	-0,67 %	-0,11 %	-0,39 %	0,71 %	1,78	2,40
	δ	11,55 %	1,69 %	15,42 %	8,57 %	11,02 %	10,52 %	12,58 %	0,72	0,66
09.09-05.01										
	σ	10,09 %	3,93 %	19,96 %	4,64 %	9,42 %	8,44 %	11,99 %	0,92	1,46
	$\tilde{\lambda}$	6,52	1,77	8,65	1,76	6,58	5,60	7,84	-1,18	2,23
	$\tilde{\kappa}$	0,15 %	0,34 %	1,28 %	0,00 %	0,00 %	0,00 %	0,07 %	2,91	8,98
	$\tilde{\upsilon}$	21,64 %	6,17 %	32,96 %	10,69 %	22,02 %	16,16 %	26,05 %	-0,08	-0,67
	$\tilde{\nu}$	55,88 %	13,42 %	75,11 %	30,04 %	60,88 %	45,92 %	66,18 %	-0,59	-0,68
	$\tilde{\alpha}$	-2,37 %	1,39 %	0,36 %	-5,28 %	-2,40 %	-3,33 %	-1,28 %	-0,04	0,25
	δ	21,29 %	5,95 %	32,03 %	10,47 %	21,75 %	16,05 %	25,60 %	-0,15	-0,66

The Ryan-Joiner test for normality is applied to the 5-day averages obtained for the jump intensity and the standard deviation of percentage jump sizes. The test is individually performed for the sample from the first sub-period and the sample from the second sample period, respectively. Exhibit E2 presents the results.

Exhibit E2 - Ryan-Joiner Test for Normality

	$\tilde{\lambda}$_1	$\tilde{\lambda}$_2	$\tilde{\upsilon}$_1	$\tilde{\upsilon}$_2
Average	4.95	6.516	0.1174	0.2164
Std.dev	1.073	1.773	0.01818	0.06167
Observations	16	16	16	16
RJ	0.972	0.95	0.949	0.991
P-Value	>0.10	0.086	0.084	>0.10

As is evident from the obtained p-values, the test does not enable the rejection of the null hypothesis, i.e. that the parameters are normally distributed. This is also supported by graphical analysis of the samples. Hence, normality is not rejected, and t-tests are thus applicable in order to determine the significance of the results.

A one-sample t-test was conducted to all samples. In all cases, the hypothesis of a mean equal to zero was rejected on a 1% significance level. A printout from the tests is presented in Exhibit E3.

Exhibit E3 – One-sample t-test of $\tilde{\lambda}$ and $\tilde{\upsilon}$ for the two sub-periods

One-Sample T: λ_1, λ_2, υ_1, υ_2

```
Test of mu = 0 vs not = 0

Variable   N      Mean    StDev  SE Mean          99% CI              T      P
λ_1       16     4.950    1.073    0.268  (  4.160,    5.741)  18.45  0.000
λ_2       16     6.516    1.773    0.443  (  5.210,    7.822)  14.70  0.000
υ_1       16   0.11739  0.01818  0.00455  (0.10399, 0.13078)  25.83  0.000
υ_2       16    0.2164   0.0617   0.0154  ( 0.1710,   0.2619)  14.04  0.000
```

In order to determine whether the assessed values of these parameters are significantly different between the two sample periods, a two-sample t-test was performed for the jump-intensity and the standard deviation of jump sizes, respectively. Both tests yielded significant results on a 99% confidence level, i.e. that 5-day averages of these variables are significantly different between the sub-periods. Exhibit E4 and E5 provides the respective test-printouts.

Exhibit E4 - Two-sample t-test of $\tilde{\lambda}$

Two-Sample T-Test and CI: λ_1, λ_2

```
Two-sample T for λ_1 vs λ_2

      N  Mean  StDev  SE Mean
λ_1  16  4.95   1.07     0.27
λ_2  16  6.52   1.77     0.44


Difference = mu (λ_1) - mu (λ_2)
Estimate for difference:  -1.566
99% CI for difference:  (-3.015, -0.117)
T-Test of difference = 0 (vs not =): T-Value = -3.02  P-Value = 0.006  DF = 24
```

Exhibit E5 - Two-sample t-test of $\tilde{\upsilon}$

Two-Sample T-Test and CI: υ_1, υ_2

```
Two-sample T for υ_1 vs υ_2

      N    Mean   StDev  SE Mean
υ_1  16  0.1174  0.0182   0.0045
υ_2  16  0.2164  0.0617    0.015


Difference = mu (υ_1) - mu (υ_2)
Estimate for difference:  -0.0990
99% CI for difference:  (-0.1456, -0.0525)
T-Test of difference = 0 (vs not =): T-Value = -6.16  P-Value = 0.000  DF = 17
```

References

Books

Benninga, Simon. 2008. *Financial Modeling. Third Edition.* London: The MIT Press.

Benth, Fred E. 2004. *Option Theory with Stochastic Analysis. An Introduction to Mathematical Finance.* Berlin: Springer.

Berk, J., & DeMarzo, P. 2011. *Corporate Finance. Second Edition.* Boston: Pearson

Brealey, Richard A., Stewart C. Myers and Alan J. Marcus. 2009. *Fundamentals of Corporate Finance. Sixth Edition.* New York: McGraw-Hill/Irwin.

Cont, Rama, and Peter Tankov. 2004a. *Financial Modelling with Jump Processes.* Paris: Chapman & Hall/CRC Press.

D'Agostino, Ralph, B. and Michael A. Stephens. 1986. *Goodnes-of-Fit Techniques.* New York: CRC Press.

Harmon, Mark. 2011. *Step-by-Step Optimization With Solver. The Excel Statistical Master.* Excel Master Series: www.ExcelMasterSeries.com.

Hull, John. C. 2012. *Options, Futures and Other Derivatives, Eighth edition.* Upper Saddle River, NJ: Pearson Education, Inc.

Keller, Gerald. 2009. *Managerial Statistics. Eighth Edition.* Canada: South-Western Cengage Learning.

Koller, T., Goedhart, M., & Wessels, D. 2010. *Valuation. Fifth Edition.* New Jersey: John Wiley & Sons.

Neftci, Salih N. 1996. *An Introduction to the Mathematics of Financial Derivatives*. San Diego: Academic Press.

Rebonato. Riccardo. 2004. *Volatility and Correlation. Second edition.* Chichester: John Wiley & Sons.

Shreve, Steven E. 2004. *Stochastic Calculus for Finance II: Continuous-Time Models.* 1st Edition. New York: Springer Finance.

Wilmott, Paul. 2007. *Paul Wilmott Introduces Quantitative Finance. Second Edition.* New Jersey: John Wiley & Sons.

Articles/Other

Bates, David S. *The Crash of '87: Was It Expected? The Evidence from Option Markets.* The Journal of Finance, Vol. 46, No. 3, July 1991, pp.1009-1044.

Black, Fisher, and Myron Scholes. *The Pricing of Options and Corporate Liabilities.* The Journal of Political Economy, 81, May/June 1973, pp. 637-59

Cassasus, Jaime, and Pierre Collin-Dufresne. *Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates*. The Journal of Finance, Vol. 60, No. 5, October 2005, pp. 2283-2331.

Cheang, Gerald H.L, and Carl Chiarella. *A Modern View on Merton´s Jump-Diffusion Model.* Quantitative Finance Research Centre, Research paper 287, January 2011.

Cont, Rama, and Peter Tankov. *Calibration of Jump-Diffusion Option Pricing Models: A Robust Non-Parametric Approach*. Journal of Computational Finance, Vol. 7, No. 3, 2004b, pp. 1-49.

Gerber, Hans U., and Elias S.W. Shiu. *Option Pricing By Esscher Transforms.* Transactions Of Society Of Actuaries, Vol. 46, 2004, pp. 99-191.

Gilli, Manfred, and Enrico Schumann. *Calibrating Option Pricing Models with Heuristics.* COMISEF Working Papers Series, WPS-030, 8th of March, 2010.

Merton, Robert C. *Option Pricing When Underlying Stock Returns Are Discontinuous.* Journal of Financial Economics, No. 3, 1976, pp. 125-144.

Miltersen, Kristian R. 2005. *Mathematical Methods and Models in Finance.* Set of lecture notes used in the course Mathematical Finance and Models in Finance arranged by SimCorp A/S.

Xu, Mingxin. 2006. *Risk Measure Pricing and Hedging in Incomplete Markets.* Annals of Finance, Vol. 2, Issue 1, pp. 51-71.

Internet

Cboe.com. 2013a. *CBOE – Product Specifications*. Available from <http://www.cboe.com/products/indexopts/spx_spec.aspx> Downloaded 1st of June, 2013.

Cboe.com. 2013b. *CBOE – CBOE Volatility Index (VIX) Options and Futures Micro Site*. Available from <http://www.cboe.com/micro/VIX/vixintro.aspx> Downloaded 12th of June, 2013.

Standardandpoors.com. 2013. *S&P | S&P500 | Americas.* Available from < http://www.standardandpoors.com/indices/sp-500/en/us/?indexId=spusa-500-usduf--p-us-l--> Downloaded 24th of May, 2013.

Solver.com. *Standard Excel Solver – GRG Nonlinear Solver Stopping Conditions | Frontline Systems*. Available from <http://www.solver.com/standard-excel-solver-grg-nonlinear-solver-stopping-conditions> Downloaded 26th of May, 2013.

Finance.yahoo.com. 2013a. *^GSPC Historical Prices | S&P500 Stock*. Available from <http://finance.yahoo.com/q/hp?s=%5EGSPC+Historical+Prices> Downloaded 14th of May, 2013.

Finance.yahoo.com. 2013b. *^IRX Historical Prices | 13-Week Treasury Bill Stock*. Available from <http://finance.yahoo.com/q/hp?s=%5EIRX+Historical+Prices> Downloaded 29th of May, 2013.