
Multi-factor Stochastic Volatility Models

A practical approach

Filip Andersson

20573@student.hhs.se

Niklas Westermark

20653@student.hhs.se

Abstract

Since the legendary Black-Scholes (1973) model was presented, both academics and practitioners have made efforts to relax its assumptions and generate option pricing models that allow for non-normal return distributions and non-constant volatility. In this thesis, we examine the performance of four structural models ranging from the single-factor stochastic volatility model of Heston (1993) to a two-factor stochastic volatility model allowing for log-normally distributed jumps in the stock return process. We apply a practical view on the models by assuming that they are all to some degree misspecified. As a result, we do not pursue the classical route of trying to find the “true” model parameters using multiple cross-sections in the model estimation, but estimate the models daily in order to find parameters that match today’s market prices as closely as possible. The structural models are benchmarked against an *ad-hoc* Black-Scholes model, popular among practitioners. Our results show that adding an additional stochastic volatility factor to the return process significantly improves pricing performance, both in- and out-of-sample. We also show that the benefits of adding jumps to the return process are negligible in our sample, partly explained by the exclusion of very short-dated options. Lastly, we also provide some evidence on the estimation and implementation difficulties that are the drawbacks of the more sophisticated models.

Tutor: Assistant professor Roméo Tédongap.

Date and time: May 12th 2009, 10:15.

Location: Room 550.

Discussants: Alok Alström and Anna Blomstrand.

Acknowledgements: We would like to thank our tutor Roméo Tédongap for helpful advice during the writing of this thesis. We are also grateful to Misha Wolynski for valuable comments and suggestions and to Jacob Niburg for inspiring discussions.

Table of Contents

1. Introduction	1
2. Purpose and research questions	3
3. Theoretical framework	5
3.1. Risk-neutral valuation	5
3.2. Stock price dynamics	6
3.3. Valuing options using characteristic functions and the Fast Fourier Transform	7
3.4. Implied volatility and the volatility surface	9
4. Previous research	11
4.1. Stochastic volatility and jump models	12
4.2. Multi-factor stochastic volatility models	13
4.3. Local volatility models	13
4.4. Other models	14
5. Model introduction	15
5.1. Stochastic volatility model (SV)	15
5.2. Stochastic volatility model with jumps (SVJ)	17
5.3. Multifactor stochastic volatility model (MFSV)	19
5.4. Multifactor stochastic volatility model with jumps (MFSVJ)	21
5.5. The Practitioner Black-Scholes model (PBS)	22
5.6. Previous empirical findings	23
6. Methodology	26
6.1. Estimation	26
6.2. Evaluation	30
7. Data description	32
8. Results	34
8.1. Parameter estimates	35
8.2. Pricing performance	40
8.2.1. In-sample performance	40
8.2.2. Out-of-sample performance	45
8.3. Sub-sample analysis	47
8.4. Estimation and implementation issues	51
9. Conclusions	53
10. References	57
Appendix A: Figures and tables	61
Appendix B: Volatility surface parameterization	78
Appendix C: Derivation of the call price formula using characteristic functions and the FFT.	81
Appendix D: Data cleaning	84
Appendix E: Estimation	86
Appendix F: The approximate IV loss function	88

1. Introduction

Over 35 years have now passed since the publication of the famous Black & Scholes (1973) paper. Since then, an immense literature on option pricing theory has emerged in order to address the inconsistencies between the Black-Scholes model and empirical findings. In particular, the assumptions of normally distributed returns and constant volatility have been shown to be the major draw-backs of the model¹. As a result, academics and practitioners have tried to develop models that allow for non-normal return distributions and non-constant volatility. Models that allow negative correlation between the underlying stock price performance and its volatility are examples of such models that have become very popular in the literature.

The development of more sophisticated models however comes at the cost of increased complexity. While the Black-Scholes model only has one unknown parameter (volatility), stochastic volatility models and further extensions often have between five and fifteen parameters. The increased parameterization imposes a risk of over-fitted models, with poor out-of-sample performance as a consequence.

Extensions of the original stochastic volatility models include multi-factor models, with two or more stochastic volatility factors. Previous literature has focused on the use of multi-factor models for capturing the variation in option prices or, equivalently, the implied volatility surface over long time periods, sometimes up to 10 years, with only one set of model parameters.

The idea of using a long time period for model estimation may seem appealing from a theoretical point of view, as we expect the estimated model parameters to converge to the true parameters as the size of the sample gets sufficiently large. Convergence to “true” model parameters, however, relies on the assumption that there actually exist some true parameters or, in other words, that the model is correctly specified. Although this assumption is sometimes necessary in order to perform a meaningful analysis, it does not necessarily hold true.

¹ See Hull (2006) for a description of the Black-Scholes model and Cont (2001) for some stylized facts on asset returns and volatility.

Christoffersen & Jacobs (2004) argue that all option pricing models are to some degree misspecified and, as a consequence, that the standard notion that a large enough sample will result in convergence to the “true” model parameters no longer is valid. The argument carries particular implications for practitioners. For traders, speculators and investors, the main objective of any option pricing model is to price options, as of today, as accurately as possible. The practical approach to option price modeling should thus be to find a model that, when incorporating all available information as of today, prices options as accurately as possible. In other words, as the notion of convergence to “true” model parameters is no longer valid, optimal model parameters should not be based on past information.

In this thesis, we bring the practical approach to option price modeling to the field of multi-factor stochastic volatility models. We explore the subject by evaluating four structural option pricing models, ranging from a single-factor stochastic volatility model to a multi-factor stochastic volatility model that allows for log-normally distributed jumps in the return process of the underlying spot price. To further emphasize the practical perspective, the sophisticated structural models are compared to an *ad-hoc* Black-Scholes model, often referred to as the Practitioner Black-Scholes model. The models are applied to a universe of 30 686 call options written on the EURO STOXX 50 index between January 1st and December 31st 2008.

From the results, several interesting conclusions can be drawn. Partially contradicting the results of Christoffersen & Jacobs (2004), we find that the *ad-hoc* Black-Scholes model is outperformed by all structural models, especially out-of-sample. Furthermore, contrary to e.g. Bates (1996a, 2000) and Bakshi, Cao & Chen (1997), we do not find significant improvements in pricing performance of the structural models through the addition of jumps to the spot price process, not even in the short-maturity category. However, the addition of jumps does not make the models over-fitted, despite non-zero estimates of the jump factor parameters, and the out-of-sample results of the jump models are very similar to the jump free counterparts. On the other hand, expanding the parameter set by introducing additional stochastic volatility factors significantly increases pricing performance both in- and out-of-sample.

Our results show that multi-factor models are not only of academic interest for explaining the long-term development of the implied volatility surface, but also carry significant interest to practitioners looking for accurate option pricing models. To complete the analysis we would

encourage further studies of multi-factor models using single cross-section estimation, in particular with regards to the topics of hedging and exotic option pricing.

2. Purpose and research questions

The purpose of this thesis is to apply a practical view on the pricing performance of four structural option pricing models and to compare their performance to an *ad-hoc* Black-Scholes model. In order to pursue this route, some limitations must be discussed.

First of all, one must decide on a finite number of models to consider. A reasonable approach is to make this choice either to include at least one model from a range of categories in order to draw conclusions about the relationship between model structure and performance. Alternatively, one could include a number of models from within the same category in order to evaluate the effect of expanding existing models. For the purpose of this thesis, we limit our attention to five option pricing models, four of which are structural stochastic volatility based models and one is a benchmark *ad-hoc* Black-Scholes model, popular among practitioners.

Second, one must decide whether to look at pricing or hedging performance or, if possible, include both aspects. Pricing refers to the models' abilities to price various options, ranging from plain vanilla calls and puts to exotic options with complicated pay-off structures². Hedging, on the other hand, refers to the models' abilities to extract hedge parameters that can be used to manage already existing positions. In other words, hedging refers to the knowledge of which off-setting positions to engage in order to neutralize an option position's sensitivity to changes in underlying variables. The two aspects are both essential: pricing allows us to know the fair price at which to buy or sell an option and hedging allows us to manage the position once the trade has settled. Hence, any decision to engage in an option position will need input with regards to both pricing and hedging.

In terms of modeling, the two characteristics are also closely connected. In simple models, such as e.g. the standard Black-Scholes model, where analytical formulas exist for the price of many options, hedge parameters can be easily obtained by differentiating the price function. In more

² See Zhang (1998) for an overview of exotic options.

complex models, where hedge parameters have to be obtained numerically, the connection between pricing and hedging is perhaps even closer since the numerical derivative of the price function is attained by re-calculating the option price after imposing small changes in the underlying variables.

In order to enable in-depth analysis within the limited scope of this thesis, we concentrate on the pricing aspects of model performance and leave hedging performance as a topic for further studies. We also restrict the analysis to the pricing performance of plain vanilla options for which reliable price data can be acquired. This can be viewed as a first step towards a complete evaluation of the models at hand, as any such evaluation must start at parameter estimation and vanilla option pricing, before engaging into the more sophisticated fields of hedging and exotic option pricing.

Option pricing models exist in various degrees of complexity and model evaluation will always be subject to a trade-off between aspects such as pricing performance, robustness, estimation difficulties, transparency and speed. In order to provide a clear and structured evaluation of the models, we focus on answering the following three research questions:

1. Does increased model complexity enhance pricing performance?
2. Do market conditions, in terms of volatility, affect the relative performance of the models?
3. What problems arise when estimating and implementing the models?

The first question focuses purely on the performance of the models with respect to pricing errors, and leads to a suggestion which model should be adapted if pricing performance is the only benchmark. The second question aims to investigate the robustness of the models, from which conclusions can be drawn about potential biases in the performance with respect to the chosen time period and the underlying index. The third question is of a more qualitative nature, as estimation difficulty and complexity are rather subjective attributes. The aim of this question is however to shed light on potential difficulties and issues arising when using the different models rather than an attempt to measure the level of complexity.

3. Theoretical framework

3.1. Risk-neutral valuation

Risk-neutral valuation dates back to Cox & Ross (1976) who extend the results of Black & Scholes (1973). Cox & Ross recognized that if it is possible to derive an analytical expression in the form of a differential or difference-differential equation that a contingent claim must satisfy, in which one model parameter does not appear, this parameter can be altered in the model to make the underlying asset earn the risk-free rate. The value of the claim can then be calculated as its expected value using the modified parameter discounted at the risk-free rate. Harrison & Kreps (1979) extended this analysis by introducing the theory of *equivalent martingale measures*. They show that Cox and Ross' method of adjusting the model parameters is equivalent to changing probability measure from the real-world probability measure $\mathbb{P}$ to an equivalent martingale measure³ $\mathbb{Q}$, also referred to as the risk-neutral probability measure. Under the risk-neutral measure, the price of a derivative can be expressed as:

$$\Pi_t = e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}}[f(S_T)] \quad (1.1)$$

where $f(S_T)$ is the pay-off function of the derivative and r is the constant risk-free rate of return⁴. We use the short-hand notation $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot | \mathcal{F}_t]$, in which $\mathcal{F}_t$ is a filtration containing all available information at time t . The existence of an equivalent martingale measure $\mathbb{Q}$ ensures that the price is arbitrage free. In case the measure is unique, we refer to the market as complete, in which case all derivatives can be replicated using other assets. This also implies that the arbitrage free price is unique (Björk, 2004).

In layman's terms, the risk-neutral probability measure can be viewed as a different approach to modeling risk. Instead of compensating for risk through the use of a higher discount rate, the probabilities of good outcomes are adjusted to be more conservative, resulting in a lower expected value. Another way of looking at the risk-neutral probability measure is to imagine a

³ The term *martingale measure* arises from the fact that under $\mathbb{Q}$, the discounted price process of the underlying asset is a martingale. Two probability measures $\mathbb{P}$ and $\mathbb{Q}$ are said to be *equivalent* if, on a measurable space $(\Omega, \mathcal{F})$, $\mathbb{P}(A) = 0 \leftrightarrow \mathbb{Q}(A) = 0 \ \forall \ A \in \mathcal{F}$.

⁴ Under stochastic interest rates, the corresponding expression is $\Pi_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r(s) ds} f(S_T) \right]$. Note that the discount factor in this case must be inside the expectation brackets, as it is unknown at time t .

parallel world where all assets have exactly the same prices as in our world, but all investors are risk-neutral. Since risk-neutral investors only care about expected value, and thus discount all investments at the risk-free rate, the expected values of risky assets must be adjusted to be lower for asset prices to be equal to prices in the real world. Hence, in the risk-neutral world, probabilities of good and bad outcomes must differ from the corresponding real-world probabilities.

The transformation from $\mathbb{P}$ to $\mathbb{Q}$ eliminates the issue of computing an appropriate discount rate to account for risk, as the risk-neutral expectation in (1.1) is discounted at the risk-free rate. The valuation problem is thus reduced to finding the distribution of S_T under the equivalent martingale measure in order to evaluate the risk-neutral expectation of $f(S_T)$.

3.2. Stock price dynamics

In order to evaluate the expectation of $f(S_T)$, some information about the distribution of S_T must be known. Rather than making any assumptions about this distribution directly, it is most often obtained from modeling the asset price as following a continuous stochastic process. One example of a simple stochastic process is the geometric Brownian motion that the stock return is assumed to follow (under $\mathbb{P}$) in the Black-Scholes model:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t^{\mathbb{P}} \quad (1.2)$$

where μ and σ denote the drift and volatility, respectively. The Wiener process $W_t^{\mathbb{P}}$ has independent normally distributed increments, $dW_t^{\mathbb{P}} \sim N(0, dt)$. Since the value of the stock is known today, the assumption of an underlying stochastic process of the stock return enables the derivation of a distribution of the stock price at some future time point.

It is important to note the link between risk-neutral valuation and modeling of stock-price dynamics. Risk-neutral valuation requires that we use the stochastic process of the asset price under the risk-neutral measure, rather than under the real-world probability measure. In other words: the use of the stochastic process for option pricing requires the knowledge of the *risk-neutral* model parameters.

Finding the risk-neutral model parameters can be approached in two different ways. One method is to assume a process of the stock price under the real-world probability measure and use historical stock price data to estimate the parameters of the model. This approach is however troublesome if the model contains some parameter that is difficult to observe, such as e.g. the market price of volatility risk. A more commonly used estimation method that mitigates this problem is to first derive the process under the risk-neutral measure and then estimate the parameters using observed option price data, disregarding the historical performance of the underlying stock price. The latter method has an advantage in particular when a model describes an incomplete market. Recall that in an incomplete market, the equivalent martingale measure is not unique and several arbitrage free prices exist. This does however not mean that derivatives can be priced arbitrarily: *conditional on the prices observed on the market*; only one arbitrage free price will exist. Hence, the real-world modeler will face the challenge of finding the particular equivalent martingale measure chosen by the market and calculate prices accordingly. However, using the latter method and calibrating the risk-neutral model directly to option prices observed results, as required, in parameters according to the markets choice of $\mathbb{Q}$.

3.3. Valuing options using characteristic functions and the Fast Fourier Transform

In order to find the distribution of S_T , or enough information about it, characteristic functions can be used. The characteristic function of a random variable X is defined as:

$$\varphi(u) = \mathbb{E}[e^{iuX}] \quad (1.3)$$

where i refers to the imaginary unit, i.e. $i = \sqrt{-1}$. The characteristic function is defined for all u and exists for all distributions. As implied by its name, the characteristic function characterizes the distribution uniquely in the sense that every random variable possesses a unique characteristic function (Gut, 2005). Hence, there is a one-to-one relationship between the characteristic function of a random variable and its distribution.

Denoting by s_T the natural logarithm of the terminal spot price of the underlying asset i.e. $s_T = \ln(S_T)$, the characteristic function of s_T under $\mathbb{Q}$ is:

$$\varphi_T(u) = \mathbb{E}^{\mathbb{Q}}[e^{ius_T}] = \int_{\mathbb{R}} e^{ius_T} q_T(s_T) ds_T \quad (1.4)$$

where $q_T(s)$ denotes the risk-neutral density of s_T . It turns out that if the characteristic function (1.4) is known analytically, semi-analytical expressions of vanilla option prices can be obtained through the application of Fourier analysis (see e.g. Bakshi, Cao & Chen, 1997; Bates, 1996a; Heston, 1993 and Scott, 1997).

Assuming that the characteristic function of the log-stock price is known analytically⁵, the price of plain vanilla options can be determined using the Fast Fourier Transform (FFT) method first presented by Carr & Madan (1999)⁶. In this approach, the call price is expressed in terms of an inverse Fourier transform of the characteristic function of the log-stock price under the assumed stochastic process. The resulting formula can then be re-formulated to enable computation using the FFT algorithm that significantly decreases computation time compared to standard numerical methods. The pricing formula for European call options using the FFT method takes the form:

$$C_T(k) = \frac{e^{-\alpha k}}{\pi} \int_0^{\infty} e^{-i\xi k} \psi_T(\xi) d\xi \quad (1.5)$$

where

$$\psi_T(\xi) = \frac{e^{-rT} \varphi_T(\xi - (\alpha + 1)i)}{\alpha^2 + \alpha - \xi^2 + i(2\alpha + 1)\xi} \quad (1.6)$$

in which $\varphi_T(\cdot)$ denotes the characteristic function of s_T , k denotes the log of the strike price and α is a damping parameter of the model.

In order to calculate call prices, (1.5) is (after some modification) computed numerically using the FFT. Put prices are obtained using the put-call parity⁷. The derivations of (1.5), (1.6) and the discrete form of (1.5) allowing for evaluation using the FFT are shown in Appendix C.

⁵ See e.g. Applebaum (2004), Carr & Madan (1999), Gatheral (2006) and Kahl & Jäckel (2005) for discussions of how to obtain the characteristic functions of different processes.

⁶ Alternative methods are suggested by e.g. Heston (1993) and Gatheral (2006) and extensions have been provided by e.g. Lee (2004) and Cont & Tankov (2004).

⁷ It is worth noting that the put-call parity relies on the assumption of no short-sale constraints. Hence, in cases when the underlying asset is a single stock or a smaller index, where short-sale possibilities are limited, methods with explicit put price formulas may be more appropriate.

3.4. Implied volatility and the volatility surface

In the context of the Black-Scholes model, the price of a European option is a function of the the spot price (S_t), strike price (K), interest rate (r), time to maturity ($T - t$), dividend yield (q) and volatility (σ). There is generally no disagreement on the values of the first five parameters, whereas the treatment of σ has become a science in itself. The Black-Scholes model assumes that σ is a constant, namely the volatility of the underlying asset.

If the assumptions of the Black-Scholes model were true, the implied volatility, i.e. the volatility that makes the Black-Scholes price coincide with the market price⁸, of options with the same underlying asset would be constant independent of both expiry time and strike price. It turns out, however, that the implied volatility varies both with regards to time to expiry and strike price. One reason for the variation in implied volatility over different maturities, referred to as the volatility term structure, is that volatility is considered to be mean-reverting (Cont, 2001). Hence, when current volatility is low with respect to historical values, the volatility term structure tends to be upward-sloping, implying that investors expect volatility to increase, and vice versa. The term structure of volatility is also event-driven in the sense that implied volatilities will be higher for short maturities when there is an upcoming event that is likely to largely affect the stock price.

Rubinstein (1994) found that the assumption of constant implied volatilities over all strike prices was fairly correct until the stock market crash in 1987. Since then, the implied volatility as a function of the strike price, called the volatility skew, typically has a form seen in Figure 1 below. Rubinstein suggested “crash-o-phobia” as an explanation to this, meaning that traders price out-of-the-money (OTM) puts and in-the-money (ITM) calls relatively higher than ITM puts and OTM calls, in order to protect themselves against the risk of a new stock market crash. Another observation, shown by e.g. Black (1976), is that the risk of a company increases with leverage. As equity decreases, the volatility increases due to the higher risk, and vice versa. In that context, the volatility is expected to be a decreasing function of price, which in turn gives rise to the common smirk shape of the volatility skew (Figlewski & Wang, 2000).

⁸ Since all other parameters are assumed to be known, the price essentially only depends on σ . Hence, for a given market price, we can solve for the value of sigma that makes the model price equal the market price.


Figure 1 below shows the volatility skew and term structure for the EURO STOXX 50 index as of July 17th 2008. Skew plots for all maturities on the same date are shown in Appendix B.

As can be seen, both plots confirm that the assumption of constant volatility over different strike prices and maturities is inconsistent with observed implied volatilities in the market. Hence, regardless of choice of σ , the Black-Scholes model will be unable to replicate market prices as a constant σ implies a horizontal line in both plots.

Figure 1

Volatility skew and term structure of EURO STOXX 50 on July 17th 2008


The left plot shows how the implied volatility decreases with strike price for call options with 32 days to maturity. The right plot shows how the implied volatility differs between ATM options with different maturities. Both plots are conflicting with the Black-Scholes assumption of constant volatility.



In order to study the implied volatility patterns in more detail, it is necessary to look at the term structure for every strike price, as well as the skew for every maturity simultaneously. To incorporate all available information with regards to both term structure and skew, we would thus need one graph for each strike price showing the term structure, as well as one graph for each maturity displaying the skew. The problem is readily solved by showing the implied volatility as a two-variable function of time and strike price in a 3D-graph. The resulting surface is referred to as the volatility surface, and shows all available information with regards to term structure and skew at a given time point. Figure 2 below shows the volatility surface of the EURO STOXX 50 index as of July 17th 2008. The surface is obtained by interpolation of the skew plots shown in Appendix B, where the calibration procedure is also described in detail.

Figure 2
Volatility surface of EURO STOXX 50 July 17th 2008

The plot shows the implied volatility (calculated from option prices) at the specific date for days to maturity and strike price. The surface is obtained by interpolating the skew plots from Appendix B.



The volatility surface plays an important role in the pricing of options. The first step towards a useful pricing model is that the model is able to replicate plain vanilla prices observed in the market. This is essentially equivalent to matching the observed implied volatilities, i.e. the market's volatility surface. Obviously, the Black-Scholes model is unable to accomplish this, as volatility in the Black-Scholes model is assumed to be constant for all maturities and strikes, implying a flat volatility surface.

4. Previous research

In this section, we present previous research on stochastic volatility models, jump models, multifactor models and local volatility models. A summary of the empirical performances of the models are presented at the end of Section 5, after the models used in this thesis have been presented in more detail.

4.1. Stochastic volatility and jump models

In stochastic volatility models, the volatility in addition to the stock price, is allowed to develop according to a stochastic process. Many different models have been proposed with the common property that volatility is modeled by its own diffusion process⁹.

In order to find a reasonable diffusion model for volatility, one must first consider some empirical facts of asset returns and volatilities. As mentioned, one of the most well-known properties of volatility is that it tends to be high in bear markets and low in bull markets, partially explained by the leverage effect. The negative correlation to asset returns is very important in the modeling of option prices, as it allows the model to generate the empirically observed volatility smirk. Additional well-documented properties that affect the prices of options and should be incorporated into any plausible stochastic volatility model, pointed out by e.g. Gatheral (2006), are volatility clustering and mean-reversion. Many stochastic volatility models, such as the Heston (1993) model indeed encompass these features. One short-coming of the stochastic volatility models is, however, their inability to capture the large short-term movements of stock prices that are observed frequently in the market. To this end, so called jump-diffusion models have been developed.

The idea of adding a jump factor to the modeling of stock prices is not a new idea, but was introduced by Merton (1976)¹⁰ short after the publication of the Black-Scholes model. The jump feature especially enables the model to explain the probabilities of large short-term moves in the stock price implied by far out-of-the money bid prices. Gatheral (2006) shows examples of 5 cent bid prices for 67 % OTM call options expiring the following morning, implying that traders are willing to pay 5 cents for options that, under normally distributed returns, have zero (to about 40 decimal places) probability of ending up in the money. Stochastic volatility models without jumps are unable to capture this implied probability of large short-term moves, and produce lower implied volatilities, and thus lower prices, for far OTM options with short maturities compared to observed prices in the market. Allowing for jumps is one way of mitigating this

⁹ See e.g. Hull & White (1987), Johnson & Shanno (1987), Melino & Turnbull (1990), Scott (1987), Stein & Stein (1991) and Wiggins (1987), although some of these models are obsolete in light of more recent models.

¹⁰ Merton's model is however a pure jump model, i.e. a model with deterministic volatility.

problem, as it will incorporate a certain probability of large instantaneous moves in the stock price.

Several different jump-models have been proposed, with and without stochastic volatility, and with different distributions of the jump size. Cox, Ross & Rubinstein (1979) suggest a pure jump model with constant jump size, whereas Merton (1976) proposes a pure jump model with log-normally distributed jump size. Extensions of the latter include Bates (1996a) who incorporates stochastic volatility as well as log-normally distributed jumps in the stock price process. Zhu (2000) conducts an extensive analysis of option pricing models, including models with log-normally distributed jumps, Pareto distributed jumps and different types of stochastic volatility diffusion processes.

4.2. Multi-factor stochastic volatility models

Bates (2000) and Christoffersen, Heston & Jacobs (2009) both propose two-factor stochastic volatility models as an alternative or extension to jump models in order to model the evolution of the implied volatility surface. The rationale behind the multi-factor model is that it is able to capture both long- and short-term movements in the volatility process. This enables the model to explain differences in both level and slope of the implied volatility surface over time.

Christoffersen, Heston & Jacobs (2009) highlight that the two-factor model has a particular advantage when estimating models using multiple cross-sections of options, as the one-factor model will suffer from structural problems when the slope and level of the implied volatility surface change simultaneously over time. The model of Bates (2000) also allows for log-normally distributed jumps in the stock price process, in addition to having two stochastic volatility factors. This extension is natural, as jumps and multiple stochastic volatility factors serve different purposes and thus should not necessarily be seen as substitutes.

4.3. Local volatility models

In local volatility models – also referred to as deterministic volatility function models – the volatility of the underlying asset is assumed to be a function of the level of the spot price and calendar time, i.e. $\sigma_t = \sigma(S_t, t)$. In continuous time, the risk-neutral stock return process in the local volatility framework is hence of the form:

$$\frac{dS_t}{S_t} = (r - q)dt + \sigma(S_t, t)dW_t^{\mathbb{Q}} \quad (4.1)$$

where r and q denote the interest and dividend yield, respectively.

The local volatility model was introduced in a discrete setting (using an implied tree method) by Derman & Kani (1994) and Rubinstein (1994), and extended to continuous time by Dupire (1994). The local volatility function $\sigma_t = \sigma(S_t, t)$ is derived to make the model consistent with observed market prices or, equivalently, consistent to the observed implied volatility surface (see e.g. Rebonato (1999) for the derivation of the local and the relation to implied volatility). Since σ_t is a function of a stochastic quantity (S_t), σ_t will also be stochastic.

Local volatility models differ from many other option pricing models in the sense that the purpose not is to model the actual evolution of the implied volatility surface, but rather provide a (not as harsh as Black & Scholes') simplification in order to enable pricing of options consistent with existing prices of vanilla options (Gatheral, 2006). The notion is confirmed by Dumas, Fleming & Whaley (1998) who conclude that the local volatility model is unable to explain the empirical dynamics of the implied volatility surface. Instead, Dumas, Fleming & Whaley propose a different type of deterministic volatility function model, in which a function of strike price and maturity, i.e. $\sigma_t = \sigma(K, T - t)$, is fitted to the observed implied volatility surface. Obviously, this function cannot be inserted into the stock price process, as doing so would lead to different processes for the same underlying stock depending on the strike price and maturity of the option at hand. Instead, the function is used to derive the implied volatility of non-traded options in order to enable pricing using the standard Black-Scholes formula.

4.4. Other models

Eraker (2004) extends the modeling of stochastic volatility to allow for jumps also in the volatility process, following in the tracks of Bates (2000) who concludes that volatility jump models are necessary for capturing the volatility shocks observed in the S&P 500 futures market. Other popular models include the variance-gamma model, proposed by Carr, Chang & Madan (1998), in which the stock price return follows a geometric Brownian motion conditional on the realization of a gamma-distributed random time. Extensions of the variance-gamma model, put

forward by e.g. Carr, Geman, Madan & Yor (2001), include models where the underlying stock price is allowed to follow other Lévy processes¹¹, driven by stochastic clocks.

5. Model introduction

In this section, we introduce the models under evaluation in more detail. For each model, we provide some intuition to the features of the model making it appealing for option pricing and, where relevant, specify the assumptions of the underlying stock price process and the corresponding characteristic function of the log-stock price. The presentation, especially with regards to the characteristic functions, is in some parts rather technical, but the reader finding it difficult to interpret the technical details may pass those parts over without any substantial loss in intuition.

For all models, we consider the risk-neutral dynamics of the stock price. We let $S = \{S_t, 0 \leq t \leq T\}$ denote the stock price process and $V = \{V_t, 0 \leq t \leq T\}$ denote the stochastic variance process. $\varphi_T(\cdot)$ denotes the characteristic function of the natural logarithm of the terminal stock price $s_T = \ln(S_T)$. The constants r and q will denote the, both constant and continuously compounded, interest rate and dividend yield, respectively. Further, we let $W_t^{\mathbb{Q}}$ denote a $\mathbb{Q}$ -Wiener process¹².

5.1. Stochastic volatility model (SV)

Allowing for the volatility of the stock price to be stochastic by itself is a well-known way of mitigating the aforementioned problems in the underlying assumptions of the Black-Scholes model. Stochastic volatility obviously allows for non-constant volatility, and also permits non-normal distributions of returns. Many different stochastic volatility models have been proposed, but we will limit our attention to the Heston (1993) stochastic volatility model, henceforth denoted SV, in which the spot price is described by the following stochastic differential equations (SDEs) under $\mathbb{Q}$:

¹¹ See Applebaum (2004) for more on applications of Lévy processes in finance.

¹² A $\mathbb{Q}$ -Wiener process is a process that fulfills the requirements of a Wiener process under the equivalent martingale measure $\mathbb{Q}$. See Björk (2004) for a more detailed description of Wiener processes.

$$\frac{dS_t}{S_t} = (r - q)dt + \sqrt{V_t}dW_t^{\mathbb{Q}(1)} \quad (5.1)$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^{\mathbb{Q}(2)} \quad (5.2)$$

$$\text{Cov}_t(dW_t^{\mathbb{Q}(1)}, dW_t^{\mathbb{Q}(2)}) = \rho dt \quad (5.3)$$

where the parameters κ , θ and σ represent the speed of mean reversion, the long-run mean and the volatility of the variance, and ρ represents the correlation between the variance and stock price processes, respectively. In addition to these parameters, the model requires the estimation of the instantaneous spot variance V_0 .

Pricing of plain-vanilla call options using the SV model can be done in several ways. Heston (1993) proposes a closed-form solution for the call price, also implemented and extended by e.g. Gatheral (2006). The closed form solution however requires numerical evaluation of the integral obtained from inversion of the characteristic function, and does thus not have the computational advantage of closed-form solutions that can be evaluated analytically (such as e.g. the Black-Scholes model). In order to minimize computation time, we will instead use the method of Carr & Madan (1999), described in Section 3 and Appendix B, and price options using the Fast Fourier Transform (FFT).

Albrecher, Mayer, Schoutens & Tistaert (2006) show that the characteristic function of s_T in the SV model requires some consideration in order to avoid numerical problems when pricing vanilla options using Fourier methods¹³. The characteristic function of the SV model, regardless of specification, includes a logarithm of complex numbers. The numerical problem, first recognized by Schöbel & Zhu (1999), arises due to the fact that the logarithm function is discontinuous in its imaginary part along the negative real axis. Hence, in order to avoid discontinuities, it is important that the argument of the logarithm function does not cross the negative real axis, which Albrecher, Mayer, Schoutens & Tistaert show can be achieved by re-formulating the characteristic function. Hence, we deviate from the original characteristic function proposed by Heston (1993) and instead use the alternative formulation proposed by Albrecher, Mayer,

¹³ Kahl & Lord (2006) provide an alternative proof using a rotation count algorithm presented by Kahl & Jäckel (2005). Their conclusion is however identical to that of Albrecher, Mayer, Schoutens & Tistaert, namely that the proposed representation mitigates the problems of the original characteristic function in Heston (1993).

Schoutens & Tistaert. Using the same representation of the parameters as in equations (5.1) – (5.3), the characteristic function of s_T takes the following form:

$$\varphi_T^{SV}(u) = S_0^{iu} f(V_0, u, T) \quad (5.4)$$

where

$$f(V_0, u, T) = \exp\{A(u, T) + B(u, T)V_0\} \quad (5.5)$$

$$A(u, T) = (r - q)iuT + \frac{\kappa\theta}{\sigma^2} \left[(\kappa - \rho\sigma iu - d)T - 2 \ln \left(\frac{1 - ge^{-dT}}{1 - g} \right) \right] \quad (5.6)$$

$$B(u, T) = \frac{\kappa - \rho\sigma iu - d}{\sigma^2} \left[\frac{1 - e^{-dT}}{1 - ge^{-dT}} \right] \quad (5.7)$$

$$d = \sqrt{(\rho\sigma iu - \kappa)^2 + \sigma^2(iu + u^2)} \quad (5.8)$$

$$g = (\kappa - \rho\sigma iu - d)/(\kappa - \rho\sigma iu + d) \quad (5.9)$$

The derivation of (5.4) is rather complicated and is thus omitted. The interested reader is referred to Gatheral (2006) or Kahl & Jäckel (2005).

Vanilla call prices in the SV model are calculated by substituting the characteristic function (5.4) into the Carr & Madan (1999) pricing formula (1.5) and evaluating using the FFT. The SV model also allows for straightforward pricing of exotic options using Monte Carlo simulation. Once the parameters have been estimated, sample paths of the process (5.1) can be simulated, allowing for the pricing of any contingent claim.

5.2. Stochastic volatility model with jumps (SVJ)

We extend the SV model in the previous section along the lines of Bates (1996a), by adding log-normally distributed jumps to the stock price process. In this model, denoted SVJ, the return process of the spot price is described by the following set of SDEs under $\mathbb{Q}$:

$$\frac{dS_t}{S_t} = (r - q - \lambda\mu_J)dt + \sqrt{V_t}dW_t^{\mathbb{Q}(1)} + J_t dY_t \quad (5.10)$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^{\mathbb{Q}(2)} \quad (5.11)$$

$$\text{Cov}_t(dW_t^{\mathbb{Q}(1)}, dW_t^{\mathbb{Q}(2)}) = \rho dt \quad (5.12)$$

where $Y = \{Y_t, 0 \leq t \leq T\}$ is a Poisson process with intensity $\lambda > 0$, i.e. $\mathbb{Q}(dY_t = 1) = \lambda dt$ and $\mathbb{Q}(dY_t = 0) = 1 - \lambda dt$, and J_t is the jump size conditional on a jump occurring. All other

parameters are defined as in (5.1) – (5.3). The subtraction of $\lambda\mu_j$ in the drift term compensates for the expected drift added by the jump component, so that the total drift of the process, as required for risk-neutral valuation, remains $(r - q)dt$.

As mentioned, the jump size is assumed to be log-normally distributed:

$$\ln(1 + J_t) \sim N\left(\ln(1 + \mu_j) - \frac{\sigma_j^2}{2}, \sigma_j^2\right) \quad (5.13)$$

Further, it is assumed that Y_t and J_t are independent of each other as well as of $W_t^{\mathbb{Q}(1)}$ and $W_t^{\mathbb{Q}(2)}$.

In the SVJ model, the total variance of the return depends both on V_t and on the variance added by the jump factor. Denoting the variance added by the jump component $V_{j,t}$, the total variance of the return process equals (Bakshi, Cao & Chen, 1997):

$$\text{Var}_t\left(\frac{dS_t}{S_t}\right) = V_t dt + V_{j,t} dt \quad (5.14)$$

where

$$V_{j,t} = \text{Var}_t(J_t dY_t) = \lambda \left(\mu_j^2 + (e^{\sigma_j^2} - 1)(1 + \mu_j)^2 \right) \quad (5.15)$$

It should also be noted that the SVJ model nests the SV model, as choosing $\lambda = \mu_j = \sigma_j = 0$ will reduce the SVJ model to the SV model¹⁴. Hence, we would expect the SVJ model to always outperform the SV model in-sample. Out-of-sample, however, its performance is not necessarily superior to the SV model due to the risk of over-parameterization (a hazard that will re-appear as we expand the parameter set even further).

Following the independence between Y_t , J_t and the two Wiener processes, it can be shown (see e.g. Gatheral, 2006 or Zhu, 2000) that the characteristic function of the SVJ model is:

$$\varphi_T^{SVJ}(u) = \varphi_T^{SV}(u) \cdot \varphi_T^J(u) \quad (5.16)$$

where:

¹⁴ In fact, setting $\lambda = 0$ or $\mu_j = \sigma_j = 0$ is sufficient, as both cases eliminate the effect of the jump component.

$$\varphi_T^J = \exp[-\lambda\mu_j iuT + \lambda T((1 + \mu_j)^{iu} \exp(\sigma_j^2 (iu/2)(iu - 1)) - 1)] \quad (5.17)$$

and $\varphi_T^{SV}(u)$ is defined as in (5.4). As in the SV model, vanilla call prices can be obtained using the FFT method and exotic option prices can be calculated using Monte Carlo simulation.

5.3. Multifactor stochastic volatility model (MFSV)

Christoffersen, Heston and Jacobs (2009) propose a two-factor stochastic volatility model as an alternative extension to the Heston (1993) SV model. They argue that the two-factor model is able to capture the time-variation in the volatility smirk better than the one-factor SV model. In particular, this will prove to be effective when the model is estimated using multiple cross-sections of options (Christoffersen, Heston & Jacobs use daily option data during one year for each estimation), as the one factor model will be unable to capture the variation in the slope and level of the volatility smile over time.

In light of the observation that the slope and level of the volatility smile often differ substantially between maturities even in a single cross-section, the multi-factor model will likely provide a better fit even in that setting. Hence, it is of interest to examine if the multi-factor model is able to outperform the SV model also in a one-dimensional cross-section. In particular, the out of sample performance will be of interest, since the addition of parameters might lead to over-parameterization.

We denote the multi-factor stochastic volatility model MFSV and let the following set of SDEs describe the return process under the risk-neutral measure:

$$\frac{dS_t}{S_t} = (r - q)dt + \sqrt{V_t^{(1)}}dW_t^{\mathbb{Q}(1)} + \sqrt{V_t^{(2)}}dW_t^{\mathbb{Q}(2)} \quad (5.18)$$

$$dV_t^{(1)} = \kappa_1 (\theta_1 - V_t^{(1)})dt + \sigma_1 \sqrt{V_t^{(1)}}dW_t^{\mathbb{Q}(3)} \quad (5.19)$$

$$dV_t^{(2)} = \kappa_2 (\theta_2 - V_t^{(2)})dt + \sigma_2 \sqrt{V_t^{(2)}}dW_t^{\mathbb{Q}(4)} \quad (5.20)$$

where the parameters have the same meaning as in (5.1) – (5.3).

The dependence structure is assumed to be as follows:

$$\text{Cov}\left(dW_t^{\mathbb{Q}(1)}, dW_t^{\mathbb{Q}(3)}\right) = \rho_1 dt \quad (5.21)$$

$$\text{Cov}\left(dW_t^{\mathbb{Q}(2)}, dW_t^{\mathbb{Q}(4)}\right) = \rho_2 dt \quad (5.22)$$

$$\text{Cov}\left(dW_t^{\mathbb{Q}(i)}, dW_t^{\mathbb{Q}(j)}\right) = 0, \quad (i, j) = (1, 2), (1, 4), (2, 3), (3, 4) \quad (5.23)$$

In other words, each variance process is correlated with the corresponding Wiener process in the return process, i.e. the diffusion term of which the respective variance process determines the magnitude. The dependence structure also implies that the total variance of the spot return equals the sum of the two variance factors, i.e.

$$\text{Var}_t\left(\frac{dS_t}{S_t}\right) = \left(V_t^{(1)} + V_t^{(2)}\right) dt \quad (5.24)$$

We obtain the characteristic function of the terminal log-stock price in the MFSV model by applying the methodology of Albrecher, Mayer, Schoutens & Tistaert (2006) to the characteristic function presented in Christoffersen, Heston & Jacobs (2009), extending it to allow for a continuous dividend yield q . The result follows by recognizing that the MFSV process (5.18) is the sum of the SV process (5.1) and an additional stochastic volatility term. By the independence of the two Wiener processes in the return process with respect to each other as well as each other's diffusion processes, the added term is independent of the nested SV model return SDE. Since the characteristic function of the sum of two independent variables is the product of their individual characteristic functions, the characteristic function of the MFSV model is determined as:

$$\varphi_T^{\text{MFSV}}(u) = \mathbb{E}_0^{\mathbb{Q}}[e^{ius_T}] = S_0^{iu} f(V_0^{(1)}, V_0^{(2)}, u, T) \quad (5.25)$$

where:

$$f(V_0^{(1)}, V_0^{(2)}, u, T) = \exp\left\{A(u, T) + B_1(u, T)V_0^{(1)} + B_2(u, T)V_0^{(2)}\right\} \quad (5.26)$$

$$A(u, T) = (r - q)iuT + \sum_{j=1}^2 \sigma_j^{-2} \kappa_j \theta_j \left[(\kappa_j - \rho_j \sigma_j iu - d_j)T - 2 \ln\left(\frac{1 - g_j e^{-d_j T}}{1 - g_j}\right) \right] \quad (5.27)$$

$$B_j(u, T) = \sigma_j^{-2} (\kappa_j - \rho_j \sigma_j iu - d_j) \left[\frac{1 - e^{-d_j T}}{1 - g_j e^{-d_j T}} \right] \quad (5.28)$$

$$g_j = \frac{\kappa_j - \rho_j \sigma_j iu - d_j}{\kappa_j - \rho_j \sigma_j iu + d_j} \quad (5.29)$$

$$d_j = \sqrt{(\rho_j \sigma_j iu - \kappa_j)^2 + \sigma_j^2 (iu + u^2)} \quad (5.30)$$

The existence of a closed form characteristic function makes pricing in the MFSV model no more difficult than in the SV and SVJ models. The potential problem, as discussed in context of the SVJ model, arises out of sample as the model might suffer from over-parameterization. It is however important to notice that the MFSV model does not nest the SVJ model. Hence, it is possible for the SVJ model to outperform the MFSV model even in-sample.

5.4. Multifactor stochastic volatility model with jumps (MFSVJ)

As explained in the context the SVJ model, jumps help the model explain the implied probability of large short-term movements in the underlying stock price. Adding jumps thus enables the model to better price far out of the money options with short expiry times. Hence, as jumps serve a different purpose than the additional stochastic volatility factor in the MFSV model, adding jumps might enhance the performance of the MFSV model. Obviously, the jump factor extends the parameter set of the model even further, and the aforementioned potential problem of over-parameterization arises once more, making out-of-sample performance vital for assessing the model's performance.

In the MFSVJ model, the risk-neutral stock price dynamics are described by the following set of SDEs:

$$\frac{dS_t}{S_t} = (r - q - \lambda\mu_J)dt + \sqrt{V_t^{(1)}}dW_t^{\mathbb{Q}(1)} + \sqrt{V_t^{(2)}}dW_t^{\mathbb{Q}(2)} + J_t dY_t \quad (5.31)$$

$$dV_t^{(1)} = \kappa_1 (\theta_1 - V_t^{(1)})dt + \sigma_1 \sqrt{V_t^{(1)}}dW_t^{\mathbb{Q}(3)} \quad (5.32)$$

$$dV_t^{(2)} = \kappa_2 (\theta_2 - V_t^{(2)})dt + \sigma_2 \sqrt{V_t^{(2)}}dW_t^{\mathbb{Q}(4)} \quad (5.33)$$

where all parameters and variables are defined as in equations (5.1) – (5.3) and (5.10). The distributions of J_t and Y_t are log-normal and Poisson, respectively, according to equations (5.10) and (5.13), and the two variables are independent, both of each other and of the four Wiener

processes. The dependence structure between the Wiener processes is the same as in the MFSV model according to equations (5.21) – (5.23).

Given the total spot return variances of the SVJ and MFSV models, the total return variance of the MFSVJ can easily be established as:

$$Var_t \left(\frac{dS_t}{S_t} \right) = \left(V_t^{(1)} + V_t^{(2)} \right) dt + V_{J,t} dt \quad (5.34)$$

where $V_{J,t}$ is defined as in equation (5.15).

Due to the independence between the added jump factor and the SDE of the MFSV model, the characteristic function of s_T is obtained in the same way as in the SVJ model, i.e. as the product of the jump-term characteristic function and the characteristic function of the MFSV model:

$$\varphi_T^{MFSVJ}(u) = \varphi_T^{MFSV}(u) \cdot \varphi_T^J(u) \quad (5.35)$$

where $\varphi_T^{MFSV}(u)$ and $\varphi_T^J(u)$ are defined in (5.25) and (5.17), respectively.

5.5. The Practitioner Black-Scholes model (PBS)

The PBS model originates from local volatility models in which the volatility is described as a deterministic function of time and the underlying stock price. Dumas, Fleming & Whaley (1998) find that local volatility models perform worse than an *ad hoc* method that smoothes implied volatilities from option data and then uses the traditional Black-Scholes pricing formula with the fitted implied volatilities. It is the latter method that is often referred to as the Practitioner Black-Scholes model (PBS), due to its popularity among practitioners. The difference between local volatility models and the PBS model is that the volatility in the PBS model is a function of strike price and time to maturity, rather than the spot price and calendar time. Christoffersen & Jacobs (2004) confirm the PBS models' validity and find that, in their sample, the PBS model actually outperforms the more advanced stochastic volatility model of Heston (1993). Berkowitz (2001) provides a mathematical justification for the use of the PBS model and shows that the PBS model, when re-calibrated sufficiently frequently to a large number of options, will become arbitrarily accurate.

The PBS model is implemented by fitting a deterministic function of strike price and time to maturity to observed implied volatilities in the market. Several functions of different complexity have been proposed, but we will constrain our study to the most general function proposed by Dumas, Fleming & Whaley (1998), also used by Christoffersen & Jacobs (2004):

$$\sigma = \alpha_0 + \alpha_1 K + \alpha_2 K^2 + \alpha_3 T + \alpha_4 T^2 + \alpha_5 KT \quad (5.36)$$

Plain vanilla call and put prices in the PBS model are simply calculated through the standard Black-Scholes formula using the implied volatility obtained from the fitted function (5.36) by inserting the strike price and time to maturity.

As the implied volatility surface is under constant change, the model must be recalibrated at certain time intervals in order assure acceptable accuracy. Due to the straight-forward pricing method using the standard Black-Scholes formula, this is fairly simple and not very computer intensive, and can be done in a matter of minutes, or even seconds, depending on the number of options at hand.

5.6. Previous empirical findings

In Table 1 below, we present a summary of previous studies on the empirical performance of the introduced models. It should be noted that the findings presented in the table are those relevant for the subject of this thesis, and thus not necessarily the main general results of the articles. In the table, the parameter time span refers to the time period used for estimation of the parameters. For example, a model estimated using one day's option data will have daily time span, whereas a model estimated using an option universe from a time period of one year will have an annual time span.

The stochastic volatility model with jumps (SVJ) was, as mentioned, introduced by Bates (1996a), and has been the focus of several succeeding studies. Papers studying jump factors often discuss the importance of jumps in both returns and volatility, where the latter jump factor will increase explanatory power for time varying volatility. Eraker, Johannes & Polson (2003) is the only paper that supports jumps in volatility, while most other papers find this jump factor redundant. Eraker (2004) is the only paper finding both jump factors redundant, while most other papers conclude that the return jump factor increases in-sample performance. However, the effect

on out-of-sample performance is found to be very small. Broadie, Chernov & Johannes (2007) show that the addition of jumps significantly improves the performance of stochastic volatility models when certain parameters are restricted based on historical estimates.

Two possible explanations to why the results regarding the jump factor are different across the previous research are given by Broadie, Chernov & Johannes (2007) and Eraker (2004). Broadie, Chernov & Johannes suggest the fact that the different papers use different sample periods, number of options per cross-section and test statistics, while Eraker points to the difference between using historical returns or option prices for model estimation.

Christoffersen, Heston & Jacobs (2009) show that the MFSV model performs better both in- and out-of-sample than the SV model, indicating that adding additional stochastic volatility factors to the underlying stock price process is desirable. They however argue that the main benefits of adding a second stochastic volatility factor arise when the model is estimated using multiple cross-sections, as the parameter estimates are then required to be valid throughout a varying volatility environment.

To the best of our knowledge, the only study elaborating on models with several stochastic volatility factors as well as jumps is Bates (2000), who however conducts his analysis using annual estimation of the model parameters, consistent with the argumentation of Christoffersen, Heston & Jacobs (2009) that multi-factor models are mainly suited for estimation using multiple cross-sections of options. As expected, the in-sample errors of the multi-factor models in Bates' study are lower than their single-factor counterpart, but he does not perform any out-of-sample analysis from which further conclusions can be drawn. The main conclusion is rather that multi-factor stochastic volatility models and jump models produce more plausible parameter estimates than single-factor stochastic volatility models, indicating that the out-of-sample performance of these models ought to be superior to the SV model.

Dumas, Fleming & Whaley (1998) find that the PBS model outperforms the binomial tree models of Rubinstein (1994) and Derman & Kani (1994), in which the trees are fitted to exactly match observed implied volatilities introducing a severe over-fitting problem. Christoffersen & Jacobs (2004) discuss the importance of the loss function in estimation and evaluation, and use the PBS and SV models to illustrate their point. Their results with respect to the relative performance of

the models are inconclusive and depend on the loss function used for estimation and evaluation, but their results still show that the PBS model is a viable competitor to stochastic volatility models. The authors however do not make any comparison of the models using the implied volatility loss function used in this thesis (presented in the next section).

Table 1
Summary of previous studies

The table below summarizes previous findings on the empirical performance of the models used in this thesis. The findings are the ones relevant for the purpose of this thesis, and not necessarily the main result of each paper.

Paper	Data / Time period	Parameter time span	Findings
Bates (1996a)	Deutsche Mark call and put options (USD) 1984 – 1991	7 years	SVJ more efficient than SV in modeling return distributions. SV cannot explain the volatility smirk, except under implausible parameters.
Bakshi, Cao & Chen (1997)	S&P 500 call options 1988 – 1991	1 day	Stochastic volatility of first importance for model (SV). Further performance improvement when jumps are added (SVJ), especially for short-term options.
Dumas, Fleming & Whaley (1998)	S&P 500 call and put options 1988 – 1993	1 week	PBS has better out-of-sample performance than DVF models that fit observed data exactly. Main reason is over-fitting problems in the DVF approach.
Bates (2000)	S&P 500 call and put options 1988 – 1993	5 years	SV gives implausible parameter values. By adding jumps, more plausible parameters are obtained (for MFSVJ and SVJ). All models exaggerated volatility during the sample period.
Andersen, Benzoni & Lund (2002)	S&P 500 index 1953-1996	1 day	Reasonable descriptive continuous time models must allow for discrete jumps and stochastic volatility (i.e. SVJ or extensions of SVJ).
Pan (2002)	S&P 500 call and put options, and index 1989-1996	1day	Jumps in returns key component to capture the smirk pattern. Jumps in volatility not as important.

Eraker, Johannes & Polson (2003)	S&P 500 and NASDAQ 100 index 1980-1999, 1985-1999	1 day	Jump components important. Including jumps in volatility to return jumps significantly increases performance.
Eraker (2004)	S&P 500 call options and index 1987-1991	1 day	Jumps in both stock price and volatility add little pricing performance compared to simple SV models.
Schoutens, Simons & Tistaert (2003)	EURO STOXX 50 call options 7 Oct. 2003	1 day	SVJ outperforms SV using four different loss functions.
Christoffersen & Jacobs (2004)	S&P 500 call options 1988 – 1991	1 day	Emphasize the importance of being consistent in loss functions when comparing models. Superior performance of PBS and SV depends on loss function.
Broadie, Chernov & Johannes (2007)	S&P 500 call options 1987-2003	1 day	SV with jumps in return improves fit with 50 %. Modest evidence for jumps in volatility.
Christoffersen, Heston & Jacobs (2009)	S&P 500 call options 1990 – 2004	1 year	MFSV outperforms SV with 24% in-sample and 23% out-of-sample. Better results from improvements in modeling of both term structure and skew.

6. Methodology

6.1. Estimation

The first step towards using the models presented above for pricing options is to find optimal parameter values. Not surprisingly, this problem becomes all the more difficult as the number of parameters increases and, in the words of Jacquier & Jarrow (2000), “the estimation method becomes as crucial as the model itself”. A deep discussion of estimation techniques is however more mathematical than financial, and lies beyond the scope of this thesis. Instead, we refer the interested reader to Brito & Ruiz (2004), Renault (1997), and the recently mentioned Jacquier & Jarrow (2000) for a detailed discussion of estimation of stochastic volatility models.

As discussed in the previous section, all models are defined under the risk-neutral measure. Hence, parameter estimates are obtained by calibrating the model to fit observed option prices (i.e. by making the model match observed option prices by altering the parameters). More formally, optimal parameter estimates under the risk-neutral measure are obtained by solving an optimization problem on the form:

$$\hat{\Theta} = \arg \min_{\Theta} \mathcal{L}(\{\hat{C}(\Theta, \Lambda)\}^n, \{C\}^n) \quad (6.1)$$

where Θ is the parameter vector and Λ the vector of spot variances¹⁵. $\{\hat{C}(\Theta, \Lambda)\}^n$ is a set of n option prices obtained from the model, $\{C\}^n$ is the corresponding set of observed option prices in the market and $\mathcal{L}(\cdot)$ is some loss function that quantifies the model's goodness of fit with respect to observed option prices. The most frequently applied loss functions in the literature are the dollar mean squared error (\$ MSE), the percentage mean squared error (% MSE) and the implied volatility mean squared error (IV MSE):

$$\text{\$ MSE}(\Theta, \Lambda) = \frac{1}{n} \sum_{i=1}^n w_i \left(C_i - \hat{C}_i(\Theta, \Lambda) \right)^2 \quad (6.2)$$

$$\text{\% MSE}(\Theta, \Lambda) = \frac{1}{n} \sum_{i=1}^n w_i \left(\frac{C_i - \hat{C}_i(\Theta, \Lambda)}{C_i} \right)^2 \quad (6.3)$$

$$\text{IV MSE}(\Theta, \Lambda) = \frac{1}{n} \sum_{i=1}^n w_i \left(\sigma_i - \hat{\sigma}_i(\Theta, \Lambda) \right)^2 \quad (6.4)$$

where σ_i is the Black-Scholes implied volatility of option i , and $\hat{\sigma}_i(\Theta, \Lambda)$ denotes the corresponding Black-Scholes implied volatility obtained using the model price as input. w_i is an appropriately chosen weight, discussed in more detail below.

The choice of loss function is important and has many implications. The \$ MSE function minimizes the squared dollar error between model prices and observed prices and will thus favor parameters that correctly price expensive options, i.e. deep ITM and long-dated options. The % MSE function, on the other hand adjusts for price level, making it less biased towards correctly pricing expensive options. On the contrary, the % MSE function will put emphasis on options

¹⁵ In our case, as the models are estimated daily, Λ will be a scalar for the SV and SVJ models (i.e. $\Lambda = V_0$). For the MFSV and MFSVJ models, we have that $\Lambda = (V_0^{(1)} V_0^{(2)})$. The PBS model does not incorporate any spot variance term.

with prices close to zero, i.e. deep OTM and short-dated options. The IV MSE function minimizes implied volatility errors, making options with higher implied volatility carry higher importance in the estimation. Due to the shape of the volatility smirk, this will in general put more weight on options with low strike prices, and less weight on options with high strike prices. There will also be a difference in weighting across maturities, depending on the shape of the term structure¹⁶.

The existing literature has focused on the choice of loss function both for evaluation purposes (e.g. Christoffersen & Jacobs, 2004), as well as for computational purposes. The reason for the latter is that most commonly proposed loss functions are non-convex and have several local (and perhaps global) minima, making standard optimization techniques unqualified (Cont & Hamida, 2005). Detlefsen & Härdle (2006) study four different loss functions for estimation of stochastic volatility models and conclude that the most suitable choice once the models of interest have been specified is an implied volatility error metric, as this best reflects the characteristics of an option pricing model that is relevant for pricing out-of-sample. Detlefsen & Härdle also show that this choice leads to good calibrations in terms of relatively good fits and stable parameters. On another technical note, the IV MSE function is sometimes preferred to the \$ MSE and % MSE loss functions also because it does not have the same problems with heteroskedasticity that can affect the estimation (Christoffersen & Jacobs, 2004).

It has also been shown, e.g. by Mikhailov & Nögel (2003), that the choice of weighting (w_i) has a large influence on the behavior of the loss function for optimization purposes, and thus must be chosen with care. Two common methods are to either include the bid-ask spread of the options as a basis for weighting or to choose weights according to the number of options within different maturity categories.

In this thesis, we have chosen to apply an implied volatility mean squared error metric using the effective bid-ask spread as weightings:

¹⁶ See Section 3 for a common shape of the volatility surface, illustrating the relationship between implied volatility and both strike price and maturity.

$$IV\ MSE(\Theta, \Lambda) = \frac{1}{n} \sum_{i=1}^n w_i (\sigma_i - \hat{\sigma}_i(\Theta, \Lambda))^2 \approx \frac{1}{n} \sum_{i=1}^n w_i \left(\frac{C_i - \hat{C}_i(\Theta, \Lambda)}{\mathcal{V}_i^{BS}} \right)^2 \quad (6.5)$$

where $\mathcal{V}_i^{BS}$ denotes the Black-Scholes Vega¹⁷ of option i and $w_i = \frac{1}{ask_i - bid_i} / \sum_j \frac{1}{ask_j - bid_j}$.

The approximation in (6.5), where the pricing error is divided by the Black-Scholes Vega, is obtained by considering the first order approximation:

$$\hat{C}_i(\Theta, \Lambda) \approx C_i + \mathcal{V}_i^{BS} \cdot (\hat{\sigma}_i(\Theta, \Lambda) - \sigma_i) \quad (6.6)$$

Assuming that the first order approximation is fairly accurate¹⁸, we get:

$$\hat{\sigma}_i(\Theta, \Lambda) - \sigma_i \approx \frac{\hat{C}_i(\Theta, \Lambda) - C_i}{\mathcal{V}_i^{BS}} \quad (6.7)$$

Similar methods are used by Christoffersen, Heston & Jacobs (2009), Carr & Wu (2007), Bakshi, Carr & Wu (2008) and Trolle & Schwartz (2008a, 2008b), among others, and significantly reduce computation time¹⁹.

The choice of w_i in (6.5) is logical. If an option is quoted with a wide bid-ask spread, there is less certainty about the true price of the option, and we assign less weight to that observation. The denominator simply rescales the weights to sum to one. A similar approach is implemented by Huang & Wu (2004) who instead account for the bid-ask spread by defining the error between the model price and the true price as zero if the price falls within the bid-ask spread. As mentioned, an additional advantage of the loss function (6.5) is that it is much better behaved than loss functions of squared dollar errors or squared percentage errors, in the sense that the optimization is faster and more stable.

The computational details of the estimation process are described in Appendix E.

¹⁷ Vega is the sensitivity of the option price with respect to volatility in the Black-Scholes model, i.e. $\mathcal{V}_i^{BS} = \partial C_i^{BS} / \partial \sigma_i$.

¹⁸ The accuracy of the approximation is discussed in Appendix F.

¹⁹ The reason for this is that no closed formula exists to calculate Black-Scholes implied volatility. Hence, the implied volatility has to be obtained numerically.

6.2. Evaluation

Evaluation refers to the different measures used for evaluating the models once optimal parameters have been obtained from the estimation procedure. As the focus of this thesis is on pricing performance, relevant metrics will relate to the models' abilities to replicate observed prices in the market.

The first category of measures is referred to as in-sample-errors. As the term implies, the in-sample-errors are calculated as the pricing errors with respect to the options that have been used in the estimation of the models. A natural starting point for this analysis is to consider the error obtained directly from the loss function used to estimate the models, i.e. the implied volatility mean squared error (IV MSE). Furthermore, as the IV MSE loss function was chosen partly with respect to optimization issues, we will not refrain from using the dollar mean squared error (\$ MSE) and percentage mean squared error (% MSE) loss functions (equations (6.2) and (6.3)) in our evaluation of the models. In a sense, this contradicts the results of Christoffersen & Jacobs (2004), who argue that it is essential to use the same loss function for estimation and evaluation. However, their results are based on *evaluating* models using the *same* loss function, when the models have been *estimated* using *different* loss functions. Nevertheless, the results under the loss functions other than the one used also for estimation should be treated with some caution.

The in-sample \$ RMSE and % RMSE were obtained by calculating the respective loss function values using the estimated parameters and spot variances from the IV MSE estimation. We also calculate categorized in-sample errors in a similar fashion, by calculating the value of the loss functions using only the options belonging to each category as input. Note that this means that we do not estimate the model to fit the option prices in the specific category, but merely calculate the pricing error in each category using the parameters obtained from estimating the models to the entire sample.

It is important to keep in mind that some of the models included in the evaluation nest other models, meaning that they include all parameters of the nested model and at least one more. As a consequence, the in-sample errors of the more complex model under the loss function used for estimation will always be less than or equal to the in-sample errors of the nested model, as the more complex model always can be reduced to the simpler form by choosing the additional

parameter values to zero in the optimization procedure. Hence, in-sample-errors will not be able to detect models that suffer from over-fitting, i.e. models that include superfluous parameters. It should be noted, however, that the over-fitting problem mainly arises when the degrees of freedom is small, i.e. when the number of parameters is close to the number of observations. Hence, if the number of observations is large, the models will be less likely to become over-fitted as redundant parameters will bear little or no significance. In order to test for over-fitting, out-of-sample evaluation is conducted.

In the out-of-sample evaluation, we calculate the IV MSE, \$ MSE and % MSE of the models with respect to today's option prices, using parameter estimates from previous days. Hence, the out-of-sample evaluation enables us to draw conclusions as to whether the models are over-fitted, in which case the redundant parameters will affect the out-of-sample errors negatively (as, in that case, the non-zero parameter estimates were only due to variations within the particular sample to which the model was estimated). Out-of-sample errors will, for the loss function used both in estimation and evaluation, by definition be higher than in-sample-errors, as the in-sample errors constitute a lower bound for the specified loss function and the given data sample. One of the most important features of the out-of-sample errors, however, is that a nested model will not necessarily have a higher out-of-sample error than the more complex model. Hence, out-of-sample evaluation constitutes an important complement to in-sample evaluation, in particular when evaluating models of varying complexity.

The out-of-sample errors were obtained by calculating the loss function values using parameter estimates corresponding to estimations one and five days prior to the option prices used as input. Note that *days* here refers to business days, so *five days* most often corresponds to seven days if weekends are included. For the structural models, we follow the method of Christoffersen, Heston & Jacobs (2009) and Huang & Wu (2004) and allow for re-estimation of the spot variance also in the out-of-sample evaluation. Recall that the spot variance is the initial value of the variance process (V_0) and thus only affects the starting value of the variance process, and not the process itself. Hence, V_0 is treated as exogenously given each day, also in the out-of-sample evaluation. The categorized out-of-sample errors were calculated in the same way as the categorized in-sample errors.

7. Data description

The data used for our analysis are European style call options written on the EURO STOXX 50 index during the period January 1st to December 31st 2008. The choice of data is interesting in several ways. First of all, the time period constitutes an exciting period in the financial markets, with volatilities rising to extreme levels subsequent to the crash of Lehman Brothers, making sub-sample analysis and tests of the models' performance with respect to changes in market conditions possible. Secondly, most previous studies have been conducted using data on the S&P 500 index. Although we would not expect our results to differ widely from previous findings, the choice of European data nevertheless constitutes a test of the models' robustness with respect to the underlying asset.

The initial data set, obtained from iVolatility.com²⁰, consists of all quoted call options on the index during 2008. For all 150 946 options in the dataset, we extract information about maturity, strike price, current index level and bid and ask quotes. From the bid and ask prices, we calculate the mid prices as simple averages. Each day we normalize all observations to correspond to an index level of 100. This way, strike prices are easily interpreted in terms of fractions of the spot price, and comparisons of dollar errors between days are not distorted by a changing index level.

To the original data set, we apply a cleaning procedure along the lines of Bakshi, Cao & Chen (1997) and Dumas, Fleming & Whaley (1998), which reduces the number of options to 30 686. The filters include removing options with no traded volume or open interest, options with extremely low prices and options with very high or very low strike prices. The cleaning procedure is described in detail in Appendix D.

²⁰ <http://www.ivolatility.com>

Table 2
Sample characteristics of EURO STOXX 50 call options

The table shows average quoted bid-ask prices for each maturity and moneyness category, together with average bid-ask spread (within brackets) and number of options in each category {in braces}. The sample period extends from January 1st 2008 through December 31st 2008, with a total of 30 686 call options. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)		Days to maturity					All
		< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	0.9496	2.0500	4.3849	7.0004	12.8845	7.3209
		(0.0545)	(0.1003)	(0.1845)	(0.2898)	(0.5594)	(0.3167)
		{937}	{1 631}	{1 853}	{2 291}	{3 728}	{10 440}
	0.94-0.97	1.4716	3.4085	6.4659	9.5319	15.3549	7.6659
		(0.0551)	(0.1084)	(0.1917)	(0.2813)	(0.5694)	(0.2566)
		{919}	{1 084}	{1 042}	{1 105}	{1 235}	{5 385}
ATM	0.97-1.00	2.5667	4.9330	7.9827	11.0068	16.5249	8.8311
		(0.0694)	(0.1223)	(0.1938)	(0.2719)	(0.5648)	(0.2526)
		{923}	{1 054}	{1 024}	{1 013}	{1 112}	{5 126}
	1.00-1.03	4.1662	6.4878	9.5773	12.5389	17.4015	9.8389
		(0.0883)	(0.1308)	(0.2111)	(0.2825)	(0.5776)	(0.2479)
		{848}	{933}	{909}	{941}	{745}	{4 376}
ITM	1.03-1.06	6.3573	8.4973	11.3103	13.9270	19.2363	10.7689
		(0.1408)	(0.1826)	(0.2702)	(0.3009)	(0.6756)	(0.2597)
		{719}	{780}	{817}	{739}	{256}	{3 311}
	1.06-1.10	8.6407	11.0374	13.7096	17.6695	26.9061	13.8129
		(0.1879)	(0.2970)	(0.3469)	(0.4358)	(0.9149)	(0.3702)
		{518}	{451}	{501}	{279}	{222}	{1 971}
All		3.5343	5.0390	7.7858	10.1667	14.9677	8.7855
		(0.0903)	(0.1363)	(0.2158)	(0.2921)	(0.5786)	(0.2828)
		{4 864}	{5 933}	{6 146}	{6 368}	{7 298}	{30 686}

Interest rates and dividend yields are obtained from Datastream. For every day in our sample, we use the expected annual dividend yield as an approximation for the continuous dividend yield of the index. We construct the yield curve every day by linear interpolation between LIBOR quotes of maturities ranging from 1 month to 6 years, in steps of 1 month. For all options with maturity less than one month, we use the 1 month LIBOR rate. The quarterly compounded LIBOR quotes are re-calculated to be continuously compounded according to $r_c = 4 \ln(1 + r_q/4)$, where r_c and r_q denote the continuously and quarterly compounded interest rates, respectively.

Table 2 above shows average mid prices, average bid-ask spread and total number of observations for each category, sorted by moneyness ($F_{t,T}/K$) and maturity. The categorization by moneyness rather than strike price is common practice, and is especially useful in a sample such as ours, with call options with a wide variety of maturities. The usefulness stems from the forward price in the numerator that makes the same moneyness category contain long-dated options with higher strike prices than short-dated options²¹. This makes sense from an economic perspective, as an option one day to maturity and strike price 110 % is much less likely to end up ITM than an option with the same strike price, but one year to maturity.

8. Results

In this section, we present the main results of the empirical study. We start out by presenting the estimated model parameters and discuss their validity. Second, we present the results of the performance evaluation, divided into in- and out-of-sample analysis. Thirdly we conduct a sub-sample analysis, where the data set is divided into high- and low volatility sub-samples. Lastly, we discuss the complications arising when implementing the various models. The four parts are closely connected to the three research questions presented in Section 2. The analysis of the parameter estimates and the performance evaluation aims to answer the question whether increased model complexity enhances model performance, whereas the sub-sample analysis is a comparison of the models' relative performance under varying market conditions. The last part provides an answer to the question of which problems that arise when estimating and implementing the models.

²¹ This holds true if $r > q$, which is the case for the vast majority of options in our sample.

8.1. Parameter estimates

The average parameter estimates and their corresponding standard deviations from the 253 daily estimations are shown in Table 3 below. Beginning with the structural models, several interesting characteristics can be observed. Firstly, the volatility filtering procedure seems to be effective, as the average spot volatilities for the four structural models all lie in the range 29–35 %, with the empirical average implied volatility²² over the 253 days being roughly 27 %.

Furthermore, we note that the correlation between return and volatility is negative in all models. The mean estimates of ρ are in all models between –81 % and –99 %, indicating significant negative skewness in the return distribution. This is in accordance with *a priori* expectations and gives rise to the well-known empirical property that volatility tends to increase in bear markets (Cont, 2001). In terms of options, this implies that the models are able to generate the observed smirk shape in the volatility skew.

The estimated long-run mean of the stochastic variance process (i.e. the long-run mean of V_t) is also reasonable in magnitude for all the models, with an average long-run mean volatility²³ in the interval 22–41 %. The width of the interval is due to the multi-factor models having a higher average long run mean volatility than the single-factor models. This is seemingly the first indication of over-parameterization of the multi-factor models with respect to the sample size, as the θ estimates, especially in the MFSVJ model, are extraordinarily high on some occasions, implying long run mean volatilities of up to 70 %. The high estimates of the long run mean volatility are in all cases a result of one theta estimate being high, whereas the second estimate is close to zero. On average, the values are however similar to the results of Christoffersen, Heston & Jacobs (2009) whose estimates of θ_1 and θ_2 in the MFSV model imply an average long-run mean volatility of 34 % during their 15 year sample period. Considering that the average observed implied volatility in our sample is 27 %, whereas the corresponding number in

²² Bates (1996b) discusses different methods to assess weighted implied volatility. As our data set has been cleaned for options with extreme strike prices, we use the method first introduced by Schmalensee & Trippi (1978) and calculate the average implied volatility each day using equal weights, i.e. $\hat{\sigma}_t = 1/N_t \sum_{i=1}^{N_t} \sigma_i$, where N_t is the total number of option contracts available at time t and σ_i is the implied volatility of option i .

²³ The long-run mean volatility is defined as $\sqrt{\theta}$ and $\sqrt{\theta_1 + \theta_2}$ for the SV and MFSV models, respectively, and as $\sqrt{\theta + \lambda^2 \sigma_f^2 + \mu_f^2 \lambda}$ and $\sqrt{\theta_1 + \theta_2 + \lambda^2 \sigma_f^2 + \mu_f^2 \lambda}$ for the SVJ and MFSVJ models, respectively.

Christoffersen, Heston & Jacob's sample is 19 %, our average long-run mean volatilities of up to 41 % are not extraordinary.

Table 3
Average parameter estimates

The average parameter estimates calculated from the sample of 253 days, together with the corresponding standard deviations (in brackets). For comparative purposes, the parameters of the PBS model have been obtained using the strike price in fractions of the spot price, making the estimates of α_1 and α_5 100 times larger and the estimate α_2 10 000 larger than the corresponding estimates if actual strike prices are used.

	κ	θ	σ	ρ	λ	μ_J	σ_J	V_0
SV	10.5781 (7.9816)	0.0748 (0.0532)	0.8072 (0.3824)	-0.9894 (0.0363)				0.1152 (0.1285)
SVJ	7.5585 (7.0616)	0.0619 (0.0587)	0.6251 (0.3969)	-0.9920 (0.0412)	1.5162 (1.4537)	-0.1165 (0.1509)	0.1800 (0.5122)	0.1067 (0.1300)
MFSV	2.0709 (1.9512)	0.0609 (0.1161)	0.8685 (0.9156)	-0.8925 (0.2148)				0.0569 (0.0852)
	12.1058 (7.8017)	0.0756 (0.1297)	1.1140 (1.0074)	-0.8089 (0.3180)				0.0557 (0.0928)
MFSVJ	1.5751 (1.6531)	0.1477 (0.2034)	1.5929 (1.7757)	-0.9400 (0.1649)	2.2912 (3.3314)	-0.0108 (0.5542)	0.8747 (1.5390)	0.0353 (0.0609)
	11.6033 (7.0161)	0.0207 (0.0429)	1.0408 (0.9279)	-0.9114 (0.2327)				0.0699 (0.0902)
	α_0	α_1	α_2	α_3	α_4	α_5		
PBS	0.4522 (0.6390)	0.1095 (1.1729)	-0.2603 (0.5601)	-0.1803 (0.1855)	0.0157 (0.0334)	0.1215 (0.1023)		

Bates (2000) calibrates (slight variations of) the MFSV and MFSVJ models to a data set of almost 40 000 options on the S&P 500 index and obtains estimates of θ implying long run means of the volatility process in the order of 240 %²⁴ and 130 %, respectively, pointing towards similar problems as encountered in our estimation. Bates however elaborates further with alternative estimation methods and successfully obtains more plausible parameter estimates, indicating that the problem might lie in the estimation technique rather than in the model specification. Bates'

²⁴ Bates uses a different (but equivalent) representation in which θ equals a fraction between the two estimated parameters α and β . However, his estimate of β_1 in the MFSV model is reported as 0.00, making us unable to deduce the estimated $\theta_1 = \alpha_1/\beta_1$. The number above constitutes a lower bound of the long-run mean volatility, assuming $\beta_1 = 0.005$.

analysis is however focused on parameter estimation, and it remains a topic for further research to examine if also the pricing performance can be enhanced through alternative estimation methods.

As for the single-factor models, the estimates of θ is in general slightly higher than corresponding estimates in e.g. Bakshi, Cao & Chen (1997) and Bates (1996a, 2000), but of similar magnitude to Christoffersen, Heston & Jacobs (2009) and Schoutens, Simons & Tistaert (2005). The discrepancy is however natural as the sample periods differ substantially in terms of observed average implied volatility.

Turning to the estimates of the speed of mean reversion (κ) we find that for both the SV and SVJ models, our estimates of the speed of mean reversion, κ , are larger than in other studies. This implies that the risk-premium of volatility risk may be smaller in our sample than in previous studies²⁵. The relationship between volatility risk and speed of mean reversion is straightforward: if the level of volatility is rapidly mean-reverting, then investors will not be as affected by volatility shocks and thus require less risk premium for carrying volatility risk. Looking more closely at the individual estimates of κ , we find that an important cause of the high mean estimates of κ is a few days with very large κ estimates. The extreme values of κ arise from the implementation of the Feller (1951) condition in the estimation procedure, discussed in Appendix D, that ensures that the variance process stays strictly positive. To impose the Feller condition, we estimate the model using the auxiliary variable $\Psi = 2\kappa\theta - \sigma^2$, instead of κ , thereafter calculating κ as $\kappa = (\Psi + \sigma^2)/2\theta$. Hence, estimations with relatively large values of σ and Ψ and a low estimate of θ can result in very large values of κ .

Similar to Christoffersen, Heston & Jacobs (2009) and Bates (2000), our average parameter estimates indicate that one stochastic volatility factor consistently has a higher κ than the other. Our average estimates of κ are however higher than in the two previous studies. In particular, our estimate of κ_1 , i.e. the speed of mean reversion in the more slowly reverting process, is significantly higher than corresponding estimates of Christoffersen, Heston & Jacobs and Bates.

²⁵ The risk premium of volatility risk is commonly defines as $\eta = \kappa^{\mathbb{P}} - \kappa^{\mathbb{Q}}$ (Eraker, 2004). Since we do not estimate $\kappa^{\mathbb{P}}$, we cannot draw any detailed conclusions about the risk-premium, but if the speed of mean-reversion is assumed to be constant under the real-world probability measure (i.e. the actual speed of mean-reversion of volatility), then η is obviously decreasing in $\kappa^{\mathbb{Q}}$.

The estimates of the volatility of the variance process, σ , are also similar to estimates in previous studies, although, as expected, of slightly larger magnitude due to the high volatility, both of the index level and the volatility itself (shown in Figure 4 in Section 8.3., where we discuss the impact of index volatility further). The pattern that the volatility of the variance factor is lower for the volatility factor with the higher speed of mean reversion, found in both Bates (2000) and Christoffersen, Heston & Jacobs (2009) is confirmed in our sample as well, shown by the mean estimates of σ_1 and σ_2 in Table 3. The magnitude of the σ estimates is higher than the estimates in Bates' study. Compared to Christoffersen, Heston & Jacobs, however, our mean estimate σ_1 is smaller and less volatile whereas our estimate of σ_2 is higher. Our results also confirm the previous finding that the absolute value of the correlation is lower for the volatility factor with the higher speed of mean reversion. Christoffersen, Heston & Jacobs suggest that this implies that this volatility factor thus is a less important driver of skewness and kurtosis in the return distribution.

For the two jump models considered, the jump frequency is on average positive and the mean jump size negative. Note however the large standard deviation of the mean jump size component in the MFSVJ model, indicating that positive estimates of the mean jump size is frequently occurring. The pattern is similar to the results in previous studies, although the parameter estimates of the jump factors seems to differ more widely between studies than the parameters of the stochastic volatility factor. For example, the estimates of Bates (1996a) point towards frequently occurring, small jumps ($\lambda = 15.01, \mu_j = -0.001$), whereas the results of Bakshi, Cao and Chen (1997), Eraker (2005) and Schoutens, Simons & Tistaert (2005) instead indicate infrequently occurring, larger jumps with jump parameters $(\lambda, \mu_j, \sigma_j)$ being $(0.59, -0.05, 0.07)$, $(0.50, -0.02, 0.06)$ and $(0.14, 0.18, 0.13)$, respectively.

In order to interpret the specific values of $\alpha_0 - \alpha_5$, recall the implied volatility function of the PBS model:

$$\sigma(K, T) = \alpha_0 + \alpha_1 K + \alpha_2 K^2 + \alpha_3 T + \alpha_4 T^2 + \alpha_5 KT \quad (8.1)$$

Beginning with the intercept α_0 , the most notable feature is the large variation in the sample. Although the mean value is 0.45, α_0 is actually negative for 59 of the 253 days. This is interesting mainly from a Black-Scholes perspective. If actual implied volatilities were constant,

or close to constant, as in the Black-Scholes model, we would anticipate the average α_0 to be positive and lie around the mean implied volatility observed in option prices. As this is obviously not the case, we can easily conclude that the PBS model is able to capture (at least some of) the variation in implied volatility over different strike prices and maturities in the parameters $\alpha_1 - \alpha_5$.

The values of α_1 , α_2 and α_5 show that, on average, implied volatility is decreasing in strike price for short and medium maturities and increasing in strike price for long maturities and high strike prices, which is consistent with the frequently observed downward sloping volatility skew, often most evident for short maturities²⁶. The negative average value of α_2 however contradicts the notion of a volatility smirk, as a negative coefficient for the quadratic term implies that the function is concave with respect to the strike price. In the relevant interval, i.e. for strikes between 80 % and 120 % of the spot price, the concave property is however fairly insignificant.

The coefficients for the time to maturity variables, $\alpha_3 - \alpha_5$ show that, on average, we have a downward sloping and convex term structure for all strike prices. This is consistent with commonly observed patterns in the market, although the volatility term structure tends to exhibit more variation and show a wider range of different shapes than the volatility skew, as the term structure to a larger extent is affected by expectations of the market volatility over different time horizons.

The standard deviations of the PBS model parameter estimates indicate that the variation in all parameters is high throughout the sample period. This pattern is confirmed by Christoffersen & Jacobs (2004), who conclude that the parameter estimates of the PBS model are especially volatile when the model is estimated using an implied volatility loss function. This poses a potential problem, especially for the out-of-sample pricing performance of the model, and we will return to the topic of parameter stability frequently in subsequent sections.

All-in-all, the parameter estimates for the structural models are in line with *a priori* expectations and empirical facts with regards to stock price return behavior as well as the results of previous studies. Hence, the analysis of the parameter estimates verifies the validity of the four structural models, although we find some indications of over-parameterization in the multi-factor models.

²⁶ See Appendix B for an example of the volatility smirk for a range of maturities.

Judging from the parameter estimates, however, the problem does not appear to be severe providing us with good hope that the multi-factor models will prove to be effective in the pricing of options using daily parameter estimation. The estimates of the PBS model indicate that the model is able to capture the slope and level of the volatility surface to some degree, but the high standard errors of the estimates indicate that the model might perform poorly out-of-sample.

To visually illustrate the properties of the models, and to show the models' abilities to generate the desired smirk shape of the volatility skew, as well as their abilities to capture the volatility term structure, we show in Appendix A the implied volatility surfaces generated by the five models, respectively, on the 17th of July 2008 (the same date for which the observed empirical implied volatility surface is shown in Figure 2 in Section 3).

8.2. Pricing performance

As discussed in section 6, we measure pricing performance through the mean squared errors of three loss functions: implied volatility mean squared error (IV MSE), dollar mean squared error (\$ MSE) and relative price mean squared error (% MSE). For clearness, we use the common practice of presenting the root mean squared errors (RMSE) rather than the MSEs, as the RMSEs are measured in the same unit as the variable subject to the loss function (i.e. percentage, dollars and percentage, respectively).

8.2.1. In-sample performance

The average in-sample errors from each of the three loss functions and the corresponding standard deviations are shown in Table 4. To further illustrate the variation over time, we show plots for the in-sample errors for each day in the sample period and each loss function in Appendix A. The results in Table 4 reveal that the in-sample results are decreasing in increased complexity in the structural models, under all loss functions. The pattern was anticipated under the IV loss function, as the more complex models nest the less complex, except for the SVJ model that is not nested by the MFSV model²⁷, whereas the fact that the relation holds under all loss functions is more interesting. The results show that none of the structural models, relative to the other models, loses significant pricing ability measured in \$ or % when estimated to minimize

²⁷ As discussed in section 6, a model that nests another model will always result in a lower in-sample error for the loss function used in both estimation and evaluation, as the more complex model can be reduced to equal the less complex model by choosing the additional parameters to equal zero.

IV MSE. Although, as pointed out by Christoffersen & Jacobs (2004), one should be careful when considering different loss functions in evaluation than in estimation, it makes sense from a practical perspective to consider several loss functions when evaluating the models. If the objective of the model is to calculate a fair price of an option, we would require from a good model that it gives at least reasonable prices for all options, which is equivalent to saying that we require the model to not “blow up” under any loss function.

Table 4
Average in-sample error for each loss function

The table shows the average in-sample errors for all models and loss functions. The figure within brackets is the standard deviation corresponding to the mean value above.

	SV	SVJ	MFSV	MFSVJ	PBS
IV RMSE	1.37 % (0.37 %)	1.29 % (0.36 %)	1.10 % (0.27 %)	1.03 % (0.23 %)	1.73 % (0.80 %)
\$ RMSE	0.3618 (0.0858)	0.3499 (0.0843)	0.3166 (0.0711)	0.2969 (0.0517)	0.5634 (0.3329)
% RMSE	9.56 % (4.32 %)	9.33 % (4.25 %)	8.46 % (4.57 %)	7.28 % (3.75 %)	9.60 % (3.25 %)

Opposite to Christoffersen & Jacobs (2004), we find that the PBS model is outperformed in-sample by all structural models, under all loss functions. The only category in which the PBS model can compete with some of the structural models is under the % MSE loss function, where it produces an in-sample RMSE of similar magnitude to the single-factor structural models. Under the \$ MSE loss function, the PBS model’s in-sample RMSE is almost 100 % worse than the MFSVJ model, and more than 55 % worse than the SV model. Both results indicate that the PBS model is poor in pricing expensive options, as the expensive options carry a higher weight in the \$ MSE loss function, whereas it is fairly good at pricing cheap options that are favored in the % MSE loss function. This indicates that the PBS model lacks some of the flexibility of the structural models, in the sense that minimizing the in-sample error with respect to a chosen loss function causes the model to perform poorly under other loss functions. The degree of this problem depends mostly on the end objective of the model. If the objective is fulfilled through good performance under a specific loss function, then this poses a small problem. As mentioned, however, in a more general setting where we would like a good model to be well-behaved in

several aspects simultaneously the lacking flexibility of the PBS model might be a considerable drawback.

Table 5
t-statistics for in-sample errors

The t-statistics are obtained by comparing the sample means of the RMSEs for all models within each loss function category. A positive t-statistic indicates that the model on the top row has a higher (inferior) sample mean. Values within brackets indicate significance on the 5 % level.

IV RMSE	SV	SVJ	MFSV	MFSVJ	PBS
SV	////	{-2.4579}	{-9.2785}	{-12.1619}	{6.3660}
SVJ	{2.4579}	////	{-6.7194}	{-9.5940}	{7.8601}
MFSV	{9.2785}	{6.7194}	////	{-3.0687}	{11.6934}
MFSVJ	{12.1619}	{9.5940}	{3.0687}	////	{13.1410}
PBS	{-6.3660}	{-7.8601}	{-11.6934}	{-13.1410}	////
\$ RMSE	SV	SVJ	MFSV	MFSVJ	PBS
SV	////	-1.5715	{-6.4466}	{-10.3124}	{9.3280}
SVJ	1.5715	////	{-4.8016}	{-8.5374}	{9.8887}
MFSV	{6.4466}	{4.8016}	////	{-3.5769}	{11.5311}
MFSVJ	{10.3124}	{8.5374}	{3.5769}	////	{12.5851}
PBS	{-9.3280}	{-9.8887}	{-11.5311}	{-12.5851}	////
% RMSE	SV	SVJ	MFSV	MFSVJ	PBS
SV	////	-0.5990	{-2.7777}	{-6.3462}	0.0994
SVJ	0.5990	////	{-2.2187}	{-5.7690}	0.7797
MFSV	{2.7777}	{2.2187}	////	{-3.1892}	{3.2133}
MFSVJ	{6.3462}	{5.7690}	{3.1892}	////	{7.4304}
PBS	-0.0994	-0.7797	{-3.2133}	{-7.4304}	////

Further, we note that the standard deviation of the average IV RMSE of the PBS model is substantially higher than in the structural models, even when adjusting for the higher mean. This implies that the PBS model is more sensitive to the characteristics of the specific daily sample. The relationship is reversed for the % RMSE, for which the PBS model in addition to providing the lowest RMSE also has the lowest standard deviation. Both patterns are easily confirmed by visual inspection of Figure A2 in Appendix A. From the graphs we can clearly see that the PBS model has severe problems during the high volatility period of the fall 2008, an aspect we will return to in the sub-sample analysis in the next section.

Table 5 shows the t-statistics when comparing the sample means of the in-sample errors between the models. A positive t-statistic indicates that the corresponding model on the top row has a

higher (inferior) sample mean than the model in the first column. The brackets indicate that the difference between the means is significant on the 5 % level (two-sided).

The t-statistics shed more light on the difference in performance between the models. Starting with the PBS model, we can conclude that the superior performance of the structural models is significant in all cases but two, namely the differences in % RMSE between the PBS model and the SV and SVJ models.

Moving on to the structural models, both multi-factor models significantly increase in-sample performance compared to the single-factor models, regardless of loss function. In a sense, this contradicts the discussion of Christoffersen, Heston & Jacobs (2009), who argue that under a daily calibration scheme, the benefits of additional stochastic volatility factors should be negligible. It should be pointed out that the fact that a model nests another model does *not* ensure that the in-sample performance will be significantly increased, but merely that the in-sample performance under the same loss function used for estimation cannot be worse than that of the nested model. Adding a supplementary variable with no true explanatory power will only lead to improvements in in-sample performance due to chance and the specific characteristics of the data set. Hence, our results point toward benefits of using multiple stochastic volatility factors, also in estimations using daily cross-sections.

The results also show that adding jumps to the stochastic process of the stock price significantly increases in-sample performance, both for the single- and multi-factor models. For the single-factor models, the difference is however only significant under the IV MSE loss function, whereas the RMSEs of the MFSVJ model are significantly lower than those of the MFSV model under all loss functions considered. An interesting conclusion that can be drawn from these results is that additional stochastic volatility factors should not be seen as substitutes to adding jump components, but rather as complements.

Tables A1 to A3 in Appendix A show the performance of the models under the three loss functions, divided into 42 categories with respect to moneyness and maturity. The tables shed some more light on the performance of the PBS model. It seems that the poor performance of the PBS model stems mostly from extraordinary inferior performance in the pricing of long-dated options, whereas in the short-maturity and far OTM categories, the performance of the PBS

model is actually superior to the SV and SVJ models on some occasions. This indicates that the poor performance of the PBS model revealed in the overall average RMSEs and the corresponding t-tests might be biased by the PBS model's extremely poor performance in pricing options with long maturities. The severity of the problem is determined by the objective of the model. If the objective is to price short- to mid-dated options, then the PBS model is clearly a viable alternative, whereas it is not suited for pricing of long-dated options.

In order to further examine the impact of long-dated options on the performance of the PBS model, we estimated the PBS model excluding all options with more than 1 year to maturity. Using this modified data set, the in-sample performance of the PBS model is drastically enhanced, and the IV RMSE is as low as 0.97 %, i.e. lower than the IV RMSEs of all the structural models from the original estimation. The improvement is however not as significant out-of-sample, where the PBS model has the highest IV RMSE (2.33 %) even in the modified sample. Note however that we did not estimate the structural models using the modified data set, and that a more thorough analysis of the models' relative performances in different sub-sets would require estimating all models using the different sub-sets and analyzing parameter estimates, as well as in- and out-of-sample performance. Such an investigation lies beyond the scope of this thesis, but would be an interesting topic for further studies.

Among the structural models, the categorized in-sample analysis does not add as much new information as for the PBS model. The MFSVJ model has the lowest RMSE in almost all categories, for all three loss functions, followed closely by the MFSV model and the single-factor models. Rather surprisingly, there does not seem to be any distinct trends with regards to the structural model's relative performance in the different categories. Especially, we would have anticipated the jump models to better capture the prices of far OTM options with short maturities (Gatheral, 2006). One explanation for this finding could be that the number of options within the short maturity category (i.e. expiries less than 60 days) that have really short maturities is rather small. The small number of really short-dated options stems from the fact that the span of maturities is discontinuous, with gaps of approximately 30 days. As we have excluded options with less than 6 days to maturity, the shortest dated option in a daily sample will on average have 21 days to expiry.

8.2.2. Out-of-sample performance

Before we present the results from the out-of-sample, some features of out-of-sample analysis must be discussed. Out-of-sample analysis may fill several purposes. For one, it constitutes a test of the stability of the models. This is mainly useful when a model is assumed to be correctly specified, in which case the out-of-sample performance of the model should be enhanced the closer the estimated parameters are to the “true” parameters. From the practical perspective, however, where we assume that the models are misspecified and do not search for any true parameters, the out-of-sample valuation has a different purpose, namely to test if the in-sample performance of the models is due to over-parameterization or if the models are actually capable of capturing the current market conditions and their effect on option prices. Ideally, we would thus like to test the models’ abilities to match market prices of options on the same day (or in the same moment) as the models were estimated. The problem is however that we, at the same time, want to incorporate all available information in the estimation of the models, thus not leaving any un-priced options for out-of-sample evaluation on the same day. Instead, as a proxy for current prices, we test the models’ performance 1 and 5 days out-of-sample, meaning that we evaluate the loss functions by pricing options using model parameters from estimations 1 and 5 days earlier, respectively. From the practical perspective, our main interest lies in the 1-day out-of-sample evaluation, as this is the closest we get to an out-of-sample evaluation under similar market conditions as when the models were estimated, whereas the 5-day out-of-sample should be seen as a complement to give an indication of the robustness of the models.

Table 6 shows the average 1-day and 5-day out-of-sample RMSEs for the five models. Plots showing the development of the RMSEs for all models and loss functions are shown in Figures A3 and A4 in Appendix A. The out-of-sample results reveal several interesting facts.

Looking first at the structural models, we can see that the multi-factor models outperform the single-factor models under all loss functions. This indicates that the multi-factor models, despite having between eight and eleven structural parameters, do not suffer from the possible over-fitting problem. The superior performance of the multi-factor models over the single-factor models is also statistically significant in all cases, indicated by the corresponding t-statistics shown in Table A4 in Appendix A, further supporting the result from the in-sample analysis that

the additional volatility factors of the multi-factor models add to the pricing performance of the models, even in daily cross-sections.

Table 6
Average 1- and 5-day out-of sample errors

The table shows the average out-of-sample errors for all models and loss functions. The figure within brackets is the standard deviation corresponding to the sample mean above.

		SV	SVJ	MFSV	MFSVJ	PBS
IV RMSE	1-day	1.49 %	1.43 %	1.24 %	1.20 %	2.98 %
		(0.44 %)	(0.44 %)	(0.39 %)	(0.38 %)	(2.78 %)
	5-day	1.65 %	1.61 %	1.36 %	1.34 %	4.26 %
		(0.58 %)	(0.60 %)	(0.48 %)	(0.51 %)	(3.81 %)
\$ RMSE	1-day	0.3933	0.3820	0.3492	0.3390	1.0570
		(0.1243)	(0.1240)	(0.1320)	(0.1592)	(1.8995)
	5-day	0.4569	0.4488	0.3836	0.3944	1.3889
		(0.2472)	(0.2486)	(0.1809)	(0.2318)	(2.2255)
% RMSE	1-day	10.55 %	10.42 %	9.37 %	8.45 %	16.77 %
		(4.70 %)	(4.67 %)	(4.72 %)	(4.12 %)	(11.81 %)
	5-day	11.58 %	11.36 %	10.04 %	9.22 %	25.99 %
		(5.31 %)	(5.17 %)	(5.20 %)	(4.38 %)	(18.34 %)

The relationship between the multi-factor models is however inconclusive. The MFSVJ model has a lower average 1-day out-of-sample error under all loss functions, whereas the relationship is reversed for the 5-day out-of-sample RMSE under the \$ MSE loss function. The difference between the models is small in both cases and the corresponding t-statistics shown in Table A4 in Appendix A testify that the differences are insignificant on the 5 % level, with the exception that the MFSVJ model's 1-day out-of-sample % RMSE is significantly lower than the corresponding % RMSE of the MFSV model. The same unambiguous results appear when comparing the RMSEs of the two single-factor models. Although the SVJ model produces lower RMSEs both 1- and 5-days out-of-sample under all loss functions, the difference is not significant on the 5 % level. Hence, our results show that the addition of jumps neither improves, nor worsens the performance of the stochastic volatility models.

As for the PBS model, the out-of-sample results confirm the results from the in-sample evaluation, that the PBS model lacks the flexibility of the structural model and is very sensitive to sample specific characteristics. Especially, we can see that the PBS model has the highest RMSE

in all categories, both 1- and 5-day out-of-sample. The instability of the PBS model is also evident by visual inspection of the plots in Figure A3 and Figure A4 in Appendix A, showing the out-of-sample errors for the whole sample period. On some occasions, the PBS model “blows up”, and produces errors of magnitudes up to eight times larger than the maximum error of any other model. This is also captured by the very large standard deviations of the average RMSEs for the PBS model as compared to the structural models. As in the in-sample evaluation, the large errors of the PBS model stems mostly from poor pricing of long-dated options, as shown in Table A5 to A10 in Appendix A, displaying the 1- and 5-day out-of-sample RMSEs divided into 42 categories by moneyness and maturity.

The main difference between the in- and out-of-sample results is, however, that the out-of-sample RMSEs of the PBS model are higher than for the structural models, also in the short- and mid-maturity categories. Hence, excluding long-dated options from the evaluation merely makes the PBS model improve from catastrophic to poor. This notion is also confirmed by the short sub-set analysis discussed above, where the PBS model was estimated using the same data set, but excluding options with maturity longer than 1 year. In that sub-set, the average out-of-sample performance of the PBS model was still significantly inferior to all four structural models.


As in the in-sample analysis, the structural models do not show any particular patterns in their relative performance with respect to the different moneyness and maturity categories. Hence, the main benefit of extending the single-factor models by adding additional factors seems to be a generally improved in- and out-of-sample performance, rather than improvements in flexibility and the pricing ability of any particular type of options.

8.3. Sub-sample analysis

In this section, we examine the relative performance of the models after dividing the days in the sample into two equal sub-sets, based on implied volatility, simply by choosing the 126 days with the highest average implied volatility to constitute the high-volatility sub-sample and the 126 days with the lowest average implied volatility to make up the low-volatility sub-sample. As volatility is the main driver of option prices, the question whether market conditions in terms of volatility affects the performance of option pricing models bears significant interest.

Figure 4**PBS in-sample IV RMSE and average implied volatility**

The left plot shows the in-sample IV RMSE of the PBS model during the sample period and the right plot shows the average implied volatility for the same period. Note the period of high volatility towards the end of 2008 and the corresponding IV RMSEs of the PBS model.



Comparing the average implied volatility of the options throughout the sample period and the in-sample IV RMSE of the PBS model, shown in Figure 4, we are lead to believe that there, at least for the PBS model, is a close connection between pricing performance and market implied volatility. Looking at Figure A2 in Appendix A, however, the pattern does not seem to be present for the structural models. Table 7 show the pair-wise correlations between the RMSEs for all models and loss functions and the market implied volatility.

Table 7**Pair-wise correlations between in-sample errors and average implied volatility**

The table shows the correlation between the in-sample errors and the average implied volatility on each day, calculated as an arithmetic average over all options observed each day.

	SV	SVJ	MFSV	MFSVJ	PBS
IV RMSE	46.20 %	54.00 %	24.75 %	15.37 %	51.85 %
\$ RMSE	65.18 %	73.56 %	38.22 %	52.64 %	63.44 %
% RMSE	-64.05 %	-66.36 %	-68.04 %	-70.12 %	-42.50 %

Looking at the pair-wise correlations, we can immediately see that the level of implied volatility in the market obviously has a large impact on the RMSEs of the models. Although the results may look surprising at first, especially the very negative correlation between implied volatility and % RMSE, they are in fact natural consequences of the properties of the loss functions. As volatility rises, prices of options go up and as options become more expensive, the squared dollar error will on average increase. The same reasoning explains the negative correlation between implied volatility and % RMSE. As option prices increase, the denominator of the % MSE loss

function increases, resulting in a lower % RMSE²⁸. The effect on the IV RMSE is not as obvious. By the same reasoning that lead to the conclusion that \$ RMSE is increasing in implied volatility due to increasing dollar prices, the IV RMSE should be increasing in implied volatility simply because the numerator of the loss function increases in magnitude. In this respect, it is important to remember that the IV RMSE differs from the other errors since it was used in the loss function used for estimating the models. Hence, the IV RMSE is the only metric of the three that had a chance to adapt properly to the changing market conditions, resulting in lower correlation to the market implied volatility. Note that using a % MSE loss function most likely would have resulted in less negative correlation between implied volatility and RMSE, not because the % RMSE would have been higher during the high volatility period, but because it would have been lower overall. For these reasons, the only relevant measure when comparing the performance of the models between the two sub-samples is the IV RMSE.

Table 8
In-sample IV RMSE for two sub-periods

The t-statistic is calculated by comparing the in-sample means of the IV RMSE for each model in the low- and high-volatility sub-samples. A positive value of the t-statistic indicates that the first low-volatility sub-sample had a higher average IV RMSE. A number in brackets indicates significance on the 5 % level.

IV RMSE	SV	SVJ	MFSV	MFSVJ	PBS
Low volatility	1.24 %	1.19 %	1.08 %	1.03 %	1.42 %
High volatility	1.50 %	1.39 %	1.13 %	1.03 %	2.03 %
t-statistic	{-6.0273}	{-4.7040}	-1.4631	0.0785	{-6.4600}

Table 8 shows the in-sample IV RMSEs of the five models divided into the two sub-samples. The t-statistic is calculated comparing the sample means of the two sub-samples where a negative value indicates a higher sample mean in the high-volatility sub-sample and a value within brackets indicates significance on the 5 % level.

As can be seen, both single-factor models and the PBS model performs worse in the high-volatility sub-sample, whereas the difference is insignificant for the multi-factor models.

²⁸ The numerator also increases due to rising option prices, but not enough to offset the effect of the denominator. This is rather obvious: if option prices on average increase by 10 %, the denominator will consequently also increase by 10 %. The numerator, however, is the difference between the market price and the model price. As the model price adjusts to accommodate the increasing market prices, the numerator will increase only by a fraction of the 10 %.

Interestingly, the MFSVJ model has a lower average IV RMSE in the high-volatility sub-sample, although the difference is very small and statistically insignificant.

To further enlighten the models' dependence on market conditions, Table 9 shows the 1-day out-of-sample IV RMSEs of the five models divided into the two sub-samples. The t-statistics are interpreted as in Table 8.

Table 9
1-day out-of-sample IV RMSE for two sub-periods

The t-statistic is calculated by comparing the 1-day out-of-sample means of the IV RMSE for each model in the low- and high-volatility sub-samples. A positive value of the t-statistic indicates that the first low-volatility sub-sample had a higher average IV RMSE. A number in brackets indicates significance on the 5 % level.

IV RMSE	SV	SVJ	MFSV	MFSVJ	PBS
Low volatility	1.32 %	1.28 %	1.14 %	1.12 %	2.00 %
High volatility	1.67 %	1.58 %	1.34 %	1.27 %	3.96 %
t-statistic	{-6.8815}	{-5.6935}	{-4.0239}	{-3.2603}	{-5.9572}

Out-of-sample, all models yield higher IV RMSEs in the high-volatility sub-sample. One important difference between the structural models and the PBS model in this respect is that the out-of-sample error calculation for the structural models involve filtering the spot volatility to match the correct day. Hence, the structural models have a certain flexibility to adjust for changes in the volatility level, whereas the estimates of the PBS model are static, making the PBS model unable to compensate even for parallel shifts in volatility from 1 day to another. The larger IV RMSE of the PBS model during the high volatility period arises rather as a consequence of more frequent and larger parallel shifts in the volatility surface, rather than a high level of volatility. This idea is also confirmed by noting that the correlation between the 1-day out-of-sample IV RMSE of the PBS model and the absolute value of the first difference of the average implied volatility²⁹ is as high as 69 %.

Interestingly, the corresponding correlations for the structural models are also relatively high, between 46 % and 53 %. One possible explanation for this is that the volatility filtering is ineffective, and unable to adjust the model to changes in spot volatility. A more plausible explanation, however, is that large changes in volatility make traders re-evaluate more parameters

²⁹ The absolute value of the first-difference of the average implied volatility is defined as $|\Delta \bar{\sigma}_t| = |\bar{\sigma}_t - \bar{\sigma}_{t-1}|$. We use the absolute value, as we are only interested in the magnitude of the change in average implied volatility, disregarding if there is an upward or downward shift.

than the spot volatility, and adjust their prices (i.e. the market prices we observe in our sample) accordingly. This makes sense from an economical perspective, as *ceteris paribus* mainly is a theoretical concept, and changes in fundamental economic variables such as volatility most often are the cause or effect of changes in related variables.

8.4. Estimation and implementation issues

The first issue arising when expanding the complexity of structural models is that the estimation procedure becomes more difficult. As the number of model parameters increase, the pricing function becomes more complex leading to increased complexity in minimizing the loss function through non-linear optimization. To mitigate this problem, the optimization of the multi-factor models requires more attention than the less parameterized single-factor models. The practical implications of this on the estimation procedure are three-fold. First, each estimation is more costly numerically the more parameters are included in the models, significantly increasing computation time. Second, as the optimization is less stable, we are required to run the optimization with an increased number of starting values for the parameter vector, in order to ensure that the local optimizer does not get stuck in a local minimum, far from the desired global minimum. Third, the number of iterations required between the volatility filtering and the parameter estimation is on average higher the more parameters are included. All three aspects make the estimation of the multi-factor models more time consuming than the single-factor models. In this respect, the PBS model has an obvious advantage. As the pricing of options in the PBS model is carried out by evaluating an analytical formula, the estimation of the PBS model is significantly faster than the structural models.

To quantify the differences we estimate all models to option prices observed on July 17th 2008 using ten different starting values. The starting values are chosen randomly on a uniformly distributed interval of $\mu_i \pm \sigma_i$, where μ_i and σ_i denote the mean and standard deviation of the parameter estimates for the whole sample period. Table A11 in Appendix A shows the average parameter estimates and their corresponding standard deviations, the mean IV RMSE, its standard deviation and the average computation time³⁰.

³⁰ The estimation was carried out in MATLAB using a computer with an Intel Core 2 Duo 2.10 GHz processor.

Beginning with the PBS model, we note the exceptional difference in estimation speed. Whereas the structural models take on average between 1 and 10 minutes per estimation, each estimation in the PBS model takes on average half a second. Notable is also that the standard deviations of the parameter estimates of the PBS model are very large in contrast to the low standard deviation of the corresponding IV RMSE. This implies that the IV MSE loss function when applied to the PBS model has several local minima, with widely differing parameter estimates, but with loss function values close to the global minimum.

The structural models indicate similar patterns, however not to the extent of the PBS model. The most interesting observation in the structural models is perhaps the estimate of κ in the SVJ model. The mean value of 9.23 and standard deviation of 20.02 reported in Table A11 does not tell the whole story. It turns out that 9 of the 10 κ estimates lie in the narrow range of 2.83 – 3.04, whereas the 10th estimate is 66.21. This illustrates the problem with local optimizers, as they run a risk of returning values from a local minimum far from the global minimum. The problem does not only affect the parameter estimates, but also distorts the loss function values. For example, the estimation providing the extraordinary estimate of κ results in an IV RMSE over 50 % higher than the average over the other 9 estimations. The SV model shows similar, but not as severe, problems. Whereas most of the κ estimates of the SV model lie around 2, the mean value of κ over the 10 estimations is 7.92 due to some estimates around 15. These estimates do however not seem to affect the IV RMSEs of the SV model, that lie in the range 1.12 % to 1.14 % for all 10 estimations. It should be pointed out that all other parameter estimates of the single-factor models are reasonably well behaved. Also for the multi-factor models, the parameter causing most problems is κ . The problem is most evident in the MFSVJ model, where the estimate of κ_1 has a mean value of 3.89 and a standard deviation of 7.70 over the 10 estimations. The high standard deviation stems from two estimates of 18.47 (i.e. the same value, independent of each other as the starting values were randomly chosen on a defined interval), whereas the mean value of the remaining 8 estimations is merely 0.24. In contrast to the SV model, the different parameter estimates of the MFSVJ model have a large impact on the IV RMSE. In both cases where $\kappa_1 = 18.47$, the IV RMSE is over 60 % higher than the lowest observed IV RMSE over the 10 estimations. This illustrates the fact that great care has to be taken when estimating the multi-factor models, as erroneous estimations can significantly affect the results.

The jump factor parameters of both the SVJ and MFSVJ models are surprisingly well-behaved with low standard errors. The main drawback of adding jumps from an estimation perspective instead seems to be that the estimation time is much higher for the jump models. The increased estimation time is however a natural result of expanding the parameter set, as it is obviously a more computer intensive task to optimize a function of 11 variables (MFSVJ) instead of 4 (SV). The high estimation time of the MFSVJ model, being on average almost 11 minutes per estimation, in addition to the difficulties arising in the resulting parameters, is potentially a drawback of the MFSVJ model. This is however mostly a problem for evaluating purposes, where we have to estimate the models to a large number of days. In practice, the severity of the problem would depend on the re-estimation frequency required to obtain desired accuracy. Our evaluation results however show that daily estimation is sufficient to obtain small pricing errors. Hence, although 10 or more estimations would be necessary to ensure that reliable parameter estimates are obtained, the fact that only one parameter set needs to be obtained, an estimation time of 1–2 hours does not pose a serious problem. It should also be pointed out that the calibration time depends on both the software and hardware used and the estimation method and that it might be possible to significantly reduce computation time through the use of different methods and applications.

9. Conclusions

This paper evaluates the performance of four structural stochastic volatility models (SV, SVJ, MFSV and MFSVJ) and one *ad-hoc* Black-Scholes benchmark model (PBS). Our results show that the parameter estimates of all four structural models are plausible and in line with previous research. This is interesting for several reasons. First, this shows that multi-factor models generate similar parameter estimates when estimated to daily cross-sections of data as when estimated to multiple cross-sections, as in previous studies by Bates (2000) and Christoffersen, Heston & Jacobs (2009). Given the superior performance of the multi-factor models, shown especially by Christoffersen, Heston & Jacobs, the parameter estimates indicate that the multi-factor models have good chance in outperforming single-factor models also in single cross-sections. The parameter estimates of the PBS model are more volatile than in the structural models, indicating that the PBS model might have difficulties pricing options out-of-sample.

The in-sample evaluation confirms the validity of the multi-factor models in single-cross sections, and the multi-factor models produce statistically significant lower average in-sample errors than the single-factor models, regardless of loss function used in evaluation. Further, the single-factor stochastic volatility models significantly outperform the PBS model that has the worst performance under all loss functions. The performance of the PBS model is however distorted by extremely poor performance in the pricing of long-dated options, and in a sub-sample excluding options with more than one year to maturity, the in-sample IV RMSE of the PBS model is actually lower than the RMSEs of the structural models from the full-sample estimation. The improved performance of the PBS model in the short-maturity sub-sample is however not persistent out-of-sample, where the 1-day IV RMSE again is higher than the IV RMSE of the structural models. The short-maturity sub-sample investigation in this thesis should however be considered a “back of an envelope” analysis, and a more thorough study of the PBS model in different sub-samples would be necessary to draw more well-founded conclusions.

The hierarchy among the models remains out-of-sample, where the multi-factor models again produce significantly superior out-of-sample errors compared to the single-factor models, both 1- and 5-day out-of-sample under all loss functions. These results indicate that the multi-factor models do not suffer from the potential over-fitting problem, and further adds to the conclusion that the addition of a second stochastic volatility factor is useful also in single cross-sections. The out-of-sample performance of the PBS model is, as indicated by the volatile parameter estimates, significantly inferior to the performance of the structural models. When considering the out-of-sample performance of the PBS model it is however important to keep in mind the purpose of the out-of-sample evaluation from a practical perspective. As our aim is not to find a correctly specified model and estimate the “true” parameters, the out-of-sample evaluation is here used as a proxy for pricing options the same day as the estimation was carried out, as accurately as possible. Hence, as conditions may change significantly from one day to another, the 1-day out-of-sample evaluation may be a rather poor proxy of the model’s performance out-of-sample in the very moment it was estimated. Hence, a more fair evaluation of the PBS model might be to leave a number of options in each cross-section out of the estimation, and consider the pricing errors of these options for out-of-sample evaluation.

The impact of adding jumps to the return process is not as unambiguous as the additional stochastic volatility factor. In-sample, the performance of the MFSVJ model is significantly superior to the MFSV model under all loss functions, whereas the SVJ model only significantly outperforms the SV model under the IV MSE loss function. Out-of-sample, however, the results are different, and the only significant improvement in out-of-sample errors is the 1-day % RMSE of the MFSVJ model compared to the MFSV model. One explanation for the poor performance of the jump models in our sample is that we have excluded options with very short maturities (less than 6 days). Hence, the main advantage of jump models that they are able to price the implied probabilities of large short-term moves in observed option prices is partially lost. Including such options however comes at the cost of a risk of introducing liquidity biases in the observed option prices, as prices of very short-dated options can be affected by traders forced to buy or sell large positions.

In a volatility based sub-sample analysis, we show that the performance of the models as measured by IV RMSE is worse in the high volatility sub-sample. The difference is especially significant for the PBS model, whose 1-day out-of-sample IV RMSE is almost 100 % higher during the high volatility sub-sample. It should be noted that a certain increase in IV RMSE was expected in the high volatility sub-sample, as the difference between model implied volatility and observed implied volatility naturally will be higher the larger the magnitude of the two quantities. One explanation for the poor out-of-sample performance of the PBS model out-of-sample is that the out-of-sample errors for the PBS model are calculated without any volatility filtering to adjust for changing levels in volatility. Hence, the out-of-sample errors of the PBS model will be very high in periods where the volatility of the volatility is high, as parallel shifts in the volatility surface will have a much larger impact on the out-of-sample performance of the PBS model compared to the structural models.

We also show that the estimation of the multi-factor models, especially the MFSVJ model, is less stable and more time consuming than the other models. In this respect, the PBS model is outstanding with an average calibration time of less than a second, as compared to 10 minutes in the MFSVJ model. This fact is of great importance in the evaluation of the models. As has been shown, the out-of-sample performance of the PBS model deteriorates significantly the further out-of-sample the model is evaluated. The extremely fast estimation of the PBS model however

ensures that the model can be re-estimated frequently to match current market conditions even on a minute-to-minute basis, making it all the more interesting to undertake more studies of the PBS model with different evaluation techniques than the standard out-of-sample evaluation.

10. References

- Albrecher, Hansjörg, Philip Mayer, Wim Schoutens and Jurgen Tistaert, 2006, The Little Heston Trap, Technical Report, Katholieke Universiteit Leuven.
- Andersen, Torben G., Luca Benzoni and Jesper Lund, 2002, An Empirical Investigation of Continuous-time Models for Equity returns, *Journal of Finance* 57, 1238-1284.
- Applebaum, David, 2004, Lévy Processes – From Probability to Finance and Quantum Groups, *Notices of the American Mathematical Society* 51, 1336-1347.
- Bakshi, Gurdip, Peter Carr and Liuren Wu, 2008, Stochastic Risk Premiums, Stochastic Skewness in Currency Options and Stochastic Discount Factors in International Economics, *Journal of Financial Economics* 87, 132-156.
- Bakshi, Gurdip, Charles Cao and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *The Journal of Finance* 52, 2003-2049.
- Bates, David S., 1996a, Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, *The Review of Financial Studies* 9, 69-107.
- Bates, David S., 1996b, Testing Option Pricing Models, in Maddala, G.S. and C.R. Rao, eds, *Handbook of Statistics*, Vol. 14, *Statistical Methods in Finance*, Elsevier, Amsterdam, 567-611.
- Bates, David S., 2000, Post-87 Crash Fears in S&P 500 Futures Options, *Journal of Econometrics* 94, 181-238.
- Berkowitz, Jeremy, 2001, Getting the Right Option Price with the Wrong Model, Working paper, University of California.
- Björk, Tomas, 2004, *Arbitrage Theory in Continuous Time*, 2nd Edition, Oxford University Press, Oxford.
- Black, Fischer, 1976, Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association*, 177-181.
- Black, Fischer and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-659.
- Borak, Szymon, Kai Detlefsen and Wolfgang Härdle, 2005, FFT Based Option Pricing, SFB 649 Discussion Paper, Humboldt University Berlin.
- Broadie, Mark, Mikhail Chernov and Michael Johannes, 2007, Model Specification and Risk Premia: Evidence from Futures Options, *Journal of Finance* 62, 1453-1490.
- Broto, Carmen and Esther Ruiz, 2004, Estimation for Stochastic Volatility Models: A survey, *Journal of Economic Surveys* 18, 613-649.

Carr, Peter, Eric C. Chang and Dilip B. Madan, 1998, The Variance Gamma Process and Option Pricing, *European Finance Review* 2, 79-105.

Carr, Peter, Hélyette Geman, Dilip B. Madan and Marc Yor, 2001, Stochastic Volatility for Lévy Processes, EFA 2002 Berlin Meetings Presented Paper.

Carr, Peter and Dilip B. Madan, 1999, Option Valuation Using the Fast Fourier Transform, *Journal of Computational Finance* 2, 61-73.

Carr, Peter and Liuren Wu, 2007, Stochastic Skew in Currency Options, *Journal of Financial Economics* 86, 213-247.

Christoffersen, Peter and Kris Jacobs, 2004, The Importance of the Loss Function in Option Pricing, *Journal of Financial Economics* 72, 291-318.

Christoffersen, Peter, Steven Heston and Kris Jacobs, 2009, The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work so Well, Working Paper, McGill University.

Cont, Rama, 2001, Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues, *Quantitative Finance* 1, 223-236.

Cont, Rama and Sana Ben Hamida, 2005, Recovering Volatility from Option Prices by Evolutionary Optimization, *Journal of Computational Finance* 8, 43-76.

Cont, Rama and Peter Tankov, 2004, *Financial Modeling With Jump Processes*, Chapman & Hall/CRC Press, Boca Ranton.

Cox, John C. and Stephen A. Ross, 1976, The Valuation of Options for Alternative Stochastic Processes, *Journal of Financial Economics* 3, 145-166.

Cox, John C., Stephen A. Ross and Mark Rubinstein, 1979, Option Pricing: a Simplified Approach, *Journal of Financial Economics* 7, 229-263.

Derman, Emanuel and Iraj Kani, 1994, The Volatility Smile and Its Implied Tree, Goldman Sachs Quantitative Strategies Research Notes.

Detlefsen, Kai and Wolfgang Härdle, 2006, Calibration Risk for Exotic Options, SFB 649 Discussion Paper, Humboldt University Berlin.

Dumas, Bernard, Jeff Fleming and Robert E. Whaley, 1998, Implied Volatility Functions: Empirical Tests, *The Journal of Finance* 53, 2059-2106.

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

Eraker, Bjørn, 2004, Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, *The Journal of Finance* 59, 1367-1404.

Eraker, Bjørn, Michael Johannes and Nicholas Polson, 2003, The Impact of Jumps in Volatility and Returns, *Journal of Finance* 58, 1269-1300.

- Feller, William, 1951, Two Singular Diffusion Problems, *The Annals of Mathematics* 54, 173-182.
- Figlewski, Stephen and Xiaozu Wang, 2000, Is the 'Leverage Effect' a Leverage Effect?, Working Paper, NYU Stern School of Business.
- Gatheral, Jim, 2004, A Parsimonious Arbitrage-free Implied Volatility Parameterization Application to the Valuation of Volatility Derivatives, Merrill Lynch Global Derivatives & Risk Management.
- Gatheral, Jim, 2006, *The Volatility Surface*, John Wiley & Sons, New Jersey.
- Gut, Allan, 1995, *An Intermediate Course in Probability*, Springer, New York.
- Harrison, Michael J. and David M. Kreps, 1979, Martingales and Arbitrage in Multiperiod Securities Markets, *Journal of Economic Theory* 20, 381-408.
- Heston, Steven L., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Application to Bond and Currency Options, *The Review of Financial Studies* 6, 327-343.
- Huang, Jing-Zhi and Liuren Wu, 2004, Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes, *Journal of Finance* 59, 1405-1440.
- Hull, John and Alan White, 1987, The Pricing of Options on Assets With Stochastic Volatilities, *Journal of Finance* 42, 281-300.
- Hull, John, 2006, *Options, Futures and Other Derivatives*, 6th Edition, Pearson Prentice Hall, New Jersey.
- Jacquier, Eric and Robert Jarrow, 2000, Bayesian Analysis of Contingent Claim Model Error, *Journal of Econometrics* 94, 145-180.
- Johnson, Herb and David Shanno, 1987, Option Pricing when the Variance is Changing, *The Journal of Financial and Quantitative Analysis* 22, 143-151.
- Kahl, Christian and Peter Jäckel, 2005, Not-So-Complex Logarithms in the Heston Model, *Wilmott Magazine* 5, 94-103.
- Kahl, Christian and Roger Lord, 2006, Why the Rotation Count Algorithm Works, Working Paper, University of Wuppertal.
- Lee, Roger W., 2004, Option Pricing by Transform Methods: Extensions, Unification and Error Control, *Journal of Computational Finance* 7, 51-86.
- Melino, Angelo and Stuart M. Turnbull, 1990, Pricing Foreign Currency Options with Stochastic Volatility, *Journal of Econometrics* 45, 239-265.
- Merton Robert C., 1976, Option Pricing when Underlying Stock Returns are Discontinuous, *Journal of Financial Economics* 3, 125-144.


- Mikhailov, Sergei and Ulrich Nögel, 2003, Heston's Stochastic Volatility Model Implementation, Calibration and Some Extensions, *Wilmott Magazine* 4, 74-79.
- Pan, Jun, 2002, The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics* 63, 3-50.
- Rebonato, Riccardo, 1999, *Volatility and Correlation in the Pricing of Equity, FX and Interest-Rate Options*, John Wiley & Sons, Chichester.
- Renault Eric, 1997, Econometric Models of Option Pricing Errors, *Advances in Economics and Econometrics: Theory and Applications: Seventh World Congress*. Cambridge University Press 3 223-78.
- Rubinstein, Mark, 1994, Implied Volatility Trees, *Journal of Finance* 49, 771-818.
- Schmalensee, Richard and Robert R. Trippi, 1978, Common Stock Volatility Expectations Implied by Option Premia, *Journal of Finance* 33, 129-147.
- Schoutens, Wim, Erwin Simons and Jurgen Tistaert, 2004, A Perfect Calibration! Now What?, *Wilmott Magazine* 2, 66-78.
- Schöbel, Rainer and Jianwei Zhu, 1999, Stochastic Volatility with an Ornstein-Uhlenbeck process: An Extension, *European Finance Review* 3, 23-46.
- Scott, Louis O., 1987, Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods, *Mathematical Finance* 7, 413-424.
- Stein, Elias M. and Jeremy C. Stein, 1991, Stock Price Distributions with Stochastic Volatility: An Analytic Approach, *The Review of Financial Studies* 4, 727-752.
- Trolle, Anders B. and Eduardo S. Schwartz, 2008a, A General Stochastic Volatility Model for the Pricing of Interest Rate Derivatives, Working Paper, UCLA.
- Trolle, Anders B. and Eduardo S. Schwartz, 2008b, Unspanned Stochastic Volatility and the Pricing of Commodity Derivatives, EFA 2008 Athens Meeting Paper.
- Wiggins, James B., 1987, Option Values under Stochastic Volatility: Theory and Empirical Estimates, *Journal of Financial Economics* 19, 351-372.
- Zhang, Peter G., 1998, *Exotic Options*, 2nd Edition, World Scientific Publishing, Singapore.
- Zhu, Jianwei, 2000, *Modular Pricing of Options*, Springer, Heidelberg.



Appendix A: Figures and tables

Figure A1

Implied volatility surfaces on July 17th, 2008

The plots show the implied volatility surface for the five models evaluated. Note that the implied volatility has the same scale in all plots to enhance comparability. All models except the PBS show a similar pattern of volatility skew and term structure.









Figure A2

In-sample RMSEs for all models and loss functions

Each model is evaluated under each loss function for each day in the sample period. The first column represents IV RMSEs, the second \$ RMSEs and the third % RMSEs. Note the high IV RMSE and \$ RMSE of the PBS model compared to the structural models during the high volatility period towards the end of 2008.




Table A1

In-sample RMSEs per category using IV MSE loss function

The table shows the average in-sample errors using the IV RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	0.99 %	0.86 %	1.07 %	1.32 %	0.72 %	1.14 %
		SVJ	0.96 %	0.90 %	1.09 %	1.29 %	0.86 %	1.12 %
		MFSV	0.79 %	0.89 %	0.86 %	1.05 %	0.76 %	1.00 %
		MFSVJ	0.58 %	0.76 %	0.87 %	0.94 %	0.71 %	0.86 %
		PBS	0.73 %	1.07 %	1.72 %	1.22 %	2.16 %	1.34 %
	0.94-0.97	SV	0.95 %	0.97 %	1.18 %	1.22 %	0.70 %	1.12 %
		SVJ	0.85 %	0.89 %	1.17 %	1.22 %	0.85 %	1.06 %
		MFSV	0.88 %	0.88 %	1.10 %	1.03 %	0.86 %	1.05 %
		MFSVJ	0.64 %	0.74 %	1.09 %	0.92 %	0.78 %	0.87 %
		PBS	1.00 %	1.23 %	1.88 %	1.02 %	1.93 %	1.38 %
	1.00-1.03	SV	0.86 %	1.27 %	1.37 %	1.14 %	0.68 %	1.19 %
		SVJ	0.74 %	1.11 %	1.36 %	1.16 %	0.86 %	1.11 %
		MFSV	0.77 %	1.00 %	1.37 %	1.02 %	0.89 %	1.09 %
		MFSVJ	0.60 %	0.88 %	1.35 %	0.92 %	0.82 %	0.95 %
		PBS	1.24 %	1.42 %	2.11 %	1.00 %	2.00 %	1.60 %
	1.03-1.06	SV	1.07 %	1.68 %	1.59 %	1.11 %	0.70 %	1.45 %
		SVJ	0.95 %	1.47 %	1.57 %	1.13 %	0.87 %	1.34 %
		MFSV	0.74 %	1.17 %	1.65 %	1.05 %	0.93 %	1.21 %
		MFSVJ	0.66 %	1.10 %	1.60 %	0.95 %	0.85 %	1.12 %
		PBS	1.47 %	1.67 %	2.36 %	1.05 %	1.76 %	1.82 %
ATM	1.06-1.10	SV	1.66 %	2.25 %	1.77 %	1.11 %	0.65 %	1.88 %
		SVJ	1.56 %	2.01 %	1.74 %	1.10 %	0.79 %	1.76 %
		MFSV	0.96 %	1.53 %	1.88 %	1.07 %	0.89 %	1.48 %
		MFSVJ	1.00 %	1.49 %	1.81 %	0.96 %	0.79 %	1.45 %
		PBS	2.15 %	2.06 %	2.65 %	1.19 %	1.49 %	2.26 %
	All	SV	1.99 %	2.85 %	1.89 %	1.20 %	0.69 %	2.35 %
		SVJ	2.01 %	2.59 %	1.84 %	1.14 %	0.74 %	2.22 %
		MFSV	1.28 %	2.01 %	2.02 %	1.13 %	0.78 %	1.89 %
		MFSVJ	1.23 %	1.96 %	1.93 %	1.06 %	0.71 %	1.81 %
		PBS	2.80 %	2.55 %	2.87 %	1.32 %	1.36 %	2.83 %
	All	SV	1.26 %	1.51 %	1.47 %	1.27 %	0.78 %	1.37 %
		SVJ	1.18 %	1.38 %	1.45 %	1.25 %	0.88 %	1.29 %
		MFSV	0.97 %	1.17 %	1.45 %	1.11 %	0.86 %	1.16 %
		MFSVJ	0.80 %	1.07 %	1.41 %	1.00 %	0.83 %	1.03 %
		PBS	1.53 %	1.56 %	2.21 %	1.22 %	2.06 %	1.73 %

Table A2

In-sample RMSEs per category using \$ MSE loss function

The table shows the average in-sample errors using the \$ RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	0.0689	0.1514	0.3503	0.6132	0.4925	0.3582
		SVJ	0.0654	0.1569	0.3509	0.5933	0.6088	0.3552
		MFSV	0.0509	0.1607	0.2775	0.4883	0.5638	0.3171
		MFSVJ	0.0400	0.1396	0.2789	0.4300	0.5373	0.3022
		PBS	0.0544	0.1875	0.5384	0.5884	1.6023	0.5712
	0.94-0.97	SV	0.0754	0.2016	0.3988	0.5761	0.4828	0.3123
		SVJ	0.0651	0.1873	0.3941	0.5733	0.6040	0.3156
		MFSV	0.0676	0.1818	0.3685	0.4891	0.6279	0.2916
		MFSVJ	0.0510	0.1573	0.3641	0.4335	0.5750	0.2692
		PBS	0.0847	0.2456	0.6170	0.4970	1.4067	0.4840
ATM	0.97-1.00	SV	0.0915	0.2795	0.4620	0.5399	0.4770	0.3444
		SVJ	0.0767	0.2472	0.4574	0.5443	0.6038	0.3451
		MFSV	0.0816	0.2206	0.4619	0.4844	0.6420	0.3290
		MFSVJ	0.0654	0.1983	0.4531	0.4309	0.5943	0.3055
		PBS	0.1203	0.3062	0.7007	0.4810	1.4627	0.5536
	1.00-1.03	SV	0.1235	0.3658	0.5256	0.5177	0.4884	0.3993
		SVJ	0.1068	0.3221	0.5192	0.5246	0.6098	0.3945
		MFSV	0.0874	0.2614	0.5453	0.4895	0.6620	0.3853
		MFSVJ	0.0775	0.2462	0.5302	0.4379	0.6026	0.3583
		PBS	0.1536	0.3619	0.7770	0.4898	1.2280	0.5735
	1.03-1.06	SV	0.1682	0.4681	0.5652	0.5046	0.4341	0.4555
		SVJ	0.1532	0.4184	0.5588	0.4996	0.5224	0.4412
		MFSV	0.0944	0.3251	0.6029	0.4831	0.5934	0.4324
		MFSVJ	0.0984	0.3146	0.5835	0.4356	0.5298	0.4075
		PBS	0.2043	0.4236	0.8474	0.5356	0.9812	0.6128
ITM	1.06-1.10	SV	0.1578	0.5488	0.5718	0.5253	0.4446	0.5334
		SVJ	0.1518	0.4967	0.5607	0.4981	0.4696	0.5032
		MFSV	0.0863	0.3903	0.6188	0.4897	0.4910	0.4913
		MFSVJ	0.0865	0.3777	0.5949	0.4636	0.4459	0.4595
		PBS	0.2122	0.4811	0.8813	0.5694	0.8393	0.6810
	All	SV	0.1126	0.3089	0.4776	0.5914	0.5462	0.3618
		SVJ	0.1013	0.2800	0.4719	0.5800	0.6190	0.3499
		MFSV	0.0814	0.2414	0.4715	0.5171	0.6080	0.3294
		MFSVJ	0.0703	0.2206	0.4578	0.4632	0.6011	0.2969
		PBS	0.1355	0.3142	0.7096	0.5736	1.5082	0.5634

Table A3

In-sample RMSEs per category using % MSE loss function

The table shows the average in-sample errors using the % RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	15.67 %	8.78 %	7.84 %	9.27 %	4.00 %	12.89 %
		SVJ	15.51 %	10.11 %	7.86 %	8.93 %	4.77 %	13.07 %
		MFSV	14.51 %	8.79 %	6.57 %	7.98 %	4.47 %	12.14 %
		MFSVJ	10.72 %	7.49 %	6.39 %	7.07 %	4.15 %	9.68 %
		PBS	8.32 %	10.84 %	12.79 %	8.63 %	11.64 %	11.25 %
	0.94-0.97	SV	12.28 %	6.31 %	6.66 %	6.80 %	3.36 %	10.63 %
		SVJ	11.47 %	5.84 %	6.44 %	6.75 %	4.14 %	10.18 %
		MFSV	13.02 %	5.89 %	6.33 %	6.13 %	4.26 %	11.05 %
		MFSVJ	9.23 %	4.92 %	6.19 %	5.45 %	3.83 %	8.40 %
		PBS	10.95 %	9.13 %	10.31 %	5.75 %	8.81 %	10.99 %
ATM	0.97-1.00	SV	6.78 %	6.06 %	6.33 %	5.37 %	2.97 %	7.18 %
		SVJ	6.05 %	5.18 %	6.13 %	5.40 %	3.77 %	6.72 %
		MFSV	6.82 %	4.95 %	6.32 %	5.01 %	3.98 %	7.12 %
		MFSVJ	4.90 %	4.37 %	6.20 %	4.48 %	3.63 %	5.66 %
		PBS	8.75 %	7.74 %	9.21 %	4.76 %	8.25 %	9.39 %
	1.00-1.03	SV	3.93 %	5.90 %	6.01 %	4.45 %	2.71 %	5.39 %
		SVJ	3.48 %	5.08 %	5.86 %	4.50 %	3.45 %	5.04 %
		MFSV	3.07 %	4.27 %	6.13 %	4.30 %	3.73 %	4.77 %
		MFSVJ	2.54 %	4.00 %	6.00 %	3.86 %	3.37 %	4.31 %
		PBS	4.81 %	6.63 %	8.44 %	4.23 %	6.79 %	6.72 %
	1.03-1.06	SV	3.12 %	5.89 %	5.51 %	3.86 %	2.32 %	4.91 %
		SVJ	2.84 %	5.20 %	5.38 %	3.80 %	2.85 %	4.58 %
		MFSV	1.83 %	4.19 %	5.70 %	3.71 %	3.21 %	4.19 %
		MFSVJ	1.95 %	4.03 %	5.57 %	3.37 %	2.89 %	4.06 %
		PBS	3.79 %	5.92 %	7.86 %	4.11 %	5.29 %	6.03 %
ITM	1.06-1.10	SV	1.95 %	5.52 %	4.79 %	3.63 %	2.08 %	4.53 %
		SVJ	1.87 %	4.96 %	4.66 %	3.43 %	2.27 %	4.20 %
		MFSV	1.12 %	4.09 %	5.00 %	3.37 %	2.41 %	3.95 %
		MFSVJ	1.14 %	3.91 %	4.85 %	3.20 %	2.19 %	3.77 %
		PBS	2.61 %	5.26 %	7.02 %	3.96 %	4.19 %	5.49 %
	All	SV	11.60 %	7.00 %	6.81 %	6.76 %	3.64 %	9.56 %
		SVJ	11.18 %	7.02 %	6.65 %	6.58 %	4.08 %	9.33 %
		MFSV	11.38 %	6.21 %	6.39 %	6.02 %	3.94 %	9.33 %
		MFSVJ	8.20 %	5.36 %	6.19 %	5.39 %	3.83 %	7.28 %
		PBS	9.60 %	8.52 %	10.04 %	6.29 %	9.36 %	9.60 %

Figure A3

1-day out-of-sample RMSEs for all models and loss functions

Each model is evaluated under each loss function for each day in the sample period using parameter estimates from the previous day. The first column represents IV RMSEs, the second \$ RMSEs and the third % RMSEs. Note the extreme variations in the RMSEs of the PBS model under all loss functions.




Figure A4**5-day out-of-sample RMSEs for all models and loss functions**

Each model is evaluated under each loss function for each day in the sample period using parameter estimates 5-days earlier. The first column represents IV RMSEs, the second \$ RMSEs and the third % RMSEs. Note the extreme variations in the RMSEs of the PBS model under all loss functions, and also the increased IV RMSEs and \$ RMSEs of the single-factor models during the high volatility period towards the end of 2008.




Table A4
t-statistics for 1- and 5-day out-of-sample RMSEs

t-statistics for each model pair for both 1- and 5-days out-of-sample. A positive t-statistic indicates that the model on the top row has a higher (inferior) sample mean. Values within brackets indicate significance at the 5 % level.

IV RMSE		SV	SVJ	MFSV	MFSVJ	PBS
SV	1-day		-1.5028	{-6.8023}	{-8.0739}	{8.3983}
	5-day		-0.7447	{-6.0670}	{-6.3880}	{10.6623}
SVJ		1.5028		{-5.1651}	{-6.4012}	{8.7290}
		0.7447		{-5.1412}	{-5.4789}	{10.8145}
MFSV		{6.8023}	{5.1651}		-1.2430	{9.8459}
		{6.0670}	{5.1412}		-0.5030	{11.8877}
MFSVJ		{8.0739}	{6.4012}	1.2430		{10.0944}
		{6.3880}	{5.4789}	0.5030		{11.9688}
PBS		{-8.3983}	{-8.7290}	{-9.8459}	{-10.0944}	
		{-10.6623}	{-10.8145}	{-11.8877}	{-11.9688}	
\$ RMSE		SV	SVJ	MFSV	MFSVJ	PBS
SV	1-day		-1.0169	{-3.8622}	{-4.2714}	{5.5354}
	5-day		-0.3679	{-3.7692}	{-2.9070}	{6.5542}
SVJ		1.0169		{-2.8800}	{-3.3897}	{5.6292}
		0.3679		{-3.3355}	{-2.5186}	{6.6113}
MFSV		{3.8622}	{2.8800}		-0.7851	{5.9016}
		{3.7692}	{3.3355}		0.5760	{7.0898}
MFSVJ		{4.2714}	{3.3897}	0.7851		{5.9803}
		{2.9070}	{2.5186}	-0.5760		{6.9993}
PBS		{-5.5354}	{-5.6292}	{-5.9016}	{-5.9803}	
		{-6.5542}	{-6.6113}	{-7.0898}	{-6.9993}	
% RMSE		SV	SVJ	MFSV	MFSVJ	PBS
SV	1-day		-0.2964	{-2.8044}	{-5.3277}	{7.7738}
	5-day		-0.4738	{-3.2615}	{-5.4035}	{11.8845}
SVJ		0.2964		{-2.5187}	{-5.0341}	{7.9359}
		0.4738		{-2.8285}	{-4.9742}	{12.0936}
MFSV		{2.8044}	{2.5187}		{-2.3331}	{9.2377}
		{3.2615}	{2.8285}		-1.9064	{13.1751}
MFSVJ		{5.3277}	{5.0341}	{2.3331}		{10.5618}
		{5.4035}	{4.9742}	1.9064		{14.0069}
PBS		{-7.7738}	{-7.9359}	{-9.2377}	{-10.5618}	
		{-11.8845}	{-12.0936}	{-13.1751}	{-14.0069}	

Table A5

1-day out-of-sample RMSEs per category using IV MSE loss function

The table shows the average 1-day out-of-sample errors using the IV RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	1.14 %	1.03 %	1.08 %	1.31 %	0.74 %	1.25 %
		SVJ	1.12 %	1.05 %	1.10 %	1.28 %	0.87 %	1.25 %
		MFSV	0.88 %	0.92 %	0.89 %	1.07 %	0.78 %	1.06 %
		MFSVJ	0.80 %	0.87 %	0.91 %	0.95 %	0.78 %	0.99 %
		PBS	1.96 %	1.69 %	1.91 %	1.36 %	4.34 %	2.67 %
	0.94-0.97	SV	1.08 %	1.09 %	1.19 %	1.21 %	0.74 %	1.21 %
		SVJ	0.99 %	1.01 %	1.18 %	1.21 %	0.89 %	1.16 %
		MFSV	0.89 %	0.89 %	1.08 %	1.04 %	0.88 %	1.06 %
		MFSVJ	0.80 %	0.82 %	1.12 %	0.94 %	0.89 %	0.99 %
		PBS	2.29 %	1.88 %	2.06 %	1.17 %	3.95 %	2.61 %
ATM	0.97-1.00	SV	0.97 %	1.36 %	1.39 %	1.14 %	0.75 %	1.27 %
		SVJ	0.84 %	1.19 %	1.37 %	1.15 %	0.92 %	1.18 %
		MFSV	0.78 %	1.02 %	1.35 %	1.02 %	0.90 %	1.09 %
		MFSVJ	0.72 %	0.94 %	1.37 %	0.94 %	0.89 %	1.03 %
		PBS	2.51 %	2.06 %	2.29 %	1.14 %	3.87 %	2.82 %
	1.00-1.03	SV	1.16 %	1.76 %	1.62 %	1.11 %	0.75 %	1.52 %
		SVJ	1.07 %	1.55 %	1.58 %	1.12 %	0.90 %	1.43 %
		MFSV	0.85 %	1.24 %	1.63 %	1.05 %	0.93 %	1.26 %
		MFSVJ	0.83 %	1.18 %	1.62 %	0.97 %	0.92 %	1.22 %
		PBS	2.78 %	2.32 %	2.56 %	1.20 %	3.57 %	3.06 %
	1.03-1.06	SV	1.82 %	2.34 %	1.79 %	1.12 %	0.69 %	1.99 %
		SVJ	1.79 %	2.10 %	1.75 %	1.10 %	0.81 %	1.89 %
		MFSV	1.30 %	1.63 %	1.87 %	1.07 %	0.87 %	1.62 %
		MFSVJ	1.33 %	1.58 %	1.83 %	1.00 %	0.85 %	1.61 %
		PBS	3.38 %	2.69 %	2.81 %	1.32 %	3.01 %	3.37 %
ITM	1.06-1.10	SV	2.23 %	2.94 %	1.92 %	1.18 %	0.73 %	2.52 %
		SVJ	2.36 %	2.66 %	1.87 %	1.13 %	0.77 %	2.45 %
		MFSV	1.69 %	2.12 %	2.03 %	1.15 %	0.79 %	2.11 %
		MFSVJ	1.71 %	2.06 %	1.97 %	1.10 %	0.78 %	2.09 %
		PBS	3.85 %	3.12 %	2.95 %	1.42 %	2.09 %	3.81 %
	All	SV	1.41 %	1.63 %	1.49 %	1.26 %	0.82 %	1.49 %
		SVJ	1.37 %	1.50 %	1.46 %	1.24 %	0.92 %	1.43 %
		MFSV	1.11 %	1.24 %	1.44 %	1.12 %	0.86 %	1.24 %
		MFSVJ	1.05 %	1.17 %	1.44 %	1.03 %	0.91 %	1.20 %
		PBS	2.78 %	2.20 %	2.41 %	1.37 %	4.09 %	2.98 %

Table A6

1-day out-of-sample RMSEs per category using \$ MSE loss function

The table shows the average 1-day out-of-sample errors using the \$ RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	0.0815	0.1786	0.3527	0.6083	0.5141	0.3670
		SVJ	0.0788	0.1800	0.3543	0.5908	0.6192	0.3657
		MFSV	0.0571	0.1651	0.2852	0.4916	0.5726	0.3216
		MFSVJ	0.0548	0.1582	0.2914	0.4379	0.5993	0.3219
		PBS	0.1576	0.2907	0.5993	0.6509	3.4055	1.1596
	0.94-0.97	SV	0.0868	0.2246	0.4039	0.5709	0.5201	0.3258
		SVJ	0.0779	0.2077	0.3977	0.5673	0.6353	0.3278
		MFSV	0.0692	0.1845	0.3619	0.4894	0.6411	0.2994
		MFSVJ	0.0642	0.1734	0.3726	0.4416	0.6633	0.2982
		PBS	0.2123	0.3640	0.6774	0.5632	3.0278	0.9306
ATM	0.97-1.00	SV	0.1009	0.2978	0.4700	0.5388	0.5277	0.3625
		SVJ	0.0861	0.2636	0.4610	0.5409	0.6468	0.3600
		MFSV	0.0849	0.2251	0.4538	0.4811	0.6516	0.3385
		MFSVJ	0.0769	0.2104	0.4581	0.4419	0.6489	0.3272
		PBS	0.2618	0.4277	0.7582	0.5443	2.9414	0.9974
	1.00-1.03	SV	0.1308	0.3839	0.5366	0.5212	0.5272	0.4176
		SVJ	0.1171	0.3386	0.5243	0.5220	0.6333	0.4102
		MFSV	0.0991	0.2731	0.5387	0.4861	0.6603	0.3913
		MFSVJ	0.0947	0.2605	0.5374	0.4502	0.6570	0.3800
		PBS	0.3045	0.4854	0.8379	0.5551	2.6525	1.0407
	1.03-1.06	SV	0.1788	0.4883	0.5725	0.5134	0.4614	0.4700
		SVJ	0.1700	0.4356	0.5607	0.5008	0.5380	0.4545
		MFSV	0.1242	0.3426	0.5985	0.4844	0.5809	0.4420
		MFSVJ	0.1262	0.3311	0.5886	0.4526	0.5747	0.4284
		PBS	0.3353	0.5414	0.8955	0.5942	2.1368	0.9979
ITM	1.06-1.10	SV	0.1710	0.5659	0.5801	0.5187	0.4670	0.5554
		SVJ	0.1724	0.5101	0.5663	0.4934	0.4891	0.5258
		MFSV	0.1170	0.4093	0.6206	0.5006	0.4993	0.5143
		MFSVJ	0.1202	0.3952	0.6032	0.4787	0.4916	0.4929
		PBS	0.3052	0.5803	0.9015	0.6130	1.3315	0.9254
	All	SV	0.1235	0.3309	0.4861	0.5891	0.5755	0.3933
		SVJ	0.1149	0.3003	0.4766	0.5751	0.6468	0.3820
		MFSV	0.0923	0.2517	0.4686	0.5178	0.6066	0.3492
		MFSVJ	0.0890	0.2390	0.4654	0.4738	0.6604	0.3390
		PBS	0.2616	0.4316	0.7717	0.6380	3.1260	1.0570

Table A7

1-day out-of-sample RMSEs per category using % MSE loss function

The table shows the average 1-day out-of-sample errors using the % RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	17.53 %	10.41 %	8.11 %	9.25 %	4.14 %	14.35 %
		SVJ	17.54 %	11.52 %	8.12 %	8.95 %	4.84 %	14.72 %
		MFSV	15.17 %	9.08 %	6.69 %	8.01 %	4.49 %	12.65 %
		MFSVJ	13.21 %	8.63 %	6.79 %	7.20 %	4.52 %	11.32 %
		PBS	20.61 %	16.11 %	13.97 %	9.28 %	24.43 %	21.80 %
	0.94-0.97	SV	13.65 %	7.10 %	6.85 %	6.78 %	3.55 %	11.61 %
		SVJ	12.86 %	6.65 %	6.63 %	6.72 %	4.29 %	11.18 %
		MFSV	12.37 %	5.56 %	6.09 %	6.12 %	4.28 %	10.58 %
		MFSVJ	10.89 %	5.54 %	6.35 %	5.57 %	4.28 %	9.57 %
		PBS	19.54 %	13.04 %	11.38 %	6.32 %	18.36 %	18.56 %
ATM	0.97-1.00	SV	7.38 %	6.44 %	6.51 %	5.39 %	3.22 %	7.66 %
		SVJ	6.49 %	5.55 %	6.27 %	5.40 %	3.96 %	7.10 %
		MFSV	6.03 %	4.78 %	6.13 %	4.98 %	3.98 %	6.52 %
		MFSVJ	5.45 %	4.65 %	6.26 %	4.60 %	3.91 %	6.10 %
		PBS	13.99 %	10.44 %	10.06 %	5.24 %	16.31 %	14.65 %
	1.00-1.03	SV	4.12 %	6.18 %	6.19 %	4.50 %	2.90 %	5.59 %
		SVJ	3.78 %	5.36 %	5.97 %	4.50 %	3.55 %	5.26 %
		MFSV	3.14 %	4.35 %	6.01 %	4.27 %	3.70 %	4.73 %
		MFSVJ	2.97 %	4.26 %	6.05 %	3.98 %	3.64 %	4.59 %
		PBS	8.30 %	8.55 %	9.17 %	4.70 %	14.18 %	10.67 %
	1.03-1.06	SV	3.35 %	6.15 %	5.62 %	3.94 %	2.45 %	5.11 %
		SVJ	3.19 %	5.43 %	5.44 %	3.83 %	2.91 %	4.79 %
		MFSV	2.39 %	4.37 %	5.64 %	3.73 %	3.14 %	4.38 %
		MFSVJ	2.46 %	4.27 %	5.61 %	3.51 %	3.09 %	4.32 %
		PBS	5.83 %	7.31 %	8.36 %	4.50 %	11.11 %	8.70 %
ITM	1.06-1.10	SV	2.13 %	5.69 %	4.89 %	3.61 %	2.16 %	4.73 %
		SVJ	2.15 %	5.09 %	4.74 %	3.41 %	2.33 %	4.41 %
		MFSV	1.48 %	4.24 %	5.01 %	3.44 %	2.43 %	4.15 %
		MFSVJ	1.54 %	4.09 %	4.91 %	3.31 %	2.37 %	4.04 %
		PBS	3.65 %	6.14 %	7.24 %	4.22 %	6.34 %	6.98 %
	All	SV	12.91 %	7.89 %	7.02 %	6.76 %	3.78 %	10.55 %
		SVJ	12.61 %	7.90 %	6.85 %	6.57 %	4.21 %	10.42 %
		MFSV	11.44 %	6.36 %	6.34 %	6.04 %	3.89 %	9.37 %
		MFSVJ	9.96 %	6.06 %	6.38 %	5.51 %	4.11 %	8.45 %
		PBS	17.44 %	12.04 %	11.06 %	6.84 %	19.14 %	16.77 %

Table A8

5-day out-of-sample RMSEs per category using IV MSE loss function

The table shows the average 5-day out-of-sample errors using the IV RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	1.37 %	1.18 %	1.15 %	1.22 %	0.85 %	1.37 %
		SVJ	1.33 %	1.16 %	1.13 %	1.19 %	0.95 %	1.35 %
		MFSV	1.04 %	0.96 %	0.96 %	1.05 %	0.82 %	1.14 %
		MFSVJ	0.98 %	0.96 %	0.98 %	0.95 %	0.89 %	1.10 %
		PBS	3.66 %	2.76 %	2.30 %	1.71 %	5.63 %	4.04 %
	0.94-0.97	SV	1.27 %	1.25 %	1.26 %	1.16 %	0.89 %	1.34 %
		SVJ	1.20 %	1.14 %	1.24 %	1.16 %	1.01 %	1.29 %
		MFSV	1.02 %	0.93 %	1.13 %	1.01 %	0.91 %	1.13 %
		MFSVJ	0.96 %	0.90 %	1.15 %	0.97 %	1.02 %	1.08 %
		PBS	4.01 %	2.91 %	2.42 %	1.55 %	5.29 %	3.97 %
ATM	0.97-1.00	SV	1.08 %	1.52 %	1.47 %	1.13 %	0.91 %	1.38 %
		SVJ	0.98 %	1.34 %	1.44 %	1.13 %	1.04 %	1.29 %
		MFSV	0.86 %	1.11 %	1.39 %	1.00 %	0.94 %	1.17 %
		MFSVJ	0.83 %	1.04 %	1.39 %	0.99 %	1.08 %	1.13 %
		PBS	4.22 %	3.09 %	2.59 %	1.47 %	5.16 %	4.07 %
	1.00-1.03	SV	1.23 %	1.93 %	1.72 %	1.15 %	0.93 %	1.63 %
		SVJ	1.17 %	1.72 %	1.67 %	1.15 %	1.04 %	1.54 %
		MFSV	0.94 %	1.38 %	1.66 %	1.04 %	0.96 %	1.36 %
		MFSVJ	0.89 %	1.30 %	1.63 %	1.04 %	1.11 %	1.31 %
		PBS	4.46 %	3.32 %	2.81 %	1.47 %	4.85 %	4.33 %
	1.03-1.06	SV	1.87 %	2.54 %	1.89 %	1.19 %	0.88 %	2.12 %
		SVJ	1.91 %	2.28 %	1.85 %	1.17 %	0.99 %	2.05 %
		MFSV	1.44 %	1.84 %	1.91 %	1.07 %	0.93 %	1.75 %
		MFSVJ	1.40 %	1.75 %	1.84 %	1.09 %	1.04 %	1.71 %
		PBS	4.98 %	3.55 %	3.01 %	1.52 %	4.24 %	4.48 %
ITM	1.06-1.10	SV	2.25 %	3.14 %	2.04 %	1.25 %	0.89 %	2.64 %
		SVJ	2.54 %	2.88 %	2.00 %	1.19 %	0.92 %	2.69 %
		MFSV	1.80 %	2.39 %	2.09 %	1.16 %	0.80 %	2.24 %
		MFSVJ	1.81 %	2.27 %	1.97 %	1.19 %	0.97 %	2.19 %
		PBS	5.15 %	3.88 %	3.19 %	1.60 %	3.15 %	4.70 %
	All	SV	1.54 %	1.79 %	1.57 %	1.25 %	1.00 %	1.65 %
		SVJ	1.54 %	1.64 %	1.53 %	1.22 %	1.05 %	1.61 %
		MFSV	1.24 %	1.36 %	1.50 %	1.11 %	0.89 %	1.36 %
		MFSVJ	1.19 %	1.30 %	1.46 %	1.07 %	1.06 %	1.34 %
		PBS	4.45 %	3.20 %	2.70 %	1.66 %	5.45 %	4.26 %

Table A9

5-day out-of-sample RMSEs per category using \$ MSE loss function

The table shows the average 5-day out-of-sample errors using the \$ RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	0.0990	0.2102	0.3710	0.5651	0.6000	0.3870
		SVJ	0.0968	0.2036	0.3632	0.5503	0.6811	0.3801
		MFSV	0.0717	0.1709	0.3115	0.4807	0.5995	0.3343
		MFSVJ	0.0716	0.1741	0.3138	0.4305	0.6700	0.3383
		PBS	0.2873	0.4667	0.7196	0.8028	4.4364	1.4920
	0.94-0.97	SV	0.1048	0.2590	0.4238	0.5500	0.6290	0.3547
		SVJ	0.0977	0.2363	0.4186	0.5461	0.7185	0.3532
		MFSV	0.0791	0.1939	0.3784	0.4761	0.6661	0.3099
		MFSVJ	0.0777	0.1909	0.3846	0.4491	0.7458	0.3197
		PBS	0.3719	0.5562	0.7891	0.7349	4.0669	1.2555
ATM	0.97-1.00	SV	0.1120	0.3324	0.4966	0.5362	0.6441	0.3978
		SVJ	0.1011	0.2957	0.4876	0.5325	0.7391	0.3914
		MFSV	0.0913	0.2419	0.4680	0.4722	0.6789	0.3499
		MFSVJ	0.0871	0.2331	0.4675	0.4589	0.7761	0.3611
		PBS	0.4469	0.6307	0.8541	0.6915	3.9294	1.3161
	1.00-1.03	SV	0.1368	0.4210	0.5680	0.5364	0.6426	0.4564
		SVJ	0.1297	0.3749	0.5540	0.5312	0.7233	0.4449
		MFSV	0.1101	0.3007	0.5499	0.4818	0.6804	0.4087
		MFSVJ	0.1015	0.2889	0.5401	0.4782	0.7812	0.4159
		PBS	0.4939	0.6823	0.9144	0.6814	3.6016	1.3804
	1.03-1.06	SV	0.1819	0.5291	0.6043	0.5409	0.5917	0.5246
		SVJ	0.1824	0.4744	0.5953	0.5275	0.6643	0.5120
		MFSV	0.1405	0.3825	0.6121	0.4842	0.6281	0.4704
		MFSVJ	0.1335	0.3674	0.5892	0.4877	0.7029	0.4716
		PBS	0.5069	0.7054	0.9546	0.6839	3.0142	1.3029
ITM	1.06-1.10	SV	0.1674	0.6055	0.6156	0.5427	0.5720	0.6021
		SVJ	0.1773	0.5536	0.6092	0.5158	0.5890	0.5781
		MFSV	0.1290	0.4574	0.6381	0.5043	0.5074	0.5290
		MFSVJ	0.1248	0.4381	0.6040	0.5145	0.6164	0.5364
		PBS	0.4205	0.7213	0.9745	0.6901	2.0511	1.1636
	All	SV	0.1339	0.3657	0.5083	0.5804	0.6989	0.4569
		SVJ	0.1298	0.3321	0.4986	0.5656	0.7384	0.4488
		MFSV	0.1054	0.2741	0.4873	0.5115	0.6287	0.3836
		MFSVJ	0.1013	0.2651	0.4734	0.4861	0.7611	0.3944
		PBS	0.4227	0.6159	0.8633	0.7681	4.1695	1.3889

Table A10

5-day out-of-sample RMSEs per category using % MSE loss function

The table shows the average 5-day out-of-sample errors using the % RMSE loss function for each maturity and moneyness category, for all models. $F_{t,T}$ denotes the forward price and K the strike price. The moneyness categories are sorted into three subgroups: out-of-the money (OTM), at-the-money (ATM) and in-the-money (ITM) options.

Moneyness ($F_{t,T}/K$)			Days to maturity					All
			< 60	60-179	180-359	360-719	>720	
OTM	0.90-0.94	SV	20.42 %	11.25 %	9.03 %	8.62 %	4.62 %	15.73 %
		SVJ	19.96 %	11.78 %	8.71 %	8.38 %	5.15 %	15.64 %
		MFSV	16.39 %	9.44 %	7.32 %	7.84 %	4.61 %	13.33 %
		MFSVJ	14.50 %	9.29 %	7.38 %	7.16 %	4.99 %	12.14 %
		PBS	41.00 %	26.88 %	16.06 %	10.92 %	31.63 %	35.95 %
	0.94-0.97	SV	15.26 %	7.95 %	7.48 %	6.50 %	4.10 %	12.77 %
		SVJ	14.57 %	7.21 %	7.29 %	6.46 %	4.68 %	12.30 %
		MFSV	13.34 %	5.73 %	6.42 %	5.95 %	4.36 %	11.20 %
		MFSVJ	11.80 %	5.78 %	6.56 %	5.65 %	4.74 %	10.18 %
		PBS	33.60 %	20.37 %	12.93 %	7.93 %	24.30 %	29.03 %
ATM	0.97-1.00	SV	7.65 %	7.08 %	7.00 %	5.33 %	3.78 %	8.09 %
		SVJ	6.88 %	6.15 %	6.79 %	5.28 %	4.38 %	7.52 %
		MFSV	6.48 %	5.07 %	6.35 %	4.88 %	4.07 %	6.92 %
		MFSVJ	6.04 %	5.02 %	6.37 %	4.76 %	4.52 %	6.68 %
		PBS	21.53 %	15.36 %	11.15 %	6.43 %	21.71 %	20.26 %
	1.00-1.03	SV	4.20 %	6.72 %	6.61 %	4.60 %	3.45 %	5.96 %
		SVJ	3.96 %	5.89 %	6.40 %	4.54 %	3.97 %	5.59 %
		MFSV	3.39 %	4.74 %	6.18 %	4.24 %	3.78 %	5.01 %
		MFSVJ	3.14 %	4.66 %	6.10 %	4.21 %	4.23 %	4.90 %
		PBS	12.94 %	11.88 %	9.88 %	5.56 %	18.96 %	14.49 %
1.03-1.06	SV	3.45 %	6.59 %	5.94 %	4.11 %	3.03 %	5.48 %	
	SVJ	3.45 %	5.85 %	5.82 %	3.99 %	3.46 %	5.21 %	
	MFSV	2.67 %	4.81 %	5.81 %	3.73 %	3.33 %	4.66 %	
	MFSVJ	2.60 %	4.68 %	5.63 %	3.76 %	3.66 %	4.60 %	
	PBS	8.64 %	9.36 %	8.79 %	5.02 %	15.33 %	11.10 %	
ITM	1.06-1.10	SV	2.12 %	6.04 %	5.17 %	3.74 %	2.59 %	4.98 %
		SVJ	2.26 %	5.49 %	5.09 %	3.53 %	2.74 %	4.76 %
		MFSV	1.62 %	4.66 %	5.16 %	3.47 %	2.45 %	4.33 %
		MFSVJ	1.60 %	4.48 %	4.93 %	3.57 %	2.89 %	4.26 %
		PBS	5.00 %	7.45 %	7.69 %	4.61 %	9.69 %	8.40 %
	All	SV	14.29 %	8.64 %	7.59 %	6.51 %	4.37 %	11.58 %
		SVJ	13.96 %	8.36 %	7.37 %	6.34 %	4.61 %	11.36 %
		MFSV	12.38 %	6.75 %	6.73 %	5.93 %	3.97 %	10.04 %
		MFSVJ	10.95 %	6.59 %	6.64 %	5.60 %	4.61 %	9.22 %
		PBS	30.48 %	19.00 %	12.40 %	8.03 %	25.25 %	25.99 %

Table A11
Average parameter estimates, IV RMSE and average computation time for 10 estimations
on July 17th 2008

The average parameter estimates calculated from the sample of 10 estimations, where the starting values of each parameter was randomly chosen on a uniformly distributed interval of $\mu_i \pm \sigma_i$, where μ_i and σ_i denote the mean and standard deviation of parameter i in the entire sample estimation. The corresponding standard deviations are shown in brackets. For comparative purposes, the parameters of the PBS model have been obtained using the strike price in fractions of the spot price, making the estimates of α_1 and α_5 100 times larger and the estimate α_2 10 000 larger than the corresponding estimates if actual strike prices are used.

	κ	θ	σ	ρ	λ	μ_J	σ_J	IV RMSE	Time (s)
SV	7.9153 (6.4225)	0.0614 (0.0048)	0.7432 (0.2264)	-0.9996 (0.0006)				1.13 % (0.01%)	53 (37)
SVJ	9.2317 (20.0212)	0.0177 (0.0063)	0.4030 (0.3002)	-0.9939 (0.0194)	0.9398 (0.2778)	-0.1584 (0.0279)	0.0754 (0.0398)	0.78 % (0.12%)	199 (115)
MFSV	0.6804 (0.4739) 14.2110 (5.1696)	0.0255 (0.0310) 0.0401 (0.0084)	1.1533 (1.3662) 0.9241 (0.6295)	-0.9552 (0.0685) -0.9238 (0.1481)				0.75 % (0.10%)	271 (209)
MFSVJ	3.8906 (7.6951) 9.4296 (8.7548)	0.0106 (0.0223) 0.0043 (0.0057)	1.5608 (1.7434) 0.0926 (0.1456)	-0.5864 (0.4418) -0.5261 (0.2554)	7.8705 (0.9875)	-0.0467 (0.0050)	0.0610 (0.0081)	0.74 % (0.10%)	638 (500)
	α_0	α_1	α_2	α_3	α_4	α_5		IV RMSE	Time (s)
PBS	-0.0587 (0.6921)	1.1306 (1.3689)	-0.8111 (0.6753)	-0.1552 (0.0493)	0.0019 (0.0002)	0.1384 (0.0473)		0.95 % (0.09%)	0.49 (0.16)

Appendix B: Volatility surface parameterization

To calibrate the volatility surface, we use the stochastic volatility inspired (SVI) method of Gatheral (2004). This means that for each maturity, we use a least-squares method to fit a function of the form

$$var(k) = a + b \left(\rho(k - m) + \sqrt{(k - m)^2 + \sigma^2} \right) \quad (B.1)$$

to each observed level of k , where $k = \log(K/F_{t,T})$ in which $F_{t,T}$ is the forward price of S , i.e. $F_{t,T} = S_t e^{(r-d)(T-t)}$. a, b, ρ, m and σ are parameters of the function. See Gatheral (2004) for more details on the SVI method, including a mathematical background and a discussion of the parameters.







The volatility surface is obtained by interpolation in the term-structure dimension using a third degree polynomial. More sophisticated interpolation techniques are often used, but for illustrative purposes, we consider the third degree polynomial to be sufficient.







Figure B5 below shows the SVI functions, with implied volatility on the y-axis and strike price on the x-axis. As can be seen from the plots, the volatility smirk is especially apparent for short maturities. The stars represent the mid prices, calculated as the average of the bid and ask price.

The resulting volatility surface is Figure 2 in Section 3 above.

Figure B5**Skew plots (SVI functions) on July 17th 2008**

The figures show the SVI function fitted to the implied volatilities (marked as stars) for different maturities. The implied volatility is plotted on the y-axis and the strike price is shown on the x-axis.





Appendix C: Derivation of the call price formula using characteristic functions and the FFT.

We show the derivation of the pricing formula for European call options using characteristic functions and the FFT in order to provide the reader with some intuition, as the FFT method is central in the estimation and evaluation of the SV, SVJ, MFSV and MFSVJ models. We follow closely the method of Carr & Madan (1999), but extend the method to allow for a continuous dividend yield. The latter is used as an approximation of calculating ex-dividend prices of the underlying indices when estimating the models.

Denote by s_T and k the natural logarithm of the terminal stock price and the strike price K , respectively. Further, let $C_T(k)$ denote the value of a European call option with pay-off function $f(S_T) = (S_T - K)^+ = (e^{s_T} - e^k)^+$ and maturity at time T . The discounted expected pay-off under $\mathbb{Q}$ is then:

$$C_T(k) = \mathbb{E}_t^{\mathbb{Q}}[e^{-rT} (S_T - K)^+] = \int_k^{\infty} e^{-rT} (e^{s_T} - e^k) q_T(s_T) ds_T \quad (\text{C.8})$$

where $q_T(s)$ is the risk-neutral density of s_T . As k tends to $-\infty$, (C.8) translates to:

$$\lim_{k \rightarrow -\infty} C_T(k) = \int_{-\infty}^{\infty} e^{-rT} e^{s_T} q_T(s_T) ds_T = e^{-rT} \mathbb{E}_t^{\mathbb{Q}}[S_T] = S_0 \quad (\text{C.9})$$

This is on the one hand reassuring, as the price of a call with zero strike should equal S_0 . On the other hand, in order to apply the Fourier transform to $C_T(k)$ it is required that the function is square integrable for all k , i.e. that $\int_{-\infty}^{\infty} |C_T(k)|^2 ds_T < \infty \forall k \in \mathbb{R}$. However, by (C.9), as k tends to $-\infty$:

$$\lim_{k \rightarrow -\infty} \int_{-\infty}^{\infty} |C_T(k)|^2 ds_T = \int_{-\infty}^{\infty} |S_0|^2 ds_T \rightarrow \infty \quad (\text{C.10})$$

showing that $C_T(k)$ is not square integrable. This problem is solved by introducing the modified call price function:

$$c_T(k) = e^{\alpha k} C_T(k) \quad (\text{C.11})$$

for some $\alpha > 0$. The modified call price function, $c_T(k)$, is then expected to be square integrable for all $k \in \mathbb{R}$, provided that α is chosen correctly. The Fourier transform of $c_T(k)$ takes the following form:

$$\mathfrak{F}\{c_T(k)\} = \int_{-\infty}^{\infty} c_T(k) e^{i\xi k} dk = \psi_T(\xi) \quad (\text{C.12})$$

Combining (C.8), (C.11) and (C.12), we obtain:

$$\begin{aligned} \psi_T(\xi) &= \int_{-\infty}^{\infty} e^{i\xi k} \int_k^{\infty} e^{\alpha k} e^{-(r-q)T} (e^{s_T} - e^k) q_T(s_T) ds_T dk \\ &= \int_{-\infty}^{\infty} e^{-rT} q_T(s_T) \int_{-\infty}^{s_T} (e^{s_T+\alpha k} - e^{(1+\alpha)k}) e^{i\xi k} dk ds_T \\ &= \int_{-\infty}^{\infty} e^{-rT} q_T(s_T) \left[\frac{e^{(\alpha+1+i\xi)s_T}}{\alpha+i\xi} - \frac{e^{(\alpha+1+i\xi)s_T}}{\alpha+1+i\xi} \right] ds_T \\ &= \frac{e^{-rT}}{\alpha^2 + \alpha - \xi^2 + i(2\alpha+1)\xi} \int_{-\infty}^{\infty} e^{(-\alpha i - i + \xi)is_T} q_T(s_T) ds_T \\ &= \frac{e^{-rT} \varphi_T(\xi - (\alpha+1)i)}{\alpha^2 + \alpha - \xi^2 + i(2\alpha+1)\xi} \end{aligned} \quad (\text{C.13})$$

where $\varphi_T(\cdot)$ denotes the characteristic function of s_T . To obtain the second equality, we use the equivalence between integrating over all $s_T > k$ with respect to s_T , i.e. $\int_k^{\infty} (\cdot) ds_T$, and integrating over all $k < s_T$ with respect to k , i.e. $\int_{-\infty}^{s_T} (\cdot) dk$. The call price can then be obtained by Fourier inversion of $\psi_T(\xi)$ and multiplication by $e^{-\alpha k}$:

$$\begin{aligned} C_T(k) &= e^{-\alpha k} \cdot \mathfrak{F}^{-1}\{\psi_T(\xi)\} = \frac{e^{-\alpha k}}{2\pi} \int_{-\infty}^{\infty} e^{-i\xi k} \psi_T(\xi) d\xi = \frac{e^{-\alpha k}}{\pi} \int_0^{\infty} e^{-i\xi k} \psi_T(\xi) d\xi \\ &\approx \frac{e^{-\alpha k}}{\pi} \sum_{j=1}^N e^{-i\xi_j k} \psi_T(\xi_j) \eta_j, \quad j = 1, \dots, N. \end{aligned} \quad (\text{C.14})$$

where $\xi_j = \eta(j - 1)$ and η is the step size in the integration grid. (C.14) can be re-written as:

$$C_T(k_u) = \frac{e^{-\alpha k_u}}{\pi} \sum_{j=1}^N e^{-i\frac{2\pi}{N}(j-1)(u-1)} e^{ib\xi_j} \psi(\xi_j) \frac{\eta}{3} (3 + (-1)^j - \mathbb{I}(j-1)_{\{0\}}) \quad (\text{C.15})$$

where $b = \pi/\eta$; $k_u = -b + (2b/N)(u - 1)$, $u = 1, \dots, N + 1$; and $\mathbb{I}(x)_{\mathcal{M}}$ is the indicator function equal to 1 if $x \in \mathcal{M}$ and 0 otherwise. The term $1/3 \cdot (3 + (-1)^j - \mathbb{I}(j-1)_{\{0\}})$ is obtained using the Simpson rule for numerical integration.

Now, the idea of writing the call price on the form (C.15) is that it enables the use of the Fast Fourier Transform (FFT). The FFT is an algorithm to efficiently evaluate summations on the form:

$$X(k) = \sum_{j=1}^N e^{-i\frac{2\pi}{N}(j-1)(k-1)} x(j), \quad k = 1, \dots, N. \quad (\text{C.16})$$

With $x_j = e^{ib\xi_j} \psi(\xi_j) \frac{\eta}{3} (3 + (-1)^j - \mathbb{I}(j-1)_{\{0\}})$, (C.15) is a special case of (C.16) and can thus be evaluated using the FFT³¹.

Since the computed call option value will be dependent on parameter choices, namely the choice of η , N and α , it is important that these are chosen carefully. For the purpose of this thesis, we chose the integration parameters N and η to be 4096 and 0.15, respectively, in order to obtain a reasonable trade-off between accuracy and speed. The optimal choice of α depends on the characteristic function of s_T for the model at hand. For all the models treated in this thesis, a value of $\alpha = 0.75$ is considered a suitable choice (see e.g. Borak, Detlefsen & Härdle, 2005 and Schoutens, Simons & Tistaert, 2003). Please refer to Carr & Madan (1999) for a more thorough discussion regarding the choice of parameters.

After deciding on appropriate values for η , N , and α , call prices are calculated by evaluating the sum in (C.15) using the FFT. The price of a put option can then be determined through put-call-parity.

³¹ The FFT is a built in function in many mathematical packages, such as e.g. MATLAB.

Appendix D: Data cleaning

In order to exclude erroneous observations that might distort the analysis, we apply several filters to the raw data before conducting our analysis. In order to save computation time, the procedure is performed step-wise and once an option has been caught in a filter, it is not examined in remaining filters. As a consequence, the number of options removed in each filter (displayed in Table D1 below) are those that breached the conditions of that particular filter, but none of the previous filters. This of course has no effect on the final data set obtained, but merely affects the interpretation of the number of options removed in each step.

Firstly, we remove all options with no traded volume or open interest, since we cannot be sure that these are valid market prices. We also remove options with shorter than six days to maturity, since options close to expiry may suffer from substantial liquidity biases due to prices being affected by traders that have to buy or sell large amounts of options before expiry (Bakshi, Cao & Chen, 1997).

In steps 3 to 6, options that violate obvious no-arbitrage conditions are removed. This includes removing options with negative prices, negative spreads and options with negative time value. Further, again following Bakshi, Cao & Chen (1997), we continue by removing all options with prices less than 10 cents in order to mitigate the effect of discrete prices in option valuation³².

Step 8 is in accordance with e.g. Dumas, Fleming & Waley (1998), and is applied as far OTM and ITM options typically have little time premium and thus contain little valuable information about the implied volatility function (which is essentially what drives option prices). Options trading close to their intrinsic value (deep ITM) or close to zero (deep OTM) contain little information about the volatility function, which is essentially what drives option prices. Hence, little information is lost by this exclusion

Additional filters are applied in order to remove illiquid observations and observations with exceptionally high implied volatility, carried out in step 9 and 10. Options that heavily violate the requirement that call prices are monotonically decreasing in strike price, are removed in step 11.

³² Bakshi, Cao & Chen remove options with price less than \$3/8. As the index level of their sample was approximately 300 as compared to 100 in our normalized sample, 10 cents is a reasonable threshold.

This step is enforced by fitting a quadratic function of strike price to the option prices for every maturity and removing options with prices further than two standard deviations from the fitted curve.

Finally, days with less than 15 options are removed to ensure that we leave at least a few degrees of freedom when fitting the models (recall that the MFSVJ model has 11 structural parameters).

Table D1
Summary of data cleaning procedures

The table summarizes all the cleaning steps used on the initial dataset. Note that the number of removed options in each step is the ones that have not breached any earlier filter. This affects the interpretation one can draw from the eliminations in each step.

Cleaning steps		Step-wise removals
1	Remove options with no traded volume or open interest	40 472
2	Remove options with shorter than 6 days to maturity	2 359
3	Remove options with negative bid or ask price	0
4	Remove options with the bid price greater than the ask price	0
5	Remove options where bid or ask price greater than the index level (S_t)	0
6	Remove options where bid or ask price less than $(F_{t,T} - K)^+$	8 033
7	Remove options with bid or ask price less than 10 cents	31 741
8	Remove options with moneyness ($F_{t,T}/K$) lower than 90 % or higher than 110 %	36 119
9	Remove options where the ask price is more than 50 % higher than bid price	32
10	Remove options with higher implied vol. than 100 %	0
11	Remove options that heavily violate the requirement that call prices are monotonically decreasing in strike	1 504
12	Remove days with less than 15 different options	0
<i>Total number of removed options</i>		120 260
Remaining options		30 686

Appendix E: Estimation

Many different methods have been proposed for the problem of estimating the parameters of stochastic volatility models, especially with regards to the filtering of spot volatilities. We adopt the method of Christoffersen, Heston & Jacobs (2009) and estimate the models using an iterative two-step procedure.

Denote by Θ the parameter vector of the given model and let Λ_t denote the spot variances at time t . For the SV and SVJ models, Λ_t will be a scalar, whereas in the MFSV and MFSVJ models $\Lambda_t = \{V_t^{(1)}, V_t^{(2)}\}$. Each model is then estimated using the following two-step procedure:

1. For a given parameter vector Θ , solve the optimization problem:

$$\hat{\Lambda}_t = \arg \min_{\Lambda_t} \frac{1}{N_t} \sum_{i=1}^{N_t} w_{it} \left(\frac{C_{it} - \hat{C}_{it}(\Theta_t, \Lambda_t)}{\mathcal{V}_{it}^{BS}} \right)^2 \quad (\text{E.1})$$

where C_{jt} is the market price of option j on day t , $\hat{C}_{jt}(\Theta_t, \Lambda_t)$ is the corresponding model price and $\mathcal{V}_{jt}$ is the Black-Scholes Vega of option C_{jt} and $w_{it} = \frac{1}{ask_{it} - bid_{it}} / \sum_i \frac{1}{ask_{it} - bid_{it}}$.

2. Using the estimated spot variances from step 1, solve the optimization problem:

$$\hat{\Theta}_t = \arg \min_{\Theta_t} \frac{1}{N_t} \sum_{i=1}^{N_t} w_{it} \left(\frac{C_{it} - \hat{C}_{it}(\Theta_t, \Lambda_t)}{\mathcal{V}_{it}^{BS}} \right)^2 \quad (\text{E.2})$$

The process is repeated until no significant improvement in the loss function in step 2 is obtained. All optimization problems are solved in MATLAB using the `lsqnonlin` function. As `lsqnonlin` is a local optimizer and the goal function is non-convex and possesses several local minima (Cont & Hamida, 2005), some care has to be taken in order to not get stuck in local solutions. One way of mitigating this problem would be to run the optimization with a large range of starting values and choose the solution that gives the smallest value of the loss function. However, such a solution would be extremely time-consuming, as one calibration in step 2 can take up to 1 minute and the total calibration scheme includes approximately 3000 iterations (i.e. a total calibration time of roughly 50 hours). In order to minimize computation time while

maintaining reasonable accuracy, the following procedure is followed for each optimization in step 2³³:

1. On the first day in the calibration, we perform 10 optimizations using different starting values randomly chosen on specified uniform intervals (centered around expected parameter values) and choose the solution associated with the smallest value of the loss function.
2. On every subsequent day, we use the previous day's optimal parameter values as starting values in the optimization.
3. If the parameter values change very little or very much, we re-run the optimization with pre-specified starting values and choose the best solution of the two.

The first step is performed in order to ensure that the optimization in the first step is not caught in a local minimum. This is particularly important on the first day, as the parameter values obtained is used as starting values in subsequent optimizations. The use of the previous day's parameter values as starting values for the optimization function significantly decreases convergence time, as it is likely that parameter values on subsequent days are of similar magnitude. However, it was noted that this method occasionally resulted in fixed parameter values over several days as a result of the optimizer getting caught in a local minimum. To avoid this trap, we calculate the squared distance between the vector of starting values and the solution, i.e. $d = ||\Theta_0 - \hat{\Theta}||^2$, and re-run the optimization with pre-specified starting values if $d > 1$ or $d < 10^{-6}$ for the single factor models and if $d > 10$ or $d < 10^{-6}$ for the multi-factor models.

For the SV and SVJ models, we also implement the so called Feller (1951) condition, namely that $2\kappa\theta - \sigma^2 > 0$. This ensures that the variance process V_t cannot reach zero. Thus, when estimating the SV and SVJ models we introduce $\Psi = 2\kappa\theta - \sigma^2$ and estimate the models using Ψ instead of κ , with the simple restriction that $\Psi > 0$. Once the estimation is finished, κ is obtained as $\kappa = (\Psi + \sigma^2)/2\theta$. In accordance with previous studies, we do not require the Feller condition to be fulfilled for multifactor models (see e.g. Christoffersen, Heston & Jacobs, 2009 and Bates, 2000) but instead assume a reflecting barrier at the origin.

³³ As the loss function in step 1 is only a function of one or two variables (Λ), the optimization is relatively well behaved and does not require as much attention.

Appendix F: The approximate IV loss function

Recall from equation (6.5) the approximate implied volatility loss function used for estimation:

$$IV\ MSE(\Theta, \Lambda) = \frac{1}{n} \sum_{i=1}^n w_i (\sigma_i - \hat{\sigma}_i(\Theta, \Lambda))^2 \approx \frac{1}{n} \sum_{i=1}^n w_i \left(\frac{C_i - \hat{C}_i(\Theta, \Lambda)}{\mathcal{V}_i^{BS}} \right)^2 \quad (F.1)$$

where $\mathcal{V}_i^{BS}$ denotes the Black-Scholes Vega of option i and $w_i = \frac{1}{ask_i - bid_i} / \sum_i \frac{1}{ask_i - bid_i}$.

The approximation is obtained by considering the first order approximation:

$$\hat{C}_i(\Theta, \Lambda) \approx C_i + \mathcal{V}_i^{BS} \cdot (\hat{\sigma}_i(\Theta, \Lambda) - \sigma_i) \quad (F.2)$$

In order to assess the accuracy of the approximation, we formulate the equation:

$$\underbrace{\frac{\hat{C}_i(\Theta, \Lambda) - C_i}{\mathcal{V}_i^{BS}}}_{\Delta \tilde{\sigma}_i} = \underbrace{\hat{\sigma}_i(\Theta, \Lambda) - \sigma_i}_{\Delta \sigma_i} + \varepsilon_i \quad (F.3)$$

where we denote by $\Delta \tilde{\sigma}$ the approximated difference between the model implied volatility and the observed implied volatility and let $\Delta \sigma$ denote the actual difference.

Rearranging the terms, we obtain the approximation error in terms of the difference between the Vega approximated implied volatility difference and the true difference between the model implied volatility and the true implied volatility:

$$\varepsilon_i = \Delta \tilde{\sigma} - \Delta \sigma \quad (F.4)$$

Examining the three components in equation (F.4), especially ε , provides information about the validity of the linear approximation using the Black-Scholes Vega. The results are summarized in Table F1 below, showing the average absolute values of ε , $\Delta \tilde{\sigma}$ and $\Delta \sigma$ for the four structural models over the 30 686 options in the sample:

Table F1**Residual components of the linear approximation of implied volatility difference**

The table shows the average absolute values of the residual components of the linear approximation of difference in model implied volatility and observed implied volatility and the pair-wise correlations between the approximate implied volatility difference and the true difference.

	$ \Delta\tilde{\sigma} $	$ \Delta\sigma $	$ \varepsilon_i = \Delta\tilde{\sigma} - \Delta\sigma $	$Corr(\Delta\tilde{\sigma}, \Delta\sigma)$
SV	1.11 %	1.15 %	0.05 %	99.79 %
SVJ	1.06 %	1.10 %	0.05 %	99.78 %
MFSV	0.93 %	0.96 %	0.04 %	99.73 %
MFSVJ	0.87 %	0.91 %	0.04 %	99.74 %

As can be seen, the linear approximation is remarkably accurate for all four models and the correlation between the approximated difference and the true difference is close to 1. Hence, using the approximation is an almost frictionless way of tremendously decreasing computation time and complexity.

As mentioned, similar methods are used by Christoffersen, Heston & Jacobs (2009), Carr & Wu (2007), Bakshi, Carr & Wu (2008) and Trolle & Schwartz (2008a, 2008b), among others.