

Empirical Tests of Option Pricing Models

Olesia Verchenko
Kyiv, Ukraine

B.A., National University "Kyiv-Mohyla Academy", 1998
M.A., National University "Kyiv-Mohyla Academy", 2000
M.A., University of Virginia, 2002

A Dissertation Presented to the Graduate
Faculty of the University of Virginia
in Candidacy for the Degree of
Doctor of Philosophy

Department of Economics

University of Virginia
December, 2008

Thomas W. Epps
Ch. O'Keefe
Genard L. Hunt, Jr.
Marco Caporale

Abstract

This dissertation examines the empirical performance of several complete and incomplete market models of stock price dynamics using S&P 500 options and stock market data. The main contribution of this work is that it suggests and implements an empirical approach to estimating a complete model with uncertain volatility, and then judges it against other popular option pricing processes. The performance of alternative models is evaluated from four perspectives: (1) in-sample fit to stock returns data, (2) in-sample fit to options data, (3) consistency of physical and risk-neutral parameter estimates and (4) out-of-sample option pricing. Overall, the complete model with uncertain volatility is found to fit the data much better than models with constant and price-level-dependent volatilities, and the variance gamma process, and its performance is comparable to that of a stochastic volatility model.

To the memory of my father

Acknowledgments

I am truly grateful to my advisor, Dr. Thomas Wake Epps, for his advice, patience and encouragement. Without his tremendous support this dissertation would not be possible.

Contents

List of Figures	vi
List of Tables	vii
1 Introduction	1
2 Overview of Stock Price Models	4
2.1 The Black-Scholes Model and Martingale Pricing	6
2.2 The Merton Model	11
2.3 The Constant Elasticity of Variance Model	13
2.4 The Heston Model	15
2.5 The Variance-Gamma Model	17
2.6 The Hobson-Rogers Model	19
3 Empirical Methodology	22
3.1 Estimation Under Measure \mathbb{P} from Underlying Prices	22
3.1.1 The Black-Scholes Model	23
3.1.2 The CEV Model	23
3.1.3 The Heston Model	25
3.1.4 The Variance Gamma Model	29
3.1.5 The Hobson-Rogers Model	30
3.1.6 Comparing Models	31
3.2 Estimation under Measure $\widetilde{\mathbb{P}}$ with Option Prices	33
3.2.1 The Nonlinear Least Squares Algorithm	35
3.2.2 Option Pricing With Simulations	36
4 The Data	37
4.1 Index Price Series	37
4.2 Options Data	39

5	Parameter Estimates and Tests of Fit	43
5.1	Maximum Likelihood Estimation from Index Returns	44
5.1.1	The Hobson-Rogers Model	44
5.1.2	The Black-Scholes Model	48
5.1.3	The CEV Model	49
5.1.4	The Heston Model	51
5.1.5	The Variance Gamma Model	53
5.1.6	Comparing the Models	54
5.2	Least Squares Estimation from Options Data	58
5.2.1	Risk-Neutral Parameter Estimates	60
5.2.2	In-Sample Option Pricing	65
5.2.3	Implied Volatility Patterns	66
5.3	Out-of-Sample Pricing Performance	72
6	Conclusion	75
	Bibliography	78
	Appendix A: The Heston Put Formula	84
	Appendix B: The VG Put Formula	85
	Appendix C: NLLS Estimates	87
	Appendix D: Implied Volatilities	91
	Appendix E: Out-Of-Sample Fit to Option Prices	94

List of Figures

2.1	An Implied Volatility "Smile" Curve	10
4.1	S&P 500 Index Levels, January 2003-January 2007	38
4.2	S&P 500 Index Daily Returns and VIX Volatility Index	38
4.3	S&P500 Returns, Implied Volatilities and VIX Index Values, March to June 2005	42
5.1	Estimated Daily Volatilities in the Hobson-Rogers Model	48
5.2	Differences of the Individual HR and Heston Log-Likelihoods	59

List of Tables

4.1	Descriptive Statistics of S&P 500 Index Returns and VIX Series)	39
4.2	Number of Contracts by Moneyness and Maturity	41
4.3	Average Prices by Moneyness and Maturity	41
4.4	Mean Implied Volatilities by Moneyness and Maturity	41
5.1	The Discrete-Time Hobson-Rogers Model: ML Estimates of Physical Parameters Using Returns Data	46
5.2	The Discrete-Time Hobson-Rogers Model: Summary Statistics of Estimated Daily Volatilities (in annual terms)	47
5.3	The Black-Scholes Model: ML Estimates of Physical Parameters Using Returns Data	49
5.4	The Discrete-Time CEV Model: ML Estimates of Physical Parameters Using Returns Data	50
5.5	The Continuous-Time CEV Model: ML Estimates of Physical Parameters Using Returns Data	51
5.6	The Discrete-Time Heston Model: ML Estimates of Physical Parameters Using Returns Data	52
5.7	The VG Model: ML Estimates of Physical Parameters Using Returns Data	54
5.8	Likelihood Ratio Tests With Respect to the GBM Model	55
5.9	The Discrete-Time Hybrid HR-CEV Model: ML Estimates of Physical Parameters Using Returns Data	56
5.10	Non-Nested Tests	57
5.11	Black-Scholes and CEV Models: NLLS Estimates of Risk-Neutral Parameters Using Options Data	60
5.12	Heston Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data	61
5.13	Variance-Gamma Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data	61
5.14	Hobson-Rogers Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data	62

5.15 In-Sample Fit to Options Data	65
5.16 Average In-Sample Implied Volatilities in the GBM, CEV and VG Models: MSEP Loss Function	68
5.17 Average In-Sample Implied Volatilities in the Heston and HR Models: MSEP Loss Function	69
5.18 Average In-Sample Implied Volatilities in the GBM, CEV and VG Models: MSEIV Loss Function	70
5.19 Average In-Sample Implied Volatilities in the Heston and HR Models: MSEIV Loss Function	71
5.20 Out-of-Sample Fit to Option Prices: MSEP Loss Function	74

Chapter 1

Introduction

Numerous models for stock price processes have been proposed in the financial literature over the past several decades. All these models can be grouped into two broad categories: complete market models and incomplete market models. Complete markets allow replicating and therefore hedging any conceivable payoff structure using a portfolio of traded assets, while in incomplete markets such replication is generally impossible. On the other hand, incomplete market models have richer structure than complete models since they contain more sources of uncertainty.

A considerable practical interest to stock price models comes from markets for derivative securities, and options in particular.¹ In general, a good model is expected to give an adequate representation of the stock price dynamics, and, at the same time, to provide a reasonable instrument for derivative pricing. In particular, as will be discussed in the next section, a good model should be able to explain such empirical anomalies as "implied volatility smile" and term structure of implied volatilities.

The goal of this dissertation is to compare the empirical performance of several option pricing models of complete and incomplete nature, including an uncertain

¹Financial derivatives are contracts whose payoffs and prices depend upon the stochastic dynamics of associated underlying financial assets.

volatility model of Hobson and Rogers (1998). There are several motives for this research. First, even though the Hobson-Rogers model is intuitively appealing and potentially able to reproduce the stylized empirical features of the market data, and there is a growing body of recent theoretical studies devoted to its analysis, this dissertation is the first known attempt to estimate rather than calibrate all model parameters using the empirical data. Second, the empirical performance of this model is compared to that of such popular stock price processes as the classic geometric Brownian motion (Black and Scholes, 1973, and Merton, 1973), the constant elasticity of variance (CEV) model (Cox, 1975, and Cox and Ross, 1976), the variance-gamma (VG) model (Madan and Seneta, 1990) and the Heston model of stochastic volatility (Heston, 1993). Third, even though a large number of papers is devoted to the VG model, its performance is usually tested only against the GBM or the CEV models, but not against more complex models of incomplete nature. The empirical evidence of this study confirms the results of other authors that the VG model outperforms the GBM or CEV processes, but yields to a stochastic volatility model.

The main contribution of this dissertation is in developing an empirical approach to estimating the parameters of the Hobson-Rogers model using stock returns and option prices data. The general strategy is to approximate a continuous-time model by a discrete-time process, which is then used either to formulate a likelihood function to estimate model parameters from stock returns, or to simulate option prices to estimate the parameters from options data.

The empirical performance of alternative models is evaluated along four dimensions. First, in-sample fit of the models to stock price data is compared. Second, in-sample fit to the options data (both option prices and implied volatilities) is evaluated. Third, out-of-sample option pricing errors are investigated, as more complex models are generally expected to offer a superior in-sample fit, but in the case of over-

fitting their out-of-sample performance is not necessarily better. Finally, the model parameter estimates obtained from the stock price data are compared to the relevant parameter estimates implicit in option prices. The reasoning is that in a correctly specified model these two sets of parameters must be consistent.

The paper is organized as follows. In Chapter 2, a brief overview of the five alternative pricing models of interest is given. Chapter 3 discusses the empirical approaches to estimating the models parameters on the stock price time-series data and option price cross-section data. Chapter 4 describes the data. Chapter 5 presents and discusses the results. Chapter 6 concludes.

Chapter 2

Overview of Stock Price Models

A large number of stock price processes have been proposed in the financial literature over the past several decades. This chapter reviews several major stock pricing models that are the most relevant to the research subject of this thesis.

To start with, the classic geometric Brownian motion (GBM) model employed by Black and Scholes (1973) and Merton (1973) in their seminal papers is discussed. This model is then used to illustrate the basic idea of the risk-neutral approach to pricing European options. Even though this modelling framework falls short of explaining a number of empirical features of stock and option market data, it still provides an important benchmark in empirical tests of more complex models.

There are numerous extensions to the GBM model in the literature. Merton (1976) suggested to include jumps in the stock dynamics to handle the observation that stock prices can change suddenly in response to news, which usually comes at random times and has varying impact. Even though Merton's model will not be used in empirical tests in this dissertation, its discussion is helpful in terms of revealing the distinction between complete and incomplete market models.

Cox (1975, 1996) and Cox and Ross (1976) allowed the stock volatility to depend

on the stock price itself. Their constant elasticity of variance (CEV) model is another important yardstick in the option pricing literature along with the GBM model.

One more appealing modification of the baseline GBM model embraced by many researchers assumes that stock volatility is stochastic rather than constant, as in the GBM model, or dependent on the current level of the stock price, as in the CEV model. A model of this class, developed by Heston (1993), is considered next. This model is one of the most popular option pricing models in the contemporaneous empirical finance literature.

The Variance-Gamma model introduced by Madan and Seneta (1990) and further extended by Madan et al. (1998) is an alternative extension to the GBM model based on a very different modeling framework. In particular, it assumes that operational time is random and follows a gamma process, to which the Brownian motion process for stock returns is subordinated. As a result, the model is able to overcome some of the empirical biases of the GBM process. Empirical studies demonstrate that this model performs better than the GBM or CEV models, but so far no research has been done to compare the VG model to other, more complex processes.

The last but by far not the least model to be discussed is the model of path-dependent volatility recently suggested by Hobson and Rogers (1998). This dissertation is the first known attempt to estimate this model empirically and compare its performance to that of other popular models. Our finding is that this complete-market model fits the data much better than the GBM, CEV and VG models, and at least as well as the Heston model.

2.1 The Black-Scholes Model and Martingale Pricing

The geometric Brownian motion (GBM) model, also known as the Black-Scholes model, stipulates that increments of the logarithm of stock price over non-overlapping time intervals of equal length are independently and identically normally distributed. In particular, the evolution process of a stock price S_t over time is described by the stochastic differential equation

$$dS_t/S_t = \mu dt + \sigma dW_t, \quad (2.1)$$

where μ and σ are constants, and $\{W_t\}_{t \geq 0}$ is a standard Brownian motion. Then, conditional on S_t , the stock price S_T , $T \geq t$, is distributed lognormally as

$$S_T \sim S_t \exp \left\{ \left(\mu - \frac{1}{2} \sigma^2 \right) (T - t) + \sigma \sqrt{T - t} Z \right\}, \quad (2.2)$$

where $Z \sim N(0, 1)$. The model can be extended to allow deterministic variation in drift and volatility processes.

A substantial interest to the GBM process from both researchers and practitioners is driven by the elegant expression for European option prices, known as the Black-Scholes formula, that results from this model. This option pricing formula can be derived using an equivalent-martingale technique, which states that if there are no opportunities for arbitrage, then there exists a probability measure under which normalized values of traded assets are martingales. A natural candidate for a numeraire asset is the money-market fund, whose price at $t \geq 0$, given the initial value of the

fund M_0 , is determined by

$$M_t = M_0 \exp \left(\int_0^t r_u du \right),$$

where $\{r_t\}_{t \geq 0}$ is the short-rate process.

More formally, assume that a probability space consists of a set of outcomes Ω , a filtration $\{F_t\}_{0 \leq t \leq T}$ of information sets that evolve on some finite time interval $[0, T]$, an overall σ -field F with $F_T \subset F$, and a (natural) probability measure \mathbb{P} on (Ω, F) that determines distributions of prices of all assets for all $t \in [0, T]$. Then, in the absence of arbitrage, there exists an equivalent measure $\widehat{\mathbb{P}}$ such that an integrable normalized price process $\{A_t^* \equiv A_t/M_t\}_{t \geq 0}$ of any traded asset is a martingale adapted to information process $\{F_t\}_{0 \leq t \leq T}$. Moreover, if such an equivalent measure does exist, then markets offer no opportunities for arbitrage.

Therefore, the (normalized) price of any traded asset can be found as

$$A_t^* = \widehat{E}_t(A_T^*),$$

where the expectation is taken under the measure $\widehat{\mathbb{P}}$ and the subscript on \widehat{E} indicates conditioning on F_t . Additionally, assuming that interest rate process $\{r_t\}_{0 \leq t \leq T}$ is known (riskless), the (original) price of the asset is determined as

$$A_t = M_t \widehat{E}_t(A_T/M_T) = M_t/M_T \widehat{E}_t(A_T) = \exp \left(- \int_t^T r_u du \right) \widehat{E}_t(A_T).$$

In this case, the probability measure $\widehat{\mathbb{P}}$ is referred to as the risk-neutral measure because the asset's expected return is the same as that of the riskless money fund.

Black and Scholes (1973) showed that in the case of constant known riskless rate

r and dividend yield δ , the stock price dynamics under the risk-neutral measure is

$$dS_t/S_t = (r - \delta)dt + \sigma d\widehat{W}_t, \quad (2.3)$$

where $\{\widehat{W}_t\}_{t \geq 0}$ is a standard Brownian motion under $\widehat{\mathbb{P}}$. Therefore, under $\widehat{\mathbb{P}}$, the distribution of the stock price S_T , conditional on S_t , $T \geq t$, is given by

$$S_T \sim S_t \exp \left\{ \left(r - \delta - \frac{1}{2}\sigma^2 \right) (T - t) + \sigma \sqrt{T - t} \widehat{Z} \right\}, \quad (2.4)$$

where $\widehat{Z} \sim N(0, 1)$.

The price of a European put option with strike price K with maturity date T can be found as

$$P_t(S_t, T - t, \delta, r, K) = e^{-r(T-t)} \widehat{E}_t(\max(K - S_T, 0)),$$

and the celebrated Black-Scholes formula results:

$$P_t(S_t, T - t, \delta, r, K) = e^{-r(T-t)} K \Phi(d_1) - e^{-\delta(T-t)} S_t \Phi(d_2), \quad (2.5)$$

where Φ is the cumulative standard normal distribution function and

$$\begin{aligned} d_1 &= \frac{\ln\left(\frac{K}{S_t}\right) - (r - \delta - \frac{1}{2}\sigma^2)(T - t)}{\sigma \sqrt{T - t}} \\ d_2 &= d_1 - \sigma(T - t). \end{aligned}$$

Notice that if the options data rather than stock price series are used in empirical analysis, the risk-neutral rather than natural parameters of the model will be estimated. While the natural and risk-neutral probability measures are related, and may

share some common parameters, they are not identical.¹ In the case of the GBM model, the volatility parameter σ is the same under both measures (expressions (2.2) and (2.4)) and therefore can be estimated using any of the two data sets, while the natural drift parameter μ can be estimated using the stock series but not the options data.

Given that the natural and risk-neutral parameters can be obtained independently using distinct data sources, comparing these estimates, whenever it is possible, is an interesting exercise, since a large discrepancy in the parameter estimates indicates model misspecification. The existing empirical evidence suggests that for many models, including the GBM process, the estimates of the physical parameters of the stock price process obtained from stock price series are inconsistent with their risk-neutral counterparts obtained from option prices.

An extensive empirical literature documents that the GBM model fails to replicate a number of stylized features of stock and options market data. Epps (2007) surveys the literature and describes the predictions of the model that are systematically violated by the empirical data. Specifically, empirical distributions of stock returns usually have thicker tails than does a normal distribution, and squared returns are generally found to be highly predictable, indicating that return volatility is variable and persistent.

In addition, when taken to the options data, the model produces the famous implied volatility "smile". Since the volatility parameter σ is the only unobservable parameter of the model, it can be inferred from option prices by inverting the Black-Scholes formula. Implied volatility is defined as the volatility that makes the theoretical option price coincide with the observed market price. As it turns out, implied volatilities computed for options with low strike prices tend to be substantially

¹These two measures would be identical only if all relevant risks had zero risk premia.

higher than those computed for contracts with strikes that are close to and above the current price of the underlying asset. A typical relation between implied volatility and the "moneyness" of options, defined as the ratio of option's strike price to current stock price, is illustrated on *Figure 2.1*.




Figure 2.1: An Implied Volatility "Smile" Curve

These smile effects are considerably more pronounced in prices of short-term options than in prices of longer-term options, and contracts with longer maturities generally yield lower implied volatilities. Thus, there is also a term structure of implied volatilities.

Even though the GBM model remains an important benchmark in the literature, these observations have motivated the development of alternative models capable of capturing at least some of the documented empirical discrepancies.

2.2 The Merton Model

Merton (1976) proposed a jump-diffusion model for stock prices. The idea was that since new information about the economy or the company may cause stock prices to change suddenly, it is reasonable to allow asset price paths to be discontinuous. The model is intuitively appealing and addresses the issue of "fat tails" in stock returns, since jumps add mass to the tails of the distribution.

The model extends the stock price dynamics (2.1) in the following way:

$$dS_t/S_t = \mu dt + \sigma dW_t + dJ_t,$$

where $\{W_t\}_{t \geq 0}$ is a standard Brownian motion, and $\{J_t\}_{t \geq 0}$ is a compound-Poisson process with a lognormal distribution of jump sizes. That is, the jumps follow a Poisson process $\{N_t\}$ with intensity λ , independent of $\{W_t\}_{t \geq 0}$. The magnitudes of jumps are determined by random variables U such that $\ln \{1 + U_j\}_{j=1}^{\infty}$ are i.i.d. normal with parameters $\ln(1 + v) - \xi^2/2$ and ξ^2 , and which are independent of both N_t and W_t . Hence, $dJ_t = U dN_t$.

The stock price process is therefore divided into two parts: "normal vibrations", modeled by a standard Brownian motion, and "abnormal vibrations", resulting from firm-specific factors, modeled by a jump process.

In the GBM model payoffs of options can be replicated with self-financing portfolios of the underlying asset and riskless bonds or money fund, but replication is not possible in the Merton framework. The reason is that randomly-arriving jumps add an extra dimension to the space of possible states of the world that cannot be spanned by existing assets, making the markets incomplete. In general, incomplete markets allow replicating derivative assets that are linear in the terminal stock price (e.g., a

forward contract), which can be achieved via a "buy-and-hold" strategy, regardless of the underlying dynamics. In contrast, reproducing a nonlinear payoff, such as that of an option, requires a dynamic portfolio strategy involving as many traded assets as there are sources of risk. This condition is violated in incomplete markets, and dynamic portfolios of existing assets can only be used to set boundaries on prices of nonlinear derivatives, but not to uniquely value them.

This indeterminacy in prices suggests that there exist infinitely many risk-neutral measures such that a price process is a martingale.² In the context of the Merton model, Epps (2007, pp.411-412) shows that all equivalent martingale measures $\hat{\mathbb{P}}$ must have $\lambda v = \hat{\lambda} \hat{v}$, where variables with "hats" describe the parameters of the risk-neutral process, and this relation is the only restriction between the two sets of parameters. Therefore, there are infinitely many ways to choose a martingale measure.

There are two approaches to picking a martingale measure (determining its parameters) among infinitely many of them. The first approach relies on the argument that since the risk of holding a nonlinear derivative security cannot be hedged away, its price must depend on risk preferences and required compensation for risk by market participants. Then, such a preference-dependent martingale measure can be found from a general equilibrium problem that relies on strong assumptions about preferences and market processes. However, this approach is likely to be of very limited practical relevance, precisely because of these strong assumptions.

An alternative approach is to infer the parameters of the martingale measure directly from the set of traded derivative prices. One limitation of this approach is that an assumption about the structure of the stock price process under the risk-neutral measure must be made, since in general it can differ from the structure of the

²At the same time, all such martingale measures price traded assets as well as assets with replicable payoffs equivalently.

physical process. In addition, in this case there is no theoretical ground for matching up the two sets of parameters. Nevertheless, this approach is the one that is most commonly used in empirical studies.

To price options using his model, Merton (1976) suggested picking a martingale measure that has exactly the same parameters as the physical measure for the diffusion and jump parts, but with different drift. This choice is consistent with the assumption that jump risk is asset-specific and therefore diversifiable. Then, using a conditioning argument, a computational formula of European options can be developed (as in Epps (2007, pp. 413-414)).

Epps (2007, p.415) shows that the Merton model can produce both an implied volatility smile and a term structure of implied volatilities. However, Bates (1996a) found that a stochastic volatility model without jumps fits the data (currency options) even better than this jump-diffusion process.³ Therefore, a stochastic volatility model by Heston (1993) rather than the Merton model will be investigated in the empirical part of this dissertation. In addition, since the primary subject of this research is the Hobson-Rogers model, which contains no discontinuities in the stock price path, taking the Merton model as a performance benchmark seems to be less appropriate.

2.3 The Constant Elasticity of Variance Model

Cox (1975) and Cox and Ross (1976) proposed an extension to the GBM model that allows volatility to change over time but maintains the continuity of stock price paths. The model is known as the constant elasticity of variance (CEV) model. Since no new source of uncertainty is introduced, the model belongs to the class of complete market models.

³Bates (1996a) suggested a model that combines jumps with stochastic volatility, which is even better able to account for the implied volatility smile.

The stock price process is described by the stochastic differential equation

$$dS_t/S_t = \mu_t dt + \sigma_0 S_t^{\gamma-1} dW_t, \quad (2.6)$$

where $\gamma - 1$ gives the elasticity of the volatility function $\sigma_t(S_t) = \sigma_0 S_t^{\gamma-1}$ with respect to the underlying price. Important special cases of the CEV model are the GBM process ($\gamma = 1$), the absolute process ($\gamma = 0$), and the square-root process ($\gamma = 0.5$). When the coefficient γ is less than unity, there is an inverse relation between the stock price and the instantaneous volatility, which is consistent with the empirical phenomenon sometimes referred to as the "leverage effect"⁴.

As for the GBM model, when riskless rate r and dividend yield δ are known and constant, the drift process μ_t in expression (2.6) is just replaced by $r - \delta$ in translating from the physical to the risk-neutral measure.

Cox (1975) shows that for $\gamma \in (0, 1)$ the price of the European call option can be found as

$$\begin{aligned} & C(S_t, T - t, \delta, r, K) \\ = & S_t e^{-\delta(T-t)} \sum_{k=1}^{\infty} g(\xi'_t; k) G\left(\theta'_t K^{2-2\gamma}; k + \frac{1}{2-2\gamma}\right) \\ & - K e^{-r(T-t)} \sum_{k=1}^{\infty} g\left(\xi'_t; k + \frac{1}{2-2\gamma}\right) G(\theta'_t K^{2-2\gamma}; k), \end{aligned} \quad (2.7)$$

⁴The "leverage effect" describes the phenomenon of increasing stock volatility when the stock price is falling. Black (1976) and Christie (1982) suggested an explanation that focuses on the effect of change in market valuation of a firm's equity on the degree of its capital leverage. An increase in leverage produces an increase in stock risk and stock volatility.

where

$$\begin{aligned}\theta'_t &\equiv \frac{r - \delta}{\sigma_0^2(1 - \gamma)(e^{2(1-\gamma)(r-\delta)(T-t)} - 1)}, \\ \xi'_t &\equiv S_t^{2-2\gamma}\theta'_t e^{2(1-\gamma)(r-\delta)(T-t)}.\end{aligned}$$

Here $g(x; \alpha)$ is the gamma p.d.f. with shape parameter α and $G(x; \alpha)$ is the complementary gamma c.d.f. Shroder (1989) states that this formula is applicable for all $\gamma < 1$. Emmanuel and MacBeth (1982) show how to extend the formula to the case of $\gamma > 1$.

The price $P(S_t, T - t, \delta, r, K)$ of a European put option can be determined from the European put-call parity relation:

$$C(S_t, T - t, \delta, r, K) - P(S_t, T - t, \delta, r, K) = S_t e^{-\delta(T-t)} - e^{-r(T-t)} K \quad (2.8)$$

This model produces thick tails in the distribution of asset returns, and can accommodate the smirk pattern of implied volatilities. However, it cannot account for the other side of the volatility smile, and fails to produce the term structure of implied volatilities.⁵ Nevertheless, researchers still frequently use it as a performance benchmark for more complex models.

2.4 The Heston Model

In the CEV model future volatility is uncertain only because future price is uncertain. What are now called “stochastic volatility” models are those in which volatility is driven by a separate source of risk. In general, these models do not yield simple formulas for prices of options; however, Heston (1993) developed a model of this class

⁵Epps (2007) provides a more detailed analysis of the model.

that does generate computationally feasible formulas. The Heston model is described by the following two processes:

$$dS_t/S_t = \mu_t dt + \sigma_t dW_{1t}, \quad (2.9)$$

and

$$d\sigma_t^2 = \xi(\bar{\sigma} - \sigma_t^2)dt + \gamma\sigma_t(\rho dW_{1t} + \bar{\rho}dW_{2t}), \quad (2.10)$$

where $\bar{\rho} \equiv \sqrt{1 - \rho^2}$ and $\{W_{1t}, W_{2t}\}$ are two independent standard Brownian motions. The model reduces to geometric Brownian motion when γ and ξ are both zero. Unlike some other models of stochastic volatility, this setup rules out negative volatility, and it allows increments in the asset's price and volatility to be correlated.

Since (2.9) and (2.10) involve two independent sources of risk, and since volatility is not a traded asset, it is impossible to replicate a derivative asset using the underlying stock and a riskless asset. Therefore, the model is incomplete and there are infinitely many equivalent martingale measures that can be used to price options. Again, to distinguish among these equivalent measures and uniquely determine prices, one could model preferences and solve a general equilibrium problem. However, a simpler approach is just to assume that the risk-neutral dynamics are of the same general form; that is,

$$dS_t/S_t = (r - \delta) dt + \sigma_t d\widehat{W}_{1t},$$

and

$$d\sigma_t^2 = (\alpha - \beta\sigma_t^2)dt + \gamma\sigma_t(\rho d\widehat{W}_{1t} + \bar{\rho}d\widehat{W}_{2t}),$$

where \widehat{W}_{1t} and \widehat{W}_{2t} are independent Brownian motions under the risk-neutral measure. Notice that this approach basically assumes that the change of measure maintains the same type of processes, with only the drift parts being adjusted. While this choice is

essentially arbitrary, it has the virtue of making it relatively easy to price European options.

The model yields the following formula for the time- t price of a European put with strike K and expiry at T :

$$P(S_t, T-t, \delta, r, K) = e^{-r(T-t)} K F(S_t, T-t, \delta, r, K) - e^{-\delta(T-t)} S_t G(S_t, T-t, \delta, r, K), \quad (2.11)$$

with the expressions for F and G given in Appendix A.

Even though several more complex models were developed on the basis of this model (e.g., Bates (1996a), Eraker (2004), Duffie et al. (2000)), the Heston model will be used to represent the stochastic volatility class in our analysis. There are several reasons: First, it has become a sort of benchmark among successors to the GBM model in the empirical finance literature. Second, the relative ease of computing option prices simplifies the process of estimating parameters from options data. Finally, the model has the same number of parameters as does the Hobson-Rogers model, which makes their comparison more equitable.

2.5 The Variance-Gamma Model

The variance-gamma (VG) model, proposed by Madan and Seneta (1990), is an extension of the GBM model that arises from a very different modeling framework. Rather than allowing volatility to vary, the geometric Brownian motion is assumed to evolve at a pace governed by a gamma process that relates operational time to calendar time. The qualitative effect is much the same, in that the marginal distributions of stock returns are skewed and thick-tailed and option prices show volatility smiles;

however, the underlying mechanism is much simpler and more elegant.

The generalized variance gamma process for stock price of Madan et al. (1998) is given by

$$S_t = S_0 \exp \{ \mu t + \gamma \mathbb{T}_t + \sigma W_{\mathbb{T}_t} \}, \quad (2.12)$$

where $\{W_s\}_{s \geq 0}$ is Brownian motion, and the operational time $\{\mathbb{T}_t\}_{t \geq 0}$ is an independent gamma process with parameters t/ν and ν and density

$$f_{\mathbb{T}_t}(\tau) = \frac{\tau^{t/\nu-1} e^{-\tau/\nu}}{\Gamma(t/\nu) \nu^{t/\nu}}, \quad t > 0, \nu > 0, \quad (2.13)$$

where $\Gamma(x)$ is the gamma function.

Madan et al. (1998) show that the skewness of the distribution is determined by sign of the parameter γ (the distribution is symmetric when $\gamma = 0$, positively skewed when $\gamma > 0$ and negatively skewed when $\gamma < 0$), and ν controls the degree of excess kurtosis. The model implies infinitely many jumps in the stock price process in any time interval. Similarly to the Merton model, jumps make it impossible to perfectly replicate nonlinear derivatives with the underlying stock and a riskless asset.

Therefore, the VG model belongs to the class of incomplete market models, which means that explicitly relating the physical and the risk-neutral measures requires making some additional assumptions about agents' preferences. Again, it is simplest just to assume that the risk-neutral process for the stock price has the same structure as (2.12) but with different parameters:

$$S_t = S_0 \exp \left\{ (r - \delta + \theta)t + \hat{\gamma} \hat{\mathbb{T}}_t + \hat{\sigma} \hat{W}_{\hat{\mathbb{T}}_t} \right\},$$

where $\theta = \ln [1 - \hat{\nu} (\hat{\gamma} + \hat{\sigma}^2/2)] / \hat{\nu}$, $\{\hat{W}_t\}_{t \geq 0}$ is a Brownian motion under $\hat{\mathbb{P}}$, and the risk-neutral operational time process $\{\hat{\mathbb{T}}_t\}_{t \geq 0}$ has parameters $t/\hat{\nu}$ and $\hat{\nu}$. Then, to

price a European put option we can follow the approach of Madan and Milne (1991); that is, exploiting the conditional lognormality of stock price S_t , express the price of a put with strike K and expiration T as

$$P(S_t, T - t, \delta, r, K) = e^{-r(T-t)} \widehat{E}_t \max(K - S_t, 0) = \int_0^\infty p(\tau) \frac{\tau^{t/\widehat{v}-1} e^{-\tau/\widehat{v}}}{\Gamma(t/\widehat{v}) \widehat{v}^{t/\widehat{v}}} d\tau,$$

with

$$p(\tau) \equiv K\Phi(d(\tau)) - S_0 e^{(r-\delta+\theta)t + \widehat{\gamma}\tau + \frac{1}{2}\widehat{\sigma}^2\tau} \Phi(d(\tau) - \widehat{\sigma}\sqrt{\tau}),$$

where

$$d(\tau) = \frac{\ln(K/S_0) - (r - \delta + \theta)t - \widehat{\gamma}\tau}{\widehat{\sigma}\sqrt{\tau}},$$

and Φ is standard normal cumulative distribution function (see Appendix B for derivation). Madan et al. (1998) derive an explicit formula for the price in terms of modified Bessel functions. Carr and Madan (1998) suggest that it is faster to invert a Fourier transform of a “damped” price function.

The VG model was found to outperform the GBM model for S&P 500 index and options data (Madan et al. (1998), Lam et al. (2002)), as well as the Merton model for foreign currency options (Daal and Madan (2005)). In particular, the model offers a superior fit to the moneyness and maturity structures of implied volatilities. However, no one has yet compared the VG model with more competitive models such as that of Heston (1993).

2.6 The Hobson-Rogers Model

Our main focus in this dissertation is on a complete-market model recently suggested by Hobson and Rogers (1998), in which volatility depends on the historical

sample path of the underlying price. The assumption is that the discounted log-price process, $\tilde{S}_t = \ln(e^{-\alpha_t t} S_t)$, where α_t is the factor that corrects for trend, is an Ito process of the form

$$d\tilde{S}_t = \mu(\Delta_t)dt + \sigma(\Delta_t)dW_t. \quad (2.14)$$

Here Δ_t is referred to as the "offset" function, defined as

$$\Delta_t = \int_0^\infty \lambda e^{-\lambda u} (\tilde{S}_t - \tilde{S}_{t-u}) du. \quad (2.15)$$

The function Δ_t resides at zero as long as the stock price does not deviate from its trend, and is "set off" by unexpected price movements.⁶ The constant $\lambda > 0$ describes the rate at which past information is discounted. Intuitively, the offset function summarizes the history of stock price shocks, assigning higher weight to more recent events and making the effect of older shocks dissipate with time.

The appeal of the model is that any functional form of $\mu(\Delta_t)$ and $\sigma(\Delta_t)$ that makes economic sense can be assumed. For instance, market turmoil is likely to result in an increase of expected volatility by market participants, which can be captured by a quadratic volatility function of the form $\sigma(\Delta_t) = \eta\sqrt{1 + \varepsilon\Delta_t^2} \wedge N$, where N is a some large constant. Hobson and Rogers (1998) have shown that this simple model specification can accomodate much richer moneyness and term structures of volatilities than the CEV model, which also has a price driven volatility process.

Despite the appeal of model, there are still relatively few research papers devoted to it. Di Francesco and Pascucci (2004) and Di Francesco, Foschi and Pascucci (2006) focus on the numerical solution to the option pricing problem in the model by several finite-differences schemes. Hallulli and Vargiolu (2005) analyze the issues related to

⁶A more general version of the model defines the discounted log-price process \tilde{S}_t in terms of the offset function of order m , denoted by Δ_t^m , where $\Delta_t^m = \int_0^\infty \lambda e^{-\lambda u} (\tilde{S}_t - \tilde{S}_{t-u})^m du$.

finite observation horizon and specification of the offset function. Among empirical studies, Foschi and Pascucci (2005) propose a calibration procedure for the volatility function $\sigma(\Delta_t)$ in (2.15) and then test it with S&P 500 option data. Platania and Rogers (2005) also calibrate the model using S&P 500 option data, and then compare its performance to the performances of the Black-Scholes model, the CEV and the Heston models. Figa-Talamanca and Guerra (2006) present a technique to estimate the discount parameter λ , and they calibrate the parameters of the volatility function $\sigma(\Delta_t)$ using sets of the S&P 500 and FTSE 100 options. However, so far as we know, no one has tried to estimate all the parameters of the model simultaneously from historical underlying price data. Our research aims at filling this gap in the literature. The problem of obtaining empirical estimates of the model parameters becomes particularly relevant in the light of the claim by Hubalek, Teichmann and Tompkins (2005) that unreasonable values for the parameters would be required to match the observed smiles and term structure in implied volatilities.⁷

⁷Hubalek, Teichmann and Tompkins (2005) suggest a generalization of the Hobson-Rogers model that, according to the authors, provides a better fit to the market data.

Chapter 3

Empirical Methodology

Following the literature, we use two different approaches to estimate the parameters of the models described above. The first approach allows estimating the parameters of the models under physical measure \mathbb{P} using the returns on the underlying asset exclusively. In the second approach, the risk-neutral parameters under measure $\hat{\mathbb{P}}$ are estimated using cross-sections of option prices. For complete-markets models, these sets of parameters can be mapped directly into each other, and therefore their comparison is straightforward. For incomplete-market models such mapping is generally possible only for a subset of parameters.

3.1 Estimation Under Measure \mathbb{P} from Underlying Prices

To estimate the physical parameters of the five models, we use the standard maximum-likelihood approach. However, since the conditional probability density functions of underlying price in the Hobson-Rogers and Heston models are not available in closed form, we must work with discrete-time versions of these two models. As

to the GBM, CEV and VG models, both discrete-time and continuous-time versions of these models can be estimated. Indeed, the two versions of the GBM and VG models are essentially identical. As to the CEV model, the two versions will be estimated separately to facilitate comparison with Hobson-Rogers and Heston. Comparing the two sets of results for CEV will shed light on the extent of discretization bias.

3.1.1 The Black-Sholes Model

Given a known starting value of underlying price S_{t-1} , the solution to equation (2.1) for price after $\Delta t = 1$ time units is

$$S_t = S_{t-1} \exp [\mu - \sigma^2/2 + \sigma (W_t - W_{t-1})], \quad (3.1)$$

or, in terms of the continuously compounded return over $(t-1, t]$,

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu - \sigma^2/2 + \sigma (W_t - W_{t-1}). \quad (3.2)$$

The likelihood function of a sample $\{R_t\}_{t=1}^n$ of returns is easily constructed, since increments $\{W_t - W_{t-1}\}$ over non-overlapping intervals of length $\Delta t = 1$ are identically and independently distributed (i.i.d.) as $N(0, 1)$. Notice that the standard Euler discretization scheme applied to (2.1) would omit the term $\sigma^2/2$ in the drift.

3.1.2 The CEV Model

The problem of estimating the parameters of the CEV process from the stock market data is not new to the literature. A number of early empirical studies, including Macbeth and Merville (1980), Beckers (1980), Christie (1982), Ang and Peterson (1984), and a more recent study by Yuen et al. (2001) employ least-squares methods

for this purpose. Ours appears to be the first application of maximum likelihood.

The discrete-time version of the model for continuously compounded returns over a unit time interval is formulated as

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu_t - \sigma_t^2/2 + \sigma_t Z_t, \quad (3.3)$$

where $Z_t \sim N(0, 1)$ and $\sigma_t = \sigma_0 S_{t-1}^{\gamma-1}$. Again, returns over nonoverlapping periods of equal length are i.i.d. We will experiment with two specifications for the $\{\mu_t\}$ process: (1) $\mu_t = \mu$, and (2) $\mu_t = \beta_1 + \beta_2 \sigma_t + \beta_3 \sigma_t^2$. The latter formulation allows expected return to depend on volatility, although it is not consistent with the usual view that it is *nondiversifiable* risk that matters. Since returns are conditionally normal in this model, it is again straightforward to express the likelihood function of sample $\{R_t\}_{t=1}^n$.

To estimate the original continuous-time version of model (2.6), we will use the expression for the risk-neutral probability density function of stock price derived by Shroder (1989) under the assumption of constant drift $\mu_t = \mu$. For this we take the approximation provided by Epps (2007), under which the conditional p.d.f. of stock price is

$$f^{CEV}(S_{t+1} | S_t) = \sum_{k=1}^{\infty} g\left(\xi_t; k + \frac{1}{2 - 2\gamma}\right) g(\theta S_{t+1}^{2-2\gamma}; k) \theta(2 - 2\gamma) S_{t+1}^{1-2\gamma},$$

where

$$\begin{aligned} \theta_t &\equiv \frac{\mu}{\sigma_0^2(1 - \gamma)(e^{2\mu(1-\gamma)} - 1)}, \\ \xi_t &\equiv S_t^{2-2\gamma} \theta_t e^{2\mu(1-\gamma)}, \end{aligned}$$

and $g(x; \alpha)$ is the gamma p.d.f. with unit scale and shape parameter α . This expression is used to set up the conditional likelihood function for the continuous-time

model.

3.1.3 The Heston Model

Different ways of discretizing stochastic volatility (s.v.) models have been proposed in the literature. The standard Euler scheme yields the following discrete-time equivalent of the Heston model:

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu_{t-1} - \sigma_{t-1}^2/2 + \sigma_{t-1}Z_{1t} \quad (3.4)$$

and

$$\sigma_t^2 - \sigma_{t-1}^2 = \xi(\sigma_\infty - \sigma_{t-1}^2) + \gamma\sigma_{t-1}(\rho Z_{1t} + \bar{\rho}Z_{2t}), \quad (3.5)$$

where Z_{1t} and Z_{2t} are i.i.d. as standard normal and $\bar{\rho} \equiv \sqrt{1 - \rho^2}$. This procedure is easy to implement, and Eraker et al. (2003) show that the resulting discretization bias is negligible with daily data. Although the process for $\{\mu_t\}$ could be modelled as a function of $\{\sigma_t^2\}$, we shall simplify the already high-dimensional estimation problem by putting $\mu_t = \mu$. We show later in the context of other models that allowing non-constant drift adds little to explanatory power.

Estimating both discrete- and continuous-time versions of a stochastic volatility model using only the time series of underlying prices poses substantial challenges. Since volatility σ_t is unobservable, expressing the likelihood function of observed returns requires this latent variable to be integrated out of the density. However, the exact closed-form expression for the marginal density function of S_t is not available. Several empirical techniques to solve this problem have been suggested in the literature. Among other authors, Durham and Gallant (2002) and Brandt and Santa-Clara (2002) use simulated maximum likelihood; Bakshi, Cao and Chen (1997) resort to cal-

ibration; Pan (2002) employs an implied-state GMM approach; Eraker *et al.* (2003) develop a likelihood-based MCMC approach; and Ait-Sahalia and Kimmel (2007) use an approximate maximum likelihood method that replaces the characteristic function of unobserved volatility by a gamma characterisitic function at each step.

Since the goal of this paper is to compare the various models, it is important that comparable estimation approaches be used. Therefore, rather than adopting any of the methods mentioned above, we will employ a technique that allows calculating the marginal likelihoods directly. To do this we discretize the support of unobservable volatility as well as time. This approach is in the same spirit as the approximate maximum likelihood method used by Bates (2006), but it computes the conditional distribution function of volatility at each time step directly from the data rather than approximating it with a formal model. Although our approach was developed independently, it is similar to the method of Friedman and Harris (1998), who apply it to a different type of s.v. model. They show in that application that it produces comparable results to those obtained with other popular estimation techniques.

The following algorithm describes the method.

Step 0: At $t = 0$, initialize $\sigma_t^2 = \sigma_0^2$, where σ_0 is treated as a parameter to be estimated.

Step 1: At $t = 1$, given σ_0^2 , an expression for σ_1^2 conditional on the data observed at $t = 1$ can be deduced from expressions (3.4) and (3.5) as

$$\sigma_1^2 - \sigma_0^2 = \xi(\sigma_\infty - \sigma_0^2) + \gamma\rho(R_1 - \mu + \sigma_0^2/2) + \gamma\sigma_0\bar{\rho}Z, \quad (3.6)$$

where $Z \sim N(0, 1)$. Then, σ_1^2 is distributed conditionally on R_1 and σ_0 as normal with mean

$$E\sigma_1^2 = \sigma_0^2 + \xi(\sigma_\infty - \sigma_0^2) + \gamma\rho(R_1 - \mu_1 + \sigma_0^2/2)$$

and variance $V\sigma_1^2 = (\gamma\sigma_0\bar{\rho})^2$. The support of this distribution is then partitioned into k points, the grid being centered at the conditional mean and stretching out to m conditional standard deviations from the mean in each direction.

The value of the probability distribution function $f_{\sigma_1^2|\sigma_0^2}$ at each cell $j \in \{1, 2, \dots, k\}$ is then computed as $f_{\sigma_1^2|\sigma_0^2}(j) = \Phi_{\sigma_1^2|\sigma_0^2}(\sigma_{1,j}^2) - \Phi_{\sigma_1^2|\sigma_0^2}(\sigma_{1,j-1}^2)$, where $\Phi_{\sigma_1^2|\sigma_0^2}(\cdot)$ is the normal cumulative distribution function with mean $E\sigma_1^2$ and variance $V\sigma_1^2$, and $f_{\sigma_1^2|\sigma_0^2}(j) = 1 - \Phi(Z_{t,k-1})$ when $j = k$. Finally, cells with negative values for σ_1^2 are assigned zero probabilities, and the probabilities in the remaining cells are normalized to sum to unity.

Step 2: At each $t \in \{2, 3, \dots, n\}$, given the mean and the variance of σ_{t-1}^2 and the data observed at t , the mean and the variance of σ_t^2 can be computed as

$$E\sigma_t^2 = E\sigma_{t-1}^2 + \xi(\sigma_\infty - E\sigma_{t-1}^2) + \gamma\rho(R_t - \mu_1 + E\sigma_{t-1}^2)$$

and

$$V\sigma_t^2 = (1 - \xi + 0.5\gamma\rho)^2 V\sigma_{t-1}^2 + (\gamma\rho)^2 E\sigma_{t-1}^2.$$

The grid for σ_t^2 is then constructed in the same way as at Step 1, and the probabilities are calculated as

$$f_{\sigma_t^2}(j) = \sum_{i=1}^k f_{\sigma_t^2|\sigma_{t-1}^2=\sigma_{t-1,i}^2}(j) f_{\sigma_{t-1}^2}(i),$$

where $f_{\sigma_t^2|\sigma_{t-1}^2=\sigma_{t-1,i}^2}$ is the probability of σ_t^2 conditional on $\sigma_{t-1}^2 = \sigma_{t-1,i}^2$. This is computed as

$$f_{\sigma_t^2|\sigma_{t-1}^2=\sigma_{t-1,i}^2}(j) = \Phi_{\sigma_t^2|\sigma_{t-1}^2=\sigma_{t-1,i}^2}(\sigma_{1,j}^2) - \Phi_{\sigma_t^2|\sigma_{t-1}^2=\sigma_{t-1,i}^2}(\sigma_{1,j-1}^2).$$

Again, zero probabilities are assigned to cells with negative values of σ_t^2 , and remaining

probabilities are normalized to sum to unity.

Step 3: Given parameter values and conditional on the past data, the value of the likelihood function for the return observed at time t is computed as

$$\sum_{i=1}^k f_{R_t|\sigma_{t-1}^2}(R_t) f_{\sigma_{t-1}^2}(i),$$

where $f_{R_t|\sigma_{t-1}^2}(R_t)$ is the probability distribution function of stock return R_t conditional on σ_{t-1}^2 , which is normal with mean and variance implied by expression (3.4).

After experimenting with the number of grid points and the spread of the distribution, the values for parameters k and m were set at 150 and 3 respectively. This setup ensures that the grid covers virtually all of the support of the conditional distribution of σ_t^2 .

For completeness, one more estimation technique for the Heston model should be mentioned. Bates (2006), Bakshi, Cao and Chen (1997), Chernov and Ghysels (2000), Pan (2002), Ait-Sahalia and Kimmel (2007) and other authors employ both option prices and underlying price series in estimation. Option prices are used to infer current levels of volatility, which are then employed in estimation as observable quantities. Two implications of this approach make it less suitable for our purposes. First, to have comparable results, all option pricing models must be estimated using the same data set. Second, we would like to be able to perform a specification test that compares the physical parameter estimates with the risk-neutral estimates from options. Since estimating a model using both stock and option price data implicitly imposes restrictions on risk-neutral and physical parameters of the model, this test becomes less meaningful.

Despite the mentioned downsides of this estimation technique for the Heston model, it is still interesting to compare its results to those obtained with our method.

Ait-Sahalia and Kimmel (2007) suggested a shortcut to this approach, which is much less computationally intensive than the exact procedure, and, according to these authors, has negligible approximation error for the Heston model. The idea is to proxy unobserved volatility with the VIX index rather than infer it from market option prices.¹ We will use this alternative technique as a complementary estimation method for the discrete-time Heston model.

3.1.4 The Variance Gamma Model

The VG model is given by the following expression

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu + \gamma\nu\Gamma_t + \sigma\sqrt{\nu\Gamma_t}Z_t, \quad (3.7)$$

where $Z_t \sim N(0, 1)$ and $\Gamma_t \sim \text{Gamma}(1)$. The conditional likelihood function can be constructed using the fact that the continuously compounded return is distributed as normal conditionally on Γ_t , then integrating the conditional p.d.f. $f(R_t|\Gamma_t)$ over the support of Γ_t . Madan et al. (1998) show that the unconditional density of stock returns is given by the following expression:

$$f_{R_t}(R) = \frac{2e^{\gamma(R-\mu)/\sigma^2}}{\nu^{1/\nu}\sqrt{2\pi\sigma^2}\Gamma(1/\nu)} \left(\left| \frac{R-\mu}{\sqrt{2\sigma^2/\nu + \gamma^2}} \right| \right)^{1/\nu-0.5} * \\ K_{1/\nu-0.5} \left(\frac{|(R-\mu)\sqrt{2\sigma^2/\nu + \gamma^2}|}{\sigma^2} \right),$$

where K is the modified Bessel function of the second kind. This expression can be used to construct the conditional likelihood function of stock returns.

¹VIX is the volatility index of the Chicago Board of Options Exchange (CBOE) introduced in 1993. The value of the index is computed using SP500 option prices of maturities up to two months.

3.1.5 The Hobson-Rogers Model

To estimate the parameters of the Hobson-Rogers model from underlying price data, we use the following discrete-time version:

$$\tilde{S}_t - \tilde{S}_{t-1} = \mu_{t-1} - \sigma_{t-1}^2/2 + \sigma_{t-1}Z_t, \quad (3.8)$$

where $Z_t \sim N(0, 1)$ and $\tilde{S}_t = \ln(e^{-\alpha t} S_t)$. Similarly, we define the discrete-time equivalent of the offset function as

$$\Delta_t = \sum_{j=1}^{\infty} \frac{\lambda e^{-\lambda j}}{A} (\tilde{S}_t - \tilde{S}_{t-j}), \quad (3.9)$$

where division by $A = \sum_{i=1}^{\infty} \lambda e^{-\lambda i} = \frac{\lambda}{e^{\lambda}-1}$ makes the weights sum to unity. This can be re-written in the recursive form

$$\Delta_t = \sum_{j=1}^{\infty} \frac{\lambda e^{-\lambda j}}{A} (\tilde{S}_t - \tilde{S}_{t-j}) = \dots = (\tilde{S}_t - \tilde{S}_{t-1}) + e^{-\lambda} \Delta_{t-1}. \quad (3.10)$$

This form is particularly useful for estimation.

We experiment with three possible specifications for the drift process $\{\mu_t\}$: (1) $\mu_t = \beta_1$, (2) $\mu_t = \beta_1 + \beta_2 \sigma_t + \beta_3 \sigma_t^2$ and (3) $\mu_t = \beta_1 + \beta_2 \Delta_t + \beta_3 \Delta_t^2$. The second specification is intended to account for the compensation for risk, while the third is a polynomial approximant of some more general function of offset. The variance σ_t^2 is assumed to be a quadratic function of offset Δ_t , as

$$\sigma_t^2(\Delta_t) = \eta_1^2 (1 + \eta_2 \Delta_t + \eta_3 \Delta_t^2). \quad (3.11)$$

This can accommodate different responses of variance to positive and negative offsets. The disadvantage is that it allows variance to be negative for certain combinations

of parameters and offset, but in the estimation stage constraints will be imposed to eliminate this problem.²

Notice that $R_{t+1} \equiv \ln(S_{t+1}/S_t)$, is distributed as normal conditionally on Δ_t ; i.e.,

$$R_{t+1} | \Delta_t \sim N(\mu(\Delta_t) + \alpha_t, \sigma^2(\Delta_t)). \quad (3.12)$$

3.1.6 Comparing Models

Our main goal is to compare the descriptive accuracy of the HR model to that of the other four models. It is easy to see that GBM is nested within HR and all the other models, so that comparison of goodness of fit under measure \mathbb{P} can be based on the standard likelihood-ratio test. Unfortunately, comparisons of the CEV, HR, Heston and VG models are much more complicated. For HR versus CEV, we can propose the following hybrid model that nests both:

$$dS_t/S_t = \mu_t dt + \sigma_t dW_t, \quad (3.13)$$

$$\sigma_t^2(\Delta_t, S_t) = \eta_1^2 (1 + \eta_2 \Delta_t + \eta_3 \Delta_t^2) S_t^{2(\gamma-1)}. \quad (3.14)$$

When $\eta_2 = \eta_3 = 0$, this reduces to CEV, and when $\gamma = 1$ it reduces to HR. This specification makes it easy to test whether the CEV term $S_t^{2(\gamma-1)}$ helps to explain the dynamics once the HR term $\eta_1^2 (1 + \eta_2 \Delta_t + \eta_3 \Delta_t^2)$ is accounted for.

Unfortunately, there is no obvious way to combine HR with Heston or VG in such a way, so we rely on formal tests of non-nested models in these cases. Common model selection criteria include Akaike's information criterion (Akaike, 1974) and Schwartz's Bayesian Information Criterion (Schwartz, 1978). However, these measures simply

²The exponential form $\sigma^2(\Delta_t) = \eta_1^2 \exp\{1 + \eta_2 \Delta_t + \eta_3 \Delta_t^2\}$ was also considered, but it resulted in a significant drop in the value of the likelihood function.

rank the models and do not constitute formal tests. Cox's test (Cox, 1961) for non-nested hypotheses is such a formal test, but it is difficult to carry out with these elaborate models.

Vuong (1989) proposed a testing procedure that is easier to implement. It relies on a popular measure of distance between distribution functions known as the Kullback-Leibler information criterion. Following Clarke's (2007) notation, consider two models, $F_\beta = f(Y_i | X_i; \beta)$ and $G_\eta = g(Y_i | Z_i; \theta)$, where f and g are probability density functions with p and q parameters respectively. Under the null hypothesis that the two models are observationally equivalent as judged by their likelihoods, Vuong's test statistic is:

$$\frac{LR^*(\hat{\beta}, \hat{\theta})}{\sqrt{n\hat{\sigma}}},$$

where $\hat{\beta}$ and $\hat{\theta}$ are maximum likelihood estimators and n is sample size; $LR^*(\hat{\beta}, \hat{\theta})$ is the ratio of likelihood functions with a correction for degrees of freedom,

$$LR^*(\hat{\beta}, \hat{\theta}) = \sum_{i=1}^n \left[\ln f(Y_i | X_i; \hat{\beta}) - \ln g(Y_i | Z_i; \hat{\theta}) \right] - \frac{1}{2} [p - q] \ln n,$$

and $\hat{\sigma}$ is the estimated standard deviation of the differences in log-likelihoods, computed as

$$\hat{\sigma}^2 \equiv \frac{1}{n} \sum_{i=1}^n \left[\ln \frac{f(Y_i | X_i; \hat{\beta})}{g(Y_i | Z_i; \hat{\theta})} \right]^2 - \left[\frac{1}{n} \sum_{i=1}^n \ln \frac{f(Y_i | X_i; \hat{\beta})}{g(Y_i | Z_i; \hat{\theta})} \right]^2.$$

Vuong (1989) shows that this statistic converges in distribution to $N(0, 1)$ under the null that the two models have equivalent explanatory power. He also shows that the test is asymptotically equivalent to the classical LR test when the two models are nested. Vuong's test will be used as a formal criterion to select the best model.

Recently, Clarke (2003) suggested a distribution-free test of non-nested models. This test utilizes the signs of the differences in the individual log-likelihoods rather

than their actual values. The advantage of this nonparametric test is that it does not rely on the asymptotic theory underlying Vuong's procedure. The test is as follows. Letting $d_i = f^*(Y_i | X_i; \hat{\beta}) - g^*(Y_i | Z_i; \hat{\theta})$, where $f^* \equiv \ln f - (\frac{p}{2n} \ln n)$ and $g^* \equiv \ln g - (\frac{q}{2n} \ln n)$ represent the individual log-likelihoods corrected for dimensionality, the test statistic is

$$B = \sum_{i=1}^n I_{(0,+\infty)}(d_i),$$

where I is the indicator function. This is distributed as binomial with parameters n and $\frac{1}{2}$ under the null that the two models have the same explanatory power. Below we report the results of both the Vuong and Clarke tests in comparing the non-nested models.

3.2 Estimation under Measure $\hat{\mathbb{P}}$ with Option Prices

The previous section described estimation of the physical, or natural, parameters of the various models. However, to value derivative assets, risk-neutral parameters are required. There are two general ways to estimate these.

The first approach relies on the theoretical relation between the physical and risk-neutral parameters. When the model is complete, the parameters are the same, so maximum likelihood estimates from the previous section can be used. However, it is well known that even in the complete-markets case the risk-neutral parameters can be quite far from those obtained from the underlying prices. For example, Christoffersen and Jacobs (2004) show that the fit of option pricing models deteriorates drastically when pricing formulas are evaluated with physical parameter estimates.

In the case of incomplete models, the two sets of parameters can be related if additional assumptions about agents' preferences are made and a general equilibrium

model is formulated. Unfortunately, the hierarchy of strong assumptions that this requires leaves one with little confidence in the results. Moreover, one must recognize that the historical price data used in estimated physical parameters are backward looking and may not reflect current beliefs about the future dynamics of the underlying asset. Therefore, even though a finding of consistency between risk-neutral and physical estimates would clearly support a model, one cannot rely on it for pricing options.

For these reasons the common approach is to estimate risk-neutral parameters directly from options data. This is usually done by applying nonlinear least squares to minimize the sum of squared differences between market prices (market implied volatilities) and prices (implied volatilities) based on a particular model. Most often, parameters are estimated separately from cross-sections of option prices having different strikes and expiration dates that are traded on the same day.

This section describes how this procedure was implemented to estimate the risk-neutral parameters of the five alternative models. Two strategies were pursued. First, the (computational) closed-form formulas for option prices available for the GBM, CEV, VG and Heston models are used to represent model-based prices. Second, since no such option pricing formula is available for HR, model-based prices have to be estimated by simulation. To facilitate comparison with the other models, separate simulation-based estimates were also obtained for all except GBM. Having dual results for three of the models makes it possible to judge the extent of bias due to discretization and the sampling error associated with simulation.

Two criteria are employed to judge the various models: (1) in-sample fit, as measured by mean-squared errors; and (2) errors from out-of-sample forecasts. For the latter, the parameter values implied by any particular day's data are used as inputs for model-based option prices one to five days ahead. In addition, consistency of physical

and risk-neutral parameters of the models is analyzed whenever this is possible.

3.2.1 The Nonlinear Least Squares Algorithm

To describe the estimation in more detail, let Θ be the vector of parameters of a given model. (This includes the initial values of volatility and offset in the Heston and HR models, respectively.) At date t we collect N_t put option prices, where N_t exceeds the maximum number of parameters of the five models. Let P_t^j be the observed price of option $j \in \{1, 2, \dots, N_t\}$ with time-to-maturity T^j and strike price K^j ; and let $\widehat{P}_t^j(\Theta_t)$ be the price determined from the model, as computed either from a formula or by simulation.

Two different loss functions are considered: mean squared errors of prices,

$$MSEP_t(\widehat{\Theta}_t) \equiv \min_{\Theta_t} \sqrt{\frac{1}{N_t} \sum_{j=1}^{N_t} [\widehat{P}_t^j(\Theta_t) - P_t^j]^2}, \quad (3.15)$$

and mean squared errors of implied volatilities,

$$MSEIV_t(\widetilde{\Theta}_t) \equiv \min_{\Theta_t} \sqrt{\frac{1}{N_t} \sum_{j=1}^{N_t} [\widehat{V}_t^j(\Theta_t) - V_t^j]^2}. \quad (3.16)$$

The mean-squared pricing error is most commonly used in empirical studies, but it puts most of the weight on options that are deep in the money those with longer maturities, both of which have relatively high prices; and these are the classes of options that the standard Black-Scholes model handles reasonably well. These effects of “moneyness” and term are mitigated by converting model and market prices to implicit volatilities. This MSEIV criterion puts more weight on options that are far from the money in either direction and especially on those with short maturities.

These are precisely the options whose prices are least satisfactorily explained by the standard Black-Scholes formulas. Since both optimization problems are highly nonlinear in the parameters and there are no analytical solutions for the gradient, numerical optimization techniques must be used. For this purpose we made use of several different routines from the IMSL and Numerical Recipes libraries for Fortran 77.

3.2.2 Option Pricing With Simulations

To compute option prices by simulation the following steps are taken. First, for each options sample at date t a model-specific risk-neutral price process was used to generate a sample path for underlying price at daily frequency, starting with S_t and extending to the longest maturity date in the sample of options. The exercise value of each option $j \in \{1, 2, \dots, N_t\}$ at the appropriate expiration date was then computed as $\max(K^j - S_{T^j}, 0)$. After M such paths were generated, the price of each option was estimated as

$$P_t^j = e^{-r(T^j-t)} \frac{1}{M} \sum_{m=1}^M \max(K^j - S_{T^j}^m, 0).$$

Our estimates were based on $M = 100,000$ sample paths. This large number was needed to give the simulated prices a degree of accuracy comparable to those obtained from the explicit formulas. The antithetic variable technique was used throughout to reduce standard errors. This involved pairing each sequence $\{Z_t\}$ of normal variates with a sequence of identical magnitude but opposite sign.

Chapter 4

The Data

The procedures described above were carried out with prices and options for the S&P 500 index, a value-weighted index of common stocks meant to be representative of various industries. Listed options on the S&P are of the European type, to which our pricing formulas apply, and they are very actively traded. For these reasons they have become the standard vehicle for testing advanced theories of option pricing. Here we describe some general features of the data. How well they are fit by the various models will be considered in the following section.

4.1 Index Price Series

The physical parameters of all models are estimated from daily closing prices of the S&P for the three years beginning on January 1, 2004—a total of 756 observations. Figures 4.1 and 4.2 depict time-series plots of index levels and daily returns, along with the time path of the volatility index VIX.

The graphs illustrate the variable and persistent nature of volatility, a feature that




Figure 4.1: S&P 500 Index Levels, January 2003-January 2007

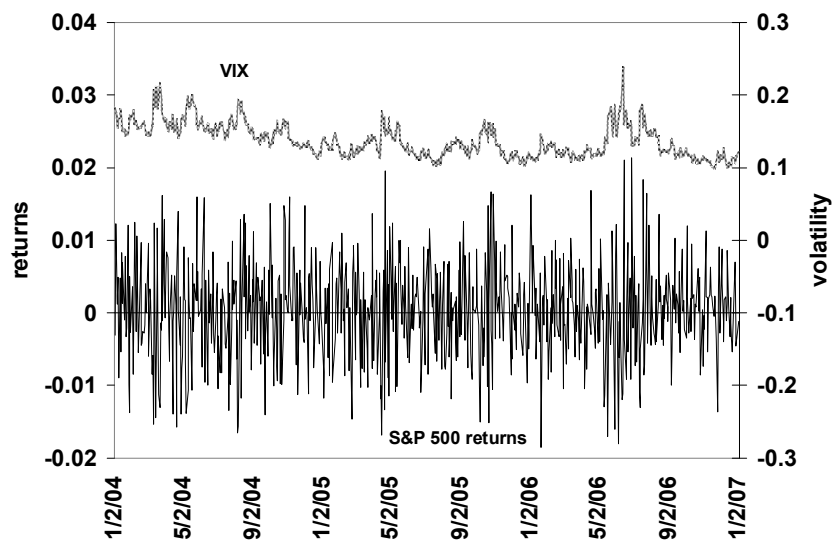


Figure 4.2: S&P 500 Index Daily Returns and VIX Volatility Index

Table 4.1: **Descriptive Statistics of S&P 500 Index Returns and VIX Series.**

The table describes daily returns and VIX volatility series (in annual terms) collected from January 1, 2004 to January 1, 2007.

	S&P 500 Daily Returns	VIX
Mean	0.081	0.137
Standard Deviation	0.105	0.023
Median	0.182	0.132
Min	-4.661	0.099
Max	5.377	0.238
1st quartile	-0.928	0.119
3rd quartile	1.073	0.152
Skewness	-0.019	0.805
Excess kurtosis	0.256	0.440
Correlation of returns and VIX	-0.205	—

several of our models are intended to capture. Table 4.1 summarizes the data.¹

4.2 Options Data

To estimate the models' risk-neutral parameters, we use daily option prices for the period from March 23, 2005 to May 23, 2005—a total of 43 days. The data were collected from the Chicago Board Options Exchange (CBOE) website. For each date the market price is taken as the average of the last reported bid and asked quotations. The data were filtered as follows. Excluded from the sample were (1) options with fewer than 20 days and more than 252 days to expiration; (2) options with trading volume below 50 contracts or open interest less than 100; (3) options with prices below \$3/8; (4) options whose calculated prices violated no-arbitrage bounds:

$$\max \{0, e^{-r(T-t)}K - e^{-\delta(T-t)}S_t\} \leq P_t \leq e^{-r(T-t)}K$$

¹At an earlier stage of this research, we have also experimented with a much longer sample, covering the period from January 1, 1997 till January 1, 2007. However, the use of this extended sample made it much harder to rationalize the consistency tests of physical and risk-neutral parameters, since the options data described below were sampled over considerably shorter time interval.

and (5) options that were far away from the money ($K/S_t < 0.8$ and $K/S_t > 1.3$). Excluding such low-priced and illiquid options gives us more confidence in the currency of price quotations and mitigates the effect of price discreteness. The same sort of filtering process has been used by Bakshi, Cao and Chen (1997), Christoffersen and Jacobs (2004) and other authors. To broaden the sample of put prices across the range of strikes, the European put-call parity relation (2.8) was used to transform calls into puts when put quotations were not available or did not meet the selection criteria.

Besides prices of the options themselves, we require data for daily Treasury-bill rates and S&P 500 futures prices in order to estimate implicit dividend rates on the index. Assuming that the riskless rate is known over the life of a futures contract with maturity T and price $F_{t,T}$, the dividend yield δ of the underlying asset can be calculated as

$$\delta = r + \frac{1}{T-t} \ln \left(\frac{F_{t,T}}{S_t} \right).$$

Treasury-bill rates and S&P 500 futures prices were obtained from Barron's.

After filtering and transforming selected calls, we were left with a total sample of 2,777 puts, the daily number ranging from 44 to 105 contracts. All options are divided into several categories according to their moneyness and maturity. Following standard classification, a put option is said to be *at-the-money* (ATM) if the ratio of its strike price to current stock price, K/S_t , is between 0.97 and 1.03, a put is *in-the-money* if $K/S_t \geq 1.03$ and *out-of-the-money* if $K/S_t \leq 0.97$. By the term to expiration, an option is *short-term* if $(T-t)$ is less than 60 days, *medium-term* if $(T-t)$ is between 60 and 180 days and *long-term* if $(T-t)$ is more than 180 days.

Tables 4.2, 4.3, and 4.4 summarize the data by moneyness and maturity and illustrate the inclusiveness of the sample.

Table 4.2: Number of Contracts by Moneyness and Maturity

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
0.80-0.94	352	138	121	205	816
0.94-0.97	142	56	34	58	290
0.97-1.00	185	67	44	97	393
1.00-1.03	236	87	46	86	455
1.03-1.06	237	78	46	56	417
1.06-1.33	132	55	63	156	406
Total	1,284	481	354	658	2,777

Table 4.3: Average Prices by Moneyness and Maturity

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
0.80-0.94	2.39	4.89	9.16	16.58	7.38
0.94-0.97	6.66	12.25	22.03	33.83	14.98
0.97-1.00	14.08	20.64	31.16	45.12	24.77
1.00-1.03	29.37	32.54	43.53	57.51	36.72
1.03-1.06	56.24	53.63	64.35	76.13	59.32
1.06-1.33	102.44	98.14	116.14	133.07	115.75
Total	46.78	52.26	39.46	87.26	39.09

Table 4.4: Mean Implied Volatilities by Moneyness and Maturity

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
0.80-0.94	0.2037	0.1767	0.1775	0.1879	0.1913
0.94-0.97	0.1551	0.1387	0.1509	0.1581	0.1520
0.97-1.00	0.1341	0.1284	0.1403	0.1497	0.1377
1.00-1.03	0.1185	0.1154	0.1287	0.1425	0.1235
1.03-1.06	0.1098	0.1031	0.1215	0.1358	0.1133
1.06-1.33	0.1274	0.1035	0.1145	0.1227	0.1204
Total	0.1475	0.1341	0.1455	0.1538	0.1464




Figure 4.3: S&P500 Returns, Implied Volatilities and VIX Index Values, March to June 2005

The maturity and moneyness patterns of implied volatilities are consistent with our discussion in section 2.1 and are similar to those found in other studies. In particular, the implied volatilities of in-the-money options and long-term options are generally higher than those of short-term and out-of-the-money contracts, and as moneyness increases, implied volatilities first go down and then slightly increase again.

Figure 4.3 plots the S&P 500 daily log returns, the daily VIX index, and the average implied volatilities of nearest-to-money options for the option sample period. Notice that although this sample covers a relatively short time interval, there is ample variation in volatilities to test the flexibility of the various models.

Chapter 5

Parameter Estimates and Tests of Fit

This section presents the empirical findings. First, the results from maximum likelihood estimation of the parameters under the physical probability measure are presented. Then, the relative performance of alternative models is evaluated using several tests of nested and non-nested models. Next, the risk-neutral parameters of the models are estimated using option data, and the in-sample fits of the models for each option sample date are compared. Finally, we look at how well models can forecast option prices one to five days ahead. All computations were done using codes in Fortran 77.

5.1 Maximum Likelihood Estimation from Index Returns

5.1.1 The Hobson-Rogers Model

Recall that the Hobson-Rogers model is described by the following process for the discounted stock price $\tilde{S}_t = \ln(e^{-\alpha t} S_t)$:

$$d\tilde{S}_t = \mu(\Delta_t)dt + \sigma(\Delta_t)dW_t,$$

where $\{W_t\}_{t \geq 0}$ is a standard Brownian motion and the offset function Δ_t is defined as

$$\Delta_t = \int_0^\infty \lambda e^{-\lambda u} (\tilde{S}_t - \tilde{S}_{t-u}) du.$$

Unlike the other models, the Hobson-Rogers model is formulated in terms of the discounted log-price process, \tilde{S}_t , rather than in terms of index price S_t itself. Accordingly, a process for discount factor α_t must be specified before estimation can proceed. Hobson and Rogers (1998) set α_t equal to a constant short rate r , and we follow that convention here. A seemingly plausible alternative would be to represent α by the long-run steady-state growth rate of $\{S_t\}$; however, when taken to the stock price data, this parameter cannot be identified separately from the constant in the drift in $\{S_t\}$. Thus, it is necessary to fix α exogenously.

The parameters to be estimated are $\beta_1, \beta_2, \beta_3, \eta_1, \eta_2, \eta_3$, and λ . In addition, the initial value of the offset function Δ_0 is estimated as one more parameter of the model.¹ The likelihood function is highly nonlinear and nonmonotone, and analytical

¹Alternatively, this variable can be computed from past stock price data given the value of the parameter λ . Experiments with this setup indicated that the difference in results is negligible, but this procedure is more involved computationally.

expressions for the maximum likelihood estimators are not available. The estimation strategy is first to “concentrate” the likelihood, expressing $\beta_1, \beta_2, \beta_3$, and η_1 in terms of η_2, η_3, λ and Δ_0 and the data using their first-order conditions, then to apply a numerical algorithm (a simplex method with simulated annealing) in an effort to find the global maximum. This approach reduces the dimension of the numerical search problem from eight to four parameters. Several different sets of starting values for η_2, η_3, λ and Δ_0 were used to help locate the global maximum.

Table 5.1 presents the results for three alternative specifications of the drift process. Entries below the parameter estimates are standard errors obtained via the BHHH algorithm and the negative inverse of the Hessian matrix, respectively.² Estimates of η_1, η_2, η_3 , and λ are statistically significant for all specifications of the drift process, while the parameters of the mean— β_1, β_2 and β_3 —are all individually insignificant. The likelihood ratio test of the joint significance of the β 's ($H_0 : \beta_2 = \beta_3 = 0$) fails to reject the null in both samples at the 10% level. Therefore, unlike the volatility of returns, it appears that expected returns do not exhibit predictable variation. The negative value of $\hat{\eta}_2$ indicates that volatility tends to be high after a period of low returns, the offset then being negative, and low after a period of high returns. The positive value of $\hat{\eta}_3$ indicates that large current price shocks of either sign predict high future volatility.

The ex-ante daily expected volatilities implied by drift specification 1 ($\mu_t = \beta_1$) are depicted in Figure 5.1, and their descriptive statistics are presented in Table 5.2.³ The average volatility of returns is 10.2%, and the volatility varies between 7.0% and 17.1% respectively. Notice that the estimated average volatility is very close to the standard deviation of sample returns reported in Table 4.1, and the estimated

²In computing these matrices both first- and second-order conditions were derived analytically.

³Specifications 2 and 3 produce virtually identical volatility series.

Table 5.1: **The Discrete-Time Hobson-Rogers Model: ML Estimates of Physical Parameters Using Returns Data**

	I	II	III
	$\mu_t = \beta_1$	$\mu_t = \beta_1 + \beta_2\sigma_t + \beta_3\sigma_t^2$	$\mu_t = \beta_1 + \beta_2\Delta_t + \beta_3\Delta_t^2$
β_1	0.025	1.309	0.035
	0.057	4.611	0.078
	0.056	0.046	0.080
β_2		-482.2	-7.257
		1,133	13.87
		1,307	13.81
β_3		42,865	129.4
		66,122	155.4
		70,127	154.2
η_1	0.105	0.104	0.104
	0.004	0.003	0.004
	0.003	0.003	0.003
η_2	-23.47	-22.21	-22.61
	4.804	2.681	4.781
	4.709	3.792	4.850
η_3	260.2	243.8	230.6
	145.3	11.06	146.0
	140.2	33.74	140.4
λ	9.498	9.082	9.597
	2.868	2.377	2.844
	2.908	1.199	3.063
Δ_0	-0.031	-0.037	-0.032
	0.023	0.019	0.023
	0.019	0.009	0.019
$-\log L$	2,611	2,610	2,611

Table 5.2: **The Discrete-Time Hobson-Rogers Model: Summary Statistics of Estimated Daily Volatilities (in annual terms)**

mean	0.102
median	0.098
standard deviation	0.021
min	0.070
max	0.171
skewness	0.794
kurtosis	3.021
1st quartile	0.086
3rd quartile	0.114

volatility of volatility is comparable to that of the VIX index.

Estimates of the decay parameter λ are particularly interesting. So far as we know, the only previous estimate from time series data is in the paper by Figa-Talamanca and Guerra (2006), who derive it from an autoregression of squared log-returns. Using 750 observations on S&P 500 index (up to February 2002), these authors obtain the parameter value of 12.6, which is argued to be a reasonable estimate then used to price options. In their original paper, Hobson and Rogers take the values of 1 and 5 for this parameter merely to illustrate different volatility patterns. Given such a wide range of the values for λ used in the literature, the formal estimation of this parameter on the basis of stock price data is therefore of particular interest.

As Table 5.1 shows, all our estimates of λ fall in the relatively narrow range from 9.08 to 9.60 and are highly significant. Decay rates of this magnitude indicate that the weights of the index returns from two months and four months earlier are only about 0.20 and 0.05, respectively, while the behavior of prices more than six months in the past has negligible effect on future volatility.




Figure 5.1: Estimated Daily Volatilities in the Hobson-Rogers Model

5.1.2 The Black-Scholes Model

The maximum likelihood estimates for the GBM model, in which the stock price process is given by $dS_t/S_t = \mu dt + \sigma dW_t$ (where $\{W_t\}_{t \geq 0}$ is a standard Brownian motion), are presented in *Table 5.3*. Entries below the parameter estimates are standard errors obtained via the BHHH algorithm and the negative inverse of the Hessian matrix, respectively. The drift parameter μ is insignificant, while the volatility parameter σ is strongly significant, which parallels the results for the Hobson-Rogers model.

Table 5.3: **The Black-Scholes Model: ML Estimates of Physical Parameters Using Returns Data**

μ	0.082
	0.060
	0.060
σ	0.104
	0.003
	0.003
$-\log L$	2,644

5.1.3 The CEV Model

Table 5.4 contains the estimates of the discrete-time CEV model. Recall that the dynamics of the return is assumed to be

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu_t - \sigma_t^2/2 + \sigma_t Z_t,$$

where $Z_t \sim N(0, 1)$ and $\sigma_t = \sigma S_{t-1}^{\gamma-1}$, and the drift process is specified either as a constant or as a quadratic function of volatility. Entries below the parameter estimates are the BHHH and inverse Hessian standard errors. The additional parameters of the drift process in the unrestricted version of the model are both individually and jointly insignificant. To help interpret volatility parameter estimates, the initial value of volatility, $\sigma_t|_{t=0} = \sigma/S_0^{\gamma-1}$, is also reported. Notice that this is the same for both drift specifications despite the differences in the estimates of σ and γ .⁴

To assess the magnitude of potential biases due to discretization of a continuous-time model, the parameter estimates from the continuous-time version of the CEV

⁴At an earlier stage of this research, a sample of S&P500 prices over the period from January 1, 1997 till January 1, 2007 was used, producing $\hat{\gamma} = 0.34$ with standard error of 0.05, suggesting that the estimates are very sensitive to the choice of the sample period. Other studies find that for individual stocks this coefficient ranges from -2 to 2.

Table 5.4: **The Discrete-Time CEV Model: ML Estimates of Physical Parameters Using Returns Data**

	I	II
	$\mu_t = \beta_1$	$\mu_t = \beta_1 + \beta_2\sigma_t + \beta_3\sigma_t^2$
β_1	0.044	3.127
	0.059	2.513
	0.058	3.495
β_2		-1,020
		868.747
		923.514
β_3		82,958
		212,312
		231,563
σ	83,976	56,021
	27,594	15,865
	32,886	19,942
γ	-0.917	-0.860
	0.387	0.233
	0.396	0.298
$-\log L$	2,632	2,631
σ_0	0.122	0.122

Table 5.5: **The Continuous-Time CEV Model: ML Estimates of Physical Parameters Using Returns Data**

β_1	0.076
	0.059
	0.058
σ	73,229
	24,960
	29,880
γ	-0.898
	0.393
	0.404
$-\log L$	2,628
σ_0	0.122

model with constant drift are reported in Table 5.5.

Notice that all estimates are very similar, in particular the elasticity of variance parameter γ . The estimates of the initial volatility are identical at the chosen precision level. The results suggest that the discrete-time approximation to a continuous time model is a reasonable alternative, at least in the context of the CEV model.

5.1.4 The Heston Model

Table 5.6 presents the estimates for the discrete-time Heston model. There are two versions, corresponding to whether the volatility is regarded as unobservable or as represented by VIX. Standard errors of estimates computed with BHHH and inverse Hessian approaches respectively are reported below the parameter values .

Several observations are noteworthy. First, the estimates under different assumptions about initial volatility are quite similar, suggesting that both estimation techniques are valid. Specifically, this similarity supports the use of VIX, thereby sim-

Table 5.6: **The Discrete-Time Heston Model: ML Estimates of Physical Parameters Using Returns Data**

	Unobserved	VIX
μ	0.045	0.054
	0.042	0.080
	0.043	0.012
ξ	7.129	10.408
	2.456	3.015
	2.746	3.872
σ_∞^2	0.011	0.013
	0.002	0.003
	0.001	0.001
γ	0.298	0.226
	0.047	0.040
	0.013	0.037
ρ	-0.835	-0.784
	0.057	0.017
	0.032	0.013
σ_0	0.144	
	0.022	
	0.023	
$-\log L$	2,611	2,624

plifying estimation.⁵ Second, the long-term value of volatility, given by $\sqrt{\sigma_\infty^2}$, in the case of unobservable volatility is equal to 0.105 matching the standard deviation of index returns reported in Table 4.1. Next, the correlation ρ between the innovations to stock price and volatility is strongly negative and quite close for both versions of the model. However, its estimate is somewhat lower in absolute value when the volatility is represented by VIX, and, as we will see later, is closer to the estimate implied from options data. This corresponds to the usual finding that the informational content of the VIX index, which itself is computed from options data, differs from that inferred from historical returns.

5.1.5 The Variance Gamma Model

Finally, *Table 5.7* reports the estimates for the VG model, in which stock returns are governed by the process

$$R_t \equiv \ln(S_t/S_{t-1}) = \mu + \gamma\nu\Gamma_t + \sigma\sqrt{\nu\Gamma_t}Z_t,$$

where $Z_t \sim N(0, 1)$ and $\Gamma_t \sim \text{Gamma}(1)$. Again, entries below the parameter estimates are standard errors obtained via the BHHH algorithm and the negative inverse of the Hessian matrix, respectively.

Again, the drift parameter is insignificant, while σ is highly significant. The hypothesis that the distribution is symmetric ($H_0 : \gamma = 0$) cannot be rejected. This supports the conclusion by Madan, Carr and Chang (1998) that S&P 500 returns are well described by a symmetric VG process. On the other hand, the kurtosis parameter ν is highly significant and yields a daily excess kurtosis coefficient of $3 \times \nu \times 252$, or ap-

⁵The estimates are also comparable to those reported by Ait-Sahalia and Kimmel (2007) for S&P500 returns for a period from January 2, 1990 till September 30, 2003.

Table 5.7: **The VG Model: ML Estimates of Physical Parameters Using Returns Data**

μ	0.094
	0.067
	0.059
σ	0.109
	0.005
	0.004
γ	-0.231
	0.156
	0.173
ν	0.003
	0.001
	0.001
$-\log L$	2,641

proximately 2 for both samples. This is much higher than the sample excess kurtosis of 0.256 reported in Table 4.1.

5.1.6 Comparing the Models

Having estimated the models, we are now interested in evaluating their relative performances. For this purpose, we first compare each model to GBM using a nested model test. Then, two non-nested tests by Vuong (1989) and Clarke (2003) are employed to compare all models pairwise. Finally, the hybrid CEV-HR model is estimated to judge the relative contributions of these two alternatives.

Since GBM is nested within the other four models, it can be compared using the likelihood-ratio test. Table 5.8 presents the log-likelihood functions for all models and the corresponding likelihood-ratio (LR) statistics versus GBM, computed as $-2(\ln L_{restricted} - \ln L_{unrestricted})$. Comparisons for CEV and HR are for the restricted versions of the drift process. The numbers beneath the LR statistics are p-values of

Table 5.8: Likelihood Ratio Tests With Respect to the GBM Model

	GBM	CEV	Heston	VG	HR
Log-likelihoods	-2,644	-2,632	-2,611	-2,641	-2,610
LR statistics	–	22.87	66.08	6.23	66.90
p-values	–	1.73e-6	1.52e-13	4.44e-02	2.97e-15

one-sided test that CEV/Heston/VG/HR model is better than GBM

Clearly, the LR tests soundly reject GBM in favor of any of these models. The likelihoods for the Hobson-Rogers and Heston models are very similar, and both are much higher than the others.

Next, we will compare the CEV, HR, Heston and VG models relatively to each other. We start with the pair of CEV and HR models. To compare them, CEV-HR hybrid model, where $dS_t/S_t = \mu_t dt + \sigma_t dW_t$ and $\sigma_t^2(\Delta_t, S_t) = \eta_1^2 (1 + \eta_2 \Delta_t + \eta_3 \Delta_t^2) S_t^{2(\gamma-1)}$, is estimated. Notice that when $\eta_2 = \eta_3 = 0$, the model reduces to the CEV model, and when $\gamma = 1$, the model reduces to the standard HR specification. The assumption of the constant drift can be easily relaxed to incorporate the dependence of the expected return process on offset or volatility. However, since a more general drift process was found to be adding little to the explanatory power of the models, the assumption of constant expected return is not crucial.

The estimation results for this hybrid model are presented in Table 5.9. Entries below the parameter estimates are standard errors obtained via 1) the BHHH algorithm and 2) the negative inverse of the Hessian matrix.

The likelihood-ratio test easily rejects the CEV model in favor of the hybrid model, while the Hobson-Rogers model cannot be rejected (the p -value of the hypothesis $\eta_2 = \eta_3 = 0$ is 0.178). In addition, we cannot reject the hypothesis of $\gamma = 1$.

Table 5.9: **The Discrete-Time Hybrid HR-CEV Model: ML Estimates of Physical Parameters Using Returns Data**

μ	0.040
	0.056
	0.056
η_1	31.13
	7.560
	10.01
η_2	-22.70
	5.401
	4.904
η_3	233.5
	159.4
	142.6
λ	11.92
	4.187
	3.932
γ	0.201
	0.464
	0.485
F_0	-0.017
	0.026
	0.022
$\log L$	-2,610
σ_0	0.168

Table 5.10: Non-Nested Tests

Panel A: Vuong's (1989) test				
Model F	Model G			
	CEV	HR	Heston	VG
GBM	1.758	2.251	2.274	0.736
	0.079	0.024	0.023	0.462
CEV		1.254	1.399	-1.772
		0.210	0.162	0.076
HR			0.475	-2.703
			0.635	0.007
Heston				-3.109
				0.002
Panel B: Clarke's (2003) "sign test"				
Model F	Model G			
	CEV	HR	Heston	VG
GBM	0.528	0.537	0.460	0.499
	0.118	0.038	0.032	0.971
CEV		0.501	0.448	0.454
		0.913	0.005	0.012
HR			0.467	0.421
			0.075	0.000
Heston				0.422
				0.000

Therefore, the results suggest that the CEV component of the hybrid model does not makes a significant contribution to its explanatory power.

Next, we compare all models with each other using formal non-nested tests by Vuong (1989) and Clarke (2003). The results are reported in Table 5.10. A p -value of each two-sided test (that Model G and Model F explain the data equally good) is given below its statistics.

For completeness, all models, including the GBM, have been compared pairwise. As expected, both these tests are less powerful than the LR test in their ability to reject the null that CEV/HR/Heston/VG models are "equal" to GBM. However, both tests imply the same ranking of the modes as the likelihood ratio statistics. Several

observations deserve mentioning. While the HR models is still clearly favored over the GBM and VG models by both tests, it cannot be differentiated from the CEV model. In addition, the HR and Heston models are indistinguishable according to Vuong's test, but the Heston model is actually doing worse in terms of frequency of observations with higher likelihoods. Interestingly, the sign test finds that the Heston model yields to the GBM and CEV models as well. The performance of the VG model is found to be equal to that of the GBM, but significantly worse than that of the other three models.

The apparent ambiguity of the results with respect to the HR and Heston models can be explained by looking at the graph of the differences of their individual log-likelihood functions (*Figure 5.2*). The graph shows that the differences in estimated probabilities of individual data points are fairly evenly distributed around zero, with the exception of several outliers. These outliers correspond to extreme market observations with particularly small likelihoods implied by both models, which are however still much higher in the Heston model than in the HR model. These outliers are driving the Vuong's statistics, as well as the likelihood-ratio statistics, but Clarke's test by construction is relatively immune to them. Therefore, even though the Heston model seems to be better able to handle large negative market shocks, the HR model performs at least as well as the Heston model under less extreme market circumstances.

5.2 Least Squares Estimation from Options Data

This section presents non-linear least squares estimation results obtained using the sample of options data. Two alternative loss functions, mean-squared error of dollar option prices (MSEP) and mean-squared error of implied volatilities (MSEIV)




Figure 5.2: Differences of the Individual HR and Heston Log-Likelihoods

Table 5.11: **Black-Scholes and CEV Models: NLLS Estimates of Risk-Neutral Parameters Using Options Data**

	Parameters		
	GBM	CEV	
Loss Function:	σ	σ	γ
MSEP	0.1386	0.1384	-3.2442
	0.0075	0.0075	0.4780
	0.0063	0.0054	0.3017
MSEIV	0.1451	0.1375	-3.6852
	0.0112	0.0082	0.0903
	0.0092	0.0078	0.1587
MSEP: simulations		0.1422	-1.7271
		0.0080	0.2040
		0.0061	0.1962

are assumed. For GBM, CEV, Heston and VG models, closed-form expressions for option prices are used to evaluate loss functions. The Hobson-Rogers model, for which no such formula is available, is estimated via simulations. To evaluate the effect of discretization bias and to make meaningful comparisons, the discrete-time versions of the CEV, Heston and VG models are also estimated via simulations under the MSEP criterion.

5.2.1 Risk-Neutral Parameter Estimates

Tables 5.11, 5.12, 5.13 and 5.14 report the daily averages and standard errors of parameters estimates under different loss functions for the HR, GBM, CEV, Heston and VG models. Two approximations of standard errors are provided beneath each parameter estimate: a sample standard deviation of the estimate and the average of sample asymptotic standard errors, taken from the asymptotic covariance matrices of the non-linear least-squares regression coefficients.

Several observations are in order. First, there is some variation in the estimates

Table 5.12: **Heston Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data**

Loss Function:	Parameters				
	σ_0^2	α	β	γ	ρ
MSEP	0.0164	0.1058	3.3792	0.6556	-0.6226
	0.0035	0.0138	0.6511	0.0748	0.0454
	0.0003	0.0096	0.4193	0.0418	0.0157
MSEIV	0.0168	0.1253	4.1547	0.8242	-0.5975
	0.0039	0.0181	0.8778	0.1363	0.0469
	0.0005	0.0087	0.3965	0.0337	0.0196
MSEP: simulations	0.0165	0.1185	4.5641	0.8017	-0.6218
	0.0039	0.0218	1.5650	0.1624	0.0493
	0.0004	0.0122	0.6748	0.0697	0.0177

Table 5.13: **Variance-Gamma Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data**

Loss Function:	Parameters		
	σ	γ	ν
MSEP	0.1293	-0.1495	0.3411
	0.0049	0.0295	0.0953
	0.0031	0.0123	0.0149
MSEIV	0.1234	-0.1818	0.2205
	0.0049	0.0295	0.0953
	0.0033	0.0177	0.0267
MSEP: simulations	0.1315	-0.1441	0.3345
	0.0082	0.0205	0.0456
	0.0032	0.0084	0.0014

Table 5.14: **Hobson-Rogers Model: NLLS Estimates of Risk-Neutral Parameters Using Options Data**

	Parameters				
Loss Function:	η_1	η_2	η_3	λ	F_0
MSEP: simulations	0.1359	-26.404	361.999	8.254	0.012
	0.0078	4.9446	125.51	1.7709	0.0074
	0.0018	1.3890	50.656	0.7072	0.0015
MSEIV: simulations	0.1332	-28.628	451.193	9.997	0.0100
	0.0055	4.6404	83.421	1.5096	0.0079
	0.0020	1.8333	66.555	1.3187	0.0017

of instantaneous volatility across models: in the Heston model, average σ_0 is 0.128 and 0.130 for the MSEP and MSEIV loss functions respectively, the average initial volatility in the HR model is 0.116 for both loss functions, and the average volatilities in the BS, CEV and VG models are given by parameters σ in Tables 5.11 and 5.13. Interestingly, the options-implied volatility estimate from the HR model is the closest to the ex-post realized volatility of 0.111 observed in the spot market, computed as the standard deviation of index returns over the four-month period centered around the option data interval.

The inconsistency between option-implied volatility and ex-post spot market volatility, which is particularly noticeable in the settings of the BS and CEV models (Table 5.11), is indicative of model misspecification.⁶ An explanation for this volatility bias is that in the case when the data-generating process involves uncertain volatility, then option prices contain additional risk-premium component, and a model that does not appropriately account for it translates this premium into higher implicit volatilities.

Second, even though the parameter estimates vary from one sample date to another, as can be seen from graphs presented in Appendix C, their averages are quite

⁶Similar findings are reported by Jackwerth and Rubinstein (1996) for the BS model and Pan (2000) for the BS and Heston models.

robust to the choice of the loss function and estimation procedure, and mostly fall into \pm two standard deviations intervals. One notable exception is the CEV model (Table 5.11), in which the elasticity of volatility parameter γ is considerably larger in absolute value when estimated using a continuous-time version of the model than in a discrete-time setup. This observation is particularly surprising in view of the findings presented in Tables 5.4 and 5.5 that suggest that the two versions of the CEV model produce quite similar results on stock data.

Third, comparing the risk-neutral parameters of the models to their physical counterparts provides a number of interesting insights. However, even though the conclusions from this exercise are potentially informative of model misspecification, they should be taken with caution since the stock index sample period extends far beyond the options sample dates.

In the context of the CEV model, the difference between the estimates of volatility elasticity γ from the options data (Table 5.11) and those from the stock index series (Table 5.4 and 5.5) is remarkable. As Epps (2007, p. 378) demonstrates, lower elasticity is associated with steeper implied volatility smile, which suggests that the structure of implied volatilities embedded in options data is considerably steeper than that in stock index series.

To formally relate the risk-neutral and physical parameters of the Heston model, the general-equilibrium models of Bakshi and Chen (1997) and Bates (1996a) can be employed. The models assume that volatility premium is proportional to the volatility risk factor, implying that the physical and risk-neutral stock return and volatility dynamics share the same stochastic structure. Moreover, the parameters γ and ρ should be the same under either probability measure.⁷ Therefore, under the

⁷In addition, Bates (1991) argues that in a representative-agent model with the risk-aversion coefficient bounded within a reasonable range, the parameters of the two distributions will not differ significantly.

hypothesis that this model is correctly specified, the average value of option-implied coefficients γ and ρ must be equal to their time-series counterparts.

The estimation results presented in Tables 5.6 and 5.12 do not support this hypothesis. In particular, the coefficient of correlation between daily returns and volatility changes of the S&P 500 index, which determines the magnitude of implied volatility skew, for the discrete-time model is equal to -0.62 when estimated from options data, and is significantly higher than the MLE estimate of -0.84 obtained from the index price series.⁸

The estimates of the volatility of variance parameter γ , which controls the convexity of volatility smile, are also significantly different : 0.802 vs. 0.298 from the options and stock data respectively. Therefore, the variance process implicit in option prices is much too volatile to be consistent with stock price series.⁹

The results for Hobson-Rogers model, presented in Tables 5.1 and 5.14, seem to be more consistent under two alternative probability measures. In particular, while the parameters of the volatility process η_1 and η_3 implied by the options data are significantly higher than those obtained from the stock series, the estimate of η_2 is indistinguishable in these two data sets. As it has already been mentioned, despite the difference in volatility process parameters the average instantaneous volatility extracted from options prices is still very close to the ex-ante realized stock volatility over the options sample period.

Another interesting result is that the estimate of the decay factor λ , at which past information is discounted in the offset function, is also statistically the same in stock returns and options samples. The finding that both data sets imply the same memory

⁸Interestingly, the correlation of S&P 500 returns and VIX series of -0.205 provided in Table 4.1 is quite different from both of these estimates. The correlation of these two series over the option sample period is -0.338 .

⁹Bakshi et al (1997) arrive to a similar conclusion using S&P 500 index options data for the period from June 1, 1988 till May 31, 1991.

Table 5.15: **In-Sample Fit to Options Data**

Model	MSEP	RMSEP	NP	MSEIV	RMSEIV	NIV
$BS_{formula}$	14.1734 (3.4601)	3.7381 (0.4528)	0.00 –	1.24E-03 (3.73E-04)	3.48E-02 (5.20E-03)	0.00 –
$CEV_{simulations}$	4.5674 (1.3414)	2.1135 (0.3207)	0.00 –	– –	– –	– –
$CEV_{formula}$	2.6091 (0.7633)	1.5971 (0.2445)	0.00 –	1.73E-04 (7.83E-05)	1.29E-02 (2.71E-03)	0.00 –
$Heston_{simulations}$	0.2071 (0.0780)	0.4468 (0.0876)	0.16 –	– –	– –	– –
$Heston_{formula}$	0.1972 (0.0694)	0.4373 (0.0785)	0.30 –	3.36E-05 (3.85E-05)	5.33E-03 (2.31E-03)	0.09 –
$VG_{simulations}$	1.6636 (0.4709)	1.2769 (0.1845)	0.00 –	– –	– –	– –
$VG_{formula}$	1.7648 (0.4505)	1.3174 (0.1728)	0.00 –	1.11E-04 (5.33E-05)	1.03E-02 (2.24E-03)	0.00 –
$HR_{simulation}$	0.1644 (0.0725)	0.3963 (0.0867)	– –	2.33E-05 (3.68E-05)	4.18E-03 (2.44E-03)	– –

process is particularly encouraging.

In summary, while all models¹⁰ fail to produce systematically consistent parameter estimates implied by the stock price series and options panel data, there is some evidence in support of the Hobson-Rogers model but not the others.

5.2.2 In-Sample Option Pricing

Table 5.15 describes the in-sample fit of alternative option pricing models to the options data. The numbers reported in the MSEP (MSEIV) column are the average mean squared errors of dollar option prices (mean squared error of implied volatilities) across the option sample and their sample standard deviations(in parentheses). The RMSEP (RMSEIV) column gives the respective average root mean squared errors

¹⁰Unfortunately, so far no general equilibrium model involving the VG process has been suggested in the literature. Therefore, this process was excluded from the discussion of consistency between physical and risk-neutral parameter estimates.

and their standard deviations (in parentheses). The NP (NIV) column contains the percentage of data points for which the model has lower MSEP (MSEIV) than the Hobson-Rogers model.

Under both loss functions, the HR and Heston models are by far performing better in terms of average mean squared errors than the three other models, which is consistent with the evidence from index returns discussed earlier. The GBM model clearly demonstrates the worst fit to the data.

As to the relative performance of the HR and Heston models, when estimated with the MSE of dollar option prices criterion the HR model is producing lower MSEs for approximately 84% of data points if compared to the discrete-time version of Heston model, and for 70% of data points if compared to the continuous-time version of Heston model. With the MSE of implied volatilities loss function, the Heston model is doing better than the HR model only in approximately 9% of cases. None of the other models is able to outperform the HR model on at least one occasion.

One more interesting result concerns the relative fit of the VG and CEV models. Recall that the fit of the VG model to stock returns data was inferior to that of the CEV model (see Table 5.10). However, the estimation results from the options data set strongly favor the VG model over the CEV model. Such reversal of relative performance is an indication of model misspecification.

5.2.3 Implied Volatility Patterns

Following Bakshi, Cao and Chen (1997), Chernov and Ghysels (2000) and other authors, we assess in-sample performance of all models from a different angle by comparing their implied-volatility patterns across both moneyness and maturity. This approach allows analyzing pricing errors without putting undue weight on expensive

options (deep-in-the-money or longer-term contracts).

For each model, in-sample implied volatilities are backed out from model-predicted option prices using the Black-Scholes option pricing formula. The basic procedure is as follows. First, at date t compute model-predicted option prices using date t parameter estimates, spot index value and interest rates. Next, for every contract find a value of parameter σ that equates an estimated option price with Black-Scholes price computed from expression (2.5), in which σ is the only unknown input.

The obtained implied volatilities are grouped into same maturity and moneyness categories as in Tables 4.2-4.4. The average model-predicted implied volatilities are reported in Tables 5.16, 5.17, 5.18 and 5.19. To interpret the results, compare these model-predicted implied volatilities to market volatilities described in Table 4.4.

The structure of implied volatilities produced by both the GBM¹¹ and the CEV models is much flatter than that implied by the data, in particular for shorter-term contracts. The short-term volatility smiles in the VG model are more pronounced than in the CEV model, but while the CEV model underprices (overprices) in-the-money (out-of-the-money) puts, the pricing errors of the VG model are reversed. The fit of both models seems to improve with options maturity.

The HR and Heston models generate virtually indistinguishable implied volatility curves, that closely follow the market. This result is not surprising given that the estimated mean squared errors of these two models are very similar and small.

To further illustrate the results, we plot market and model implied volatilities for options traded on an average-volatility day (March 29, 2005) in Appendix D.

¹¹The differences in GBM implied volatilities in Tables 5.16 and 5.18 across cells is due to averaging.

Table 5.16: Average In-Sample Implied Volatilities in the GBM, CEV and VG Models: MSEP Loss Function

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
Panel A: GBM Model					
0.80-0.94	0.1395	0.1274	0.1307	0.1387	0.1359
0.94-0.97	0.1374	0.1245	0.1362	0.1367	0.1346
0.97-1.00	0.1388	0.1303	0.1357	0.1391	0.1371
1.00-1.03	0.1375	0.1276	0.1332	0.1396	0.1356
1.03-1.06	0.1404	0.1248	0.1359	0.1398	0.1369
1.06-1.33	0.1436	0.1287	0.1395	0.1395	0.1394
Total	0.1394	0.1272	0.1344	0.1390	0.1365
Panel B: CEV Model					
0.80-0.94	0.1795	0.1647	0.1729	0.1888	0.1784
0.94-0.97	0.1518	0.1374	0.1503	0.1521	0.1489
0.97-1.00	0.1441	0.1350	0.1413	0.1457	0.1426
1.00-1.03	0.1342	0.1249	0.1310	0.1383	0.1329
1.03-1.06	0.1289	0.1145	0.1255	0.1304	0.1260
1.06-1.33	0.1318	0.1105	0.1158	0.1139	0.1196
Total	0.1488	0.1359	0.1450	0.1499	0.1463
Panel C: VG Model					
0.80-0.94	0.2228	0.1824	0.1717	0.1746	0.1963
0.94-0.97	0.1692	0.1442	0.1525	0.1525	0.1591
0.97-1.00	0.1377	0.1325	0.1425	0.1487	0.1401
1.00-1.03	0.1044	0.1134	0.1309	0.1438	0.1163
1.03-1.06	0.1100	0.1009	0.1232	0.1379	0.1135
1.06-1.33	0.1346	0.1111	0.1194	0.1277	0.1264
Total	0.1530	0.1372	0.1453	0.1506	0.1487

Table 5.17: **Average In-Sample Implied Volatilities in the Heston and HR Models: MSEF Loss Function**

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
Panel A: Heston Model					
0.80-0.94	0.1966	0.1739	0.1768	0.1870	0.1874
0.94-0.97	0.1567	0.1403	0.1512	0.1585	0.1533
0.97-1.00	0.1363	0.1286	0.1397	0.1499	0.1387
1.00-1.03	0.1186	0.1137	0.1276	0.1429	0.1232
1.03-1.06	0.1099	0.1014	0.1202	0.1354	0.1129
1.06-1.33	0.1226	0.1051	0.1158	0.1238	0.1196
Total	0.1456	0.1332	0.1451	0.1539	0.1453
Panel B: Hobson-Rogers Model					
0.80-0.94	0.1994	0.1784	0.1801	0.1879	0.1901
0.94-0.97	0.1537	0.1396	0.1512	0.1579	0.1515
0.97-1.00	0.1341	0.1285	0.1401	0.1496	0.1376
1.00-1.03	0.1188	0.1149	0.1288	0.1431	0.1236
1.03-1.06	0.1092	0.1028	0.1218	0.1358	0.1130
1.06-1.33	0.1131	0.1036	0.1141	0.1226	0.1156
Total	0.1446	0.1346	0.1463	0.1538	0.1453

Table 5.18: **Average In-Sample Implied Volatilities in the GBM, CEV and VG Models: MSEIV Loss Function**

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
Panel A: GBM Model					
0.80-0.94	0.1469	0.1330	0.1365	0.1454	0.1426
0.94-0.97	0.1441	0.1293	0.1430	0.1428	0.1408
0.97-1.00	0.1453	0.1352	0.1413	0.1454	0.1431
1.00-1.03	0.1441	0.1326	0.1392	0.1462	0.1418
1.03-1.06	0.1474	0.1294	0.1418	0.1474	0.1434
1.06-1.33	0.1517	0.1335	0.1460	0.1469	0.1465
Total	0.1464	0.1323	0.1405	0.1458	0.1431
Panel B: CEV Model					
0.80-0.94	0.1751	0.1676	0.1782	0.1944	0.1791
0.94-0.97	0.1517	0.1368	0.1503	0.1523	0.1488
0.97-1.00	0.1428	0.1333	0.1398	0.1444	0.1413
1.00-1.03	0.1320	0.1226	0.1288	0.1363	0.1307
1.03-1.06	0.1255	0.1111	0.1215	0.1269	0.1226
1.06-1.33	0.1283	0.1068	0.1101	0.1086	0.1150
Total	0.1460	0.1350	0.1449	0.1497	0.1448
Panel C: VG Model					
0.80-0.94	0.2127	0.1736	0.1649	0.1670	0.1875
0.94-0.97	0.1656	0.1403	0.1493	0.1482	0.1553
0.97-1.00	0.1396	0.1322	0.1408	0.1458	0.1400
1.00-1.03	0.1101	0.1156	0.1312	0.1422	0.1194
1.03-1.06	0.1087	0.1020	0.1250	0.1374	0.1131
1.06-1.33	0.1308	0.1073	0.1199	0.1280	0.1248
Total	0.1505	0.1343	0.1428	0.1472	0.1459

Table 5.19: **Average In-Sample Implied Volatilities in the Heston and HR Models: MSEIV Loss Function**

Moneyness (K/S)	Maturity (days)				Total
	21-44	45-89	90-134	135-252	
Panel A: Heston Model					
0.80-0.94	0.2029	0.1778	0.1792	0.1867	0.1911
0.94-0.97	0.1596	0.1405	0.1513	0.1580	0.1546
0.97-1.00	0.1359	0.1274	0.1388	0.1509	0.1385
1.00-1.03	0.1167	0.1119	0.1262	0.1440	0.1219
1.03-1.06	0.1104	0.1003	0.1192	0.1363	0.1130
1.06-1.33	0.1269	0.1059	0.1166	0.1243	0.1214
Total	0.1476	0.1337	0.1457	0.1542	0.1465
Panel B: Hobson-Rogers Model					
0.80-0.94	0.2030	0.1795	0.1802	0.1867	0.1915
0.94-0.97	0.1566	0.1407	0.1517	0.1580	0.1532
0.97-1.00	0.1360	0.1293	0.1408	0.1501	0.1389
1.00-1.03	0.1198	0.1156	0.1296	0.1439	0.1245
1.03-1.06	0.1104	0.1035	0.1235	0.1376	0.1142
1.06-1.33	0.1202	0.1042	0.1172	0.1252	0.1195
Total	0.1473	0.1355	0.1474	0.1544	0.1470

5.3 Out-of-Sample Pricing Performance

The in-sample fit of the models to daily option prices becomes increasingly better as we move from the GBM model to the CEV model, and then to the VG, Heston and HR models. Since the number of parameters in each consecutive model is increasing¹², the issue of overfitting might be raised. To address this issue, the out-of-sample pricing performance of the models is examined next. If extra parameters cause overfitting but do not improve the structural performance of the model, then its forecasting ability will not necessarily be superior to that of more parsimonious models.

To price out-of-sample, the day t model parameter estimates are used as inputs to compute date $t+1, \dots, t+5$ model-based option prices. In the Heston model, the value of initial volatility $\{\sigma_{t+s}^2\}_{s=1}^5$ is computed from the estimate of σ_t^2 according to the following naive undating rule: $\sigma_{t+1}^2 = \sigma_t^2 + (\alpha - \beta\sigma_t^2) + \gamma\rho(\ln(S_{t+1}/S_t) - r + 0.5\sigma_t^2)$. This rule gives the expected conditional variance of the model under the risk-neutral measure.¹³ The value of the initial offset in the HR model is updated using expression (3.10).

Table 5.20 reports average and median mean squared errors and root mean squared errors of forecasted option prices and their sample standard deviations computed with parameter estimates obtained under the MSEP loss function. In addition, the percentage of sample dates on which a model is producing lower mean squared errors of forecasts than its closest alternative (in terms of in-sample MSEP) is reported. The MSE column reports the average (AVG), median (MED) and standard deviation (SD) of daily mean squared errors of forecasts. The RMSE column reports the average (AVG), median (MED) and standard deviation (SD) of daily root mean squared errors of forecasts. The % column reports the number of days on which the MSE of the

¹²the Heston model and the HR model have the same number of estimated parameters

¹³This rule closely resembles the conditional variance process (3.6) under the physical measure.

GBM/CEV/VG/Heston model was lower than than of the CEV/VG/Heston/HR model.

The pricing errors of all model progressively deteriorate as forecasting horizon increases, but their ranking measured by mean squared errors of prices is preserved. The GBM model shows the worst fit, and is able to outperform the CEV model in less than 5% of cases only for 4- and 5-days-ahead forecasts. Interestingly, even though the CEV model shows considerably worse in-sample fit and its average forecast MSEs are still consistently higher, it is producing lower mean squared forecast errors than the VG model in about 30% of cases over all forecasting horizons. The VG model is almost never able to outperform the Heston model.

The HR model, which delivered lower in-sample pricing errors than the Heston model (Table 5.15), is producing better average predictions only over one day forecasting horizon, but even in this case yields to the Heston model on 49% of sample days. The relative forecasting power of the HR model over all other forecasting horizons is progressively worse. Similar conclusions can be drawn from *Appendix E*, which reports the out-of-sample fit results for the MSEIV loss function.

In summary, the relative quality of forecasts produced by alternative option pricing models is closely related to their in-sample performance. However, while the HR model outperforms the Heston model in terms of in-sample fit, the Heston model demonstrates superior out-of-sample forecasting ability over longer-term horizons.

Table 5.20: **Out-of-Sample Fit to Option Prices: MSEP Loss Function**

	horizon	MSE			RMSE			%
		AVG	MED	SD	AVG	MED	SD	
GBM Model	1 day	14.970	13.884	3.648	3.842	3.726	0.463	0.000
	2 days	16.060	15.418	4.547	3.969	3.927	0.561	0.000
	3 days	16.962	16.290	5.719	4.065	4.036	0.668	0.000
	4 days	18.027	17.678	6.270	4.187	4.204	0.711	0.023
	5 days	18.823	17.933	6.559	4.277	4.235	0.737	0.047
CEV Model	1 day	4.325	3.793	2.317	2.021	1.948	0.498	0.279
	2 days	4.798	4.036	2.250	2.139	2.009	0.478	0.279
	3 days	5.469	4.498	3.343	2.254	2.121	0.632	0.326
	4 days	6.790	4.831	6.092	2.450	2.198	0.899	0.349
	5 days	7.105	4.496	6.255	2.493	2.120	0.955	0.372
VG Model	1 day	2.980	2.556	1.559	1.679	1.599	0.408	0.047
	2 days	3.873	2.885	2.879	1.874	1.698	0.610	0.023
	3 days	4.870	3.489	4.397	2.056	1.868	0.811	0.000
	4 days	5.861	3.951	5.039	2.249	1.988	0.907	0.023
	5 days	6.649	3.553	5.512	2.387	1.885	0.988	0.000
Heston Model	1 day	0.613	0.469	0.633	0.726	0.685	0.297	0.488
	2 days	0.831	0.482	0.788	0.839	0.695	0.362	0.605
	3 days	0.923	0.776	0.638	0.908	0.881	0.318	0.605
	4 days	1.036	0.853	0.688	0.964	0.923	0.331	0.628
	5 days	1.067	1.025	0.649	0.985	1.013	0.317	0.651
HR Model	1 day	0.556	0.393	0.454	0.696	0.627	0.272	–
	2 days	0.891	0.512	0.956	0.846	0.715	0.423	–
	3 days	1.151	0.926	0.971	0.997	0.962	0.400	–
	4 days	1.427	1.039	1.130	1.105	1.019	0.458	–
	5 days	1.591	1.144	1.218	1.169	1.070	0.478	–

Chapter 6

Conclusion

The motivation for this dissertation was to empirically investigate a complete-market uncertain volatility option pricing model recently proposed by Hobson and Rogers (1998), and compare its performance to that of several other well-known models of complete and incomplete nature. For this purpose, a methodology whereby this model can be estimated using stock price series and options cross-sections was developed and implemented. The Black-Scholes (BS) and the Constant Elasticity of Variance (CEV) models were taken as complete-market benchmarks, and the Heston and Variance Gamma (VG) models were used as incomplete-market alternatives.

The appeal of complete markets is that any possible payoff structure can be replicated using a (dynamic) portfolio of traded assets, which is essential for modern hedging and portfolio management techniques. However, the existing empirical evidence demonstrates that such popular complete-market models as the BS and CEV models by far yield in terms of both in-sample and out-of-sample fit to the data to more complex incomplete-market processes, including the Heston and VG models. In the light of such results, testing a new complete-market model potentially able to generate a rich structure of volatility patterns becomes particularly interesting. Even

though a number of papers tried to calibrate the Hobson-Rogers model using different approaches, the author is not aware of any attempts in the literature to rigorously estimate the model by maximum-likelihood or nonlinear least squares techniques.

The empirical strategy is to estimate the physical and risk-neutral parameters of all models separately on S&P 500 index returns and options data using both discrete-time and continuous time versions of the models whenever feasible. This approach is common to the literature and therefore permits comparing the results, in particular for the BS, CEV, Heston and VG models, with those obtained by other authors. Also, in addition to being technically more straightforward than the one employing both available data sets simultaneously, this approach produces two, rather than one, sets of parameter estimates that can then be used in model misspecification tests.

The results indicate that the Hobson-Rogers model consistently outperforms the BS, CEV and VG models on both returns and option prices data. As to the Heston model, the returns data cannot discriminate generally between the two processes. The evidence from options data supports the Hobson-Rogers model in-sample and out-of-sample with one-day-ahead forecasts, while the Heston model produces better forecasts over longer horizons.

The general conclusion is that this type of complete-market uncertain volatility models is indeed capable of capturing a number of stylized features of empirical data, such as implied volatility smiles and volatility term structure, and is a viable alternative to an acknowledged stochastic volatility process. In addition, misspecification tests based on consistency of physical and risk-neutral parameter estimates seem to offer more support to this model than to any other model being investigated in this dissertation.

The methods and results of this research can be refined in a number of ways. First, taking a longer sample of option data and bringing it closer to the returns

sample would make the conclusions much more substantiated. Second, developing an approach to pricing options in the Hobson-Rogers model that does not rely on simulations would permit to drastically reduce the required computational time. Third, since both the returns and options data contain information about the stock price process, estimating the models using these two data sets simultaneously, and Chernov and Ghysels (2000) did with the Heston model, would make the results more precise, more consistent and easier to interpret.

In the light of the empirical success of the Hobson-Rogers model, an interesting question is how comparable this model is to yet more complex incomplete market models proven to fit the data even better than the Heston model, such as, for example, stochastic volatility models with jumps. Moreover, extending the model to allow discontinuities in volatility (and stock) price process seems to be a particularly promising direction for future research.

Bibliography

Ait-Sahalia, Y. and R. Kimmel (2007). Maximum Likelihood estimation of stochastic volatility models, *Journal of Financial Economics* 83, 413-452.

Akaike, H. (1974) A new look at the statistical model identification, *IEEE Transactions on Automatic Control AC*, 19, 716-23.

Ang, J. and Peterson, P.P. (1984). The leasing puzzle, *Journal of Finance* 39, 1055-1065.

Bakshi, G., C. Cao and Z. Chen (1997). Empirical performance of alternative option pricing models, *Journal of Finance* 52, 2003-2049.

Bakshi, G., and Z. Chen (1997). An alternative valuation model for contingent claims, *Journal of Financial Economics* 44, 123-165.

Bates, D.S. (1996a). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, *Review of Financial Studies* 9 (1), 69-107.

Bates, D.S. (1996b). Testing Option Pricing Models, in *Statistical Methods in Finance (Handbook of Statistics, v. 14)*, G.S. Maddala and C. R. Rao, eds, 567-611.

Bates, D.S. (2006). Maximum Likelihood Estimation of Latent Affine Processes, *Review of Financial Studies* 19(3) 909-965.

- Beckers, S. (1980) The constant elasticity of variance model and its implications for option pricing, *The Journal of Finance* 35(3), 661-673.
- Black, F. (1976) Studies of stock price volatility changes. *Proceedings of the 1976 Meeting of the American Statistical Association, Business and Economic Statistics*, 177-181.
- Black, F. and M. Scholes (1973) The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.
- Brandt, M. and P. Santa-Clara (2002). Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* 63, 161-210.
- Carr, P. and D. Madan (1998). Option valuation using the fast Fourier transform, *Journal of Computational Finance* 2, 61-73.
- Chernov, M. and E. Ghysels (2000). Towards a unified approach to the joint estimation of objective and risk-neutral measures for the purpose of options valuation, *Journal of Financial Economics* 56, 407-458.
- Christie, A.A. (1982) The stochastic behavior of common stock variances - value, leverage and interest rate effects, *Journal of Financial Economics*, 3, 145-166.
- Christoffersen, P. and K. Jacobs (2004). Which GARCH model for option valuation? *Management Science*, 50(9), 1204-1221.
- Clarke, K. A. (2003) Nonparametric Model Discrimination in International Relations, *Journal of Conflict Resolution* 47, 72-93.
- Clarke, K.A. (2007) A simple distribution-free test for nonnested model selection, *Political Analysis* 15(3), 347-363.

- Cox, D.R. (1961) Tests of separate families of hypotheses. *Proc. 4th Berkeley Symp.* 6, 105-123.
- Cox, J. (1975). Notes of option pricing I: Constant elasticity of diffusions, *Unpublished Manuscript*, Stanford University.
- Cox J. (1996) The constant elasticity of variance option pricing model, *Journal of Portfolio Management, Special Issue*, 15-17.
- Cox, J. and S. Ross (1976). The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 145-166.
- Daal, E.A. and D.B. Madan (2005). An empirical examination of the Variance-Gamma model for foreign currency options. *Journal of Business* 78 (6), 2121-2152.
- Di Francesco, M. and A. Pascucci (2004) On the complete market model with stochastic volatility by Hobson and Rogers, *Proceedings of The Royal Society of London. Series A. Mathematical, Physical and Engineering Sciences*, 460 (2051), 3327–3338.
- Di Francesco, M , P. Foschi and A. Pascucci (2006). Analysis of an uncertain volatility model, *Journal of Applied Mathematics and Decision Sciences*, 2006, Article ID 15609, 17 pages.
- Durham, G.B. and A.R. Gallant (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes, *Journal of Business and Economic Statistics* 20(3), 297-316.
- Duffie, D., J. Pan and K.J. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions, *Econometrics* 68, 1343-1376.

- Emmanuel, D.C. and J.D. MacBeth (1982). Further results on the constant elasticity of variance call options pricing model, *Journal of Financial and Quantitative Analysis* 17, 533-554.
- Epps T.W. (2007). Pricing Derivative Securities, 2nd Edition. World Scientific.
- Eraker, B., M. Johannes and N. Polson (2003). The impact of jumps in volatility and returns. *Journal of Finance* 58(3), 1269-1300.
- Eraker, B. (2004) Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59 (3), 1367-1403.
- Foschi, P. and A. Pascucci (2005) Calibration of the Hobson-Rogers model: empirical tests, *Working Paper*.
- Figa-Talamanca, G. and M.L. Guerra (2006) Fitting prices with a complete model. *Journal of Banking and Finance* 30, 247-258.
- Friedman, M. and L. Harris (1998). A maximum likelihood approach for non-gaussian stochastic volatility models, *Journal of Business and Economic Statistics* 16(3), 282-291.
- Hallulli, V.B. and T. Vargiolu (2005). Robustness of the Hobson-Rogers model with respect to the offset function, *Seminar on Stochastic Analysis, Random Fields and Application V, Progress in Probability*, 59, 469 - 492.
- Heston, S.L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327-343.
- Hobson, D.G., and L.C.D. Rogers (1998). Complete market models with stochastic volatility, *Mathematical Finance* 8, 27-48.

- Hollander, M. and D.A. Wolfe (1999) *Nonparametric Statistical Methods*, 2nd Edition.
- Hubalek, F., J. Teichmann and R. Tompkins (2005). Flexible complete models with stochastic volatility: generalizing Hobson-Rogers. *Working Paper*.
- Lam, K., E. Chang and M.C. Lee (2002). An empirical test of the variance gamma option pricing model. *Pacific-Basin Finance Journal* 10, 267-285.
- Macbeth, J.D and L.J. Merville (1980). Tests of the Black-Scholes and Cox call option valuation models, *The Journal of Finance* 35(2), 285-301.
- Madan, D.B. and E. Seneta (1990). The Variance Gamma model for share market returns, *Journal of Business* 63 (4), 511-524.
- Madan, D.B, P. Carr and E.C. Chang (1998). The Variance Gamma process and option pricing. *European Finance Review* 2, 79-105.
- Merton, R.C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141-183.
- Merton, R.C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125-144.
- Pan, J. (2002). The jump-risk premia implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3-50.
- Platania A, and L.C.D. Rogers (2005) Putting the Hobson-Rogers model to the test. *Working Paper*.
- Shroder, M. (1989) Computing the constant elasticity of variance option pricing formula, *Journal of Finance* 44 (1), 211-219.

Schwarz, G. (1978) Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.

Vuong, Q.H. (1989) Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. *Econometrica*, 57(2), 307-333.

Yuen, K.C., H. Yang and K.L. Chu (2001). Estimation in the constant elasticity of variance model, *British Actuarial Journal* 7(2), 275-292.

Appendix A: The Heston Put Formula

Cumulative distribution functions F and G in expression (2.11) can be computed from the following formula:

$$J(x) = \frac{1}{2} - \lim_{c \rightarrow \infty} \int_{-c}^c \frac{e^{-i\zeta x}}{2\pi i \zeta} \Psi_J(\zeta) d\zeta, \quad J \in \{F, G\},$$

where

$$\begin{aligned} \Psi_F(\zeta) &= \exp \left\{ i\zeta \ln S_t + g(T-t; \zeta) + h(T-t; \zeta) \sigma_t^2 \right\}, \\ \Psi_G(\zeta) &= \Psi_F(\zeta - i) / \Psi_F(-i), \\ g(T-t; \zeta) &= i\zeta r(T-t) + \frac{\alpha}{\gamma^2} \left[(-B + C)(T-t) - 2 \ln \left(\frac{1 - D e^{C(T-t)}}{1 - D} \right) \right], \gamma \neq 0, \\ h(T-t; \zeta) &= \frac{B - C}{\gamma^2} \frac{e^{C(T-t)} - 1}{1 - D e^{C(T-t)}}, \\ A &\equiv -(i\zeta + \zeta^2) / 2 \\ B &\equiv i\zeta \rho \gamma - \beta \\ C &\equiv \begin{cases} \sqrt{B^2 - 2A\gamma^2}, & \text{when } \gamma \zeta \neq 0 \\ -\beta, & \text{otherwise} \end{cases} \\ D &\equiv (B - C) / (B + C). \end{aligned}$$

Appendix B: The VG Put Formula

In the VG model, the price of a European put option with strike K , which matures t periods from now, can be found as

$$\begin{aligned} P_0(S_0, K, t) &= e^{-rt} \widehat{E}_0(\max(K - S_t, 0)) = e^{-rt} \int_{S_t < K} (K - S) f_{S_t}(S) dS \\ &= e^{-rt} \int_{\mathbb{T}_t} \int_{S_t < K} (K - S) f_{S_t | \mathbb{T}_t = \tau}(S) dS f_{\mathbb{T}_t}(\tau) d\tau = e^{-rt} \int_{\mathbb{T}_t} p(\tau) f_{\mathbb{T}_t}(\tau) d\tau, \end{aligned}$$

where

$$p(\tau) = \int_{S_t < K} (K - S) f_{S_t | \mathbb{T}_t = \tau}(S) dS.$$

Then, the conditional distribution of the stock price S_t given $\mathbb{T}_t = \tau$ is

$$S_t \sim S_0 \exp \left\{ (r - \delta + \theta)t + \widehat{\gamma}\tau + \widehat{\sigma}\sqrt{\tau}Z \right\},$$

where $Z \sim N(0, 1)$, and it follows that


$$\begin{aligned}
p(\tau) &= \int_{Z < d} \left(K - S_0 e^{(r-\delta+\theta)t + \hat{\gamma}\tau + \hat{\sigma}\sqrt{\tau}Z} \right) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}Z^2} dZ \\
&= K\Phi(d) - S_0 e^{(r-\delta+\theta)t + \hat{\gamma}\tau} \int_{Z < d} e^{\hat{\sigma}\sqrt{\tau}Z} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}Z^2} dZ \\
&= K\Phi(d) - S_0 e^{(r-\delta+\theta)t + \hat{\gamma}\tau + \frac{1}{2}\hat{\sigma}^2\tau} \int_{X < d - \hat{\sigma}\sqrt{\tau}} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}X^2} dX \\
&= K\Phi(d) - S_0 e^{(r-\delta+\theta)t + \hat{\gamma}\tau + \frac{1}{2}\hat{\sigma}^2\tau} \Phi(d - \hat{\sigma}\sqrt{\tau}),
\end{aligned}$$

where


$$d = \frac{\ln(K/S_0) - (r - \delta + \theta)t - \hat{\gamma}\tau}{\hat{\sigma}\sqrt{\tau}}$$

and Φ is standard normal cumulative distribution function.

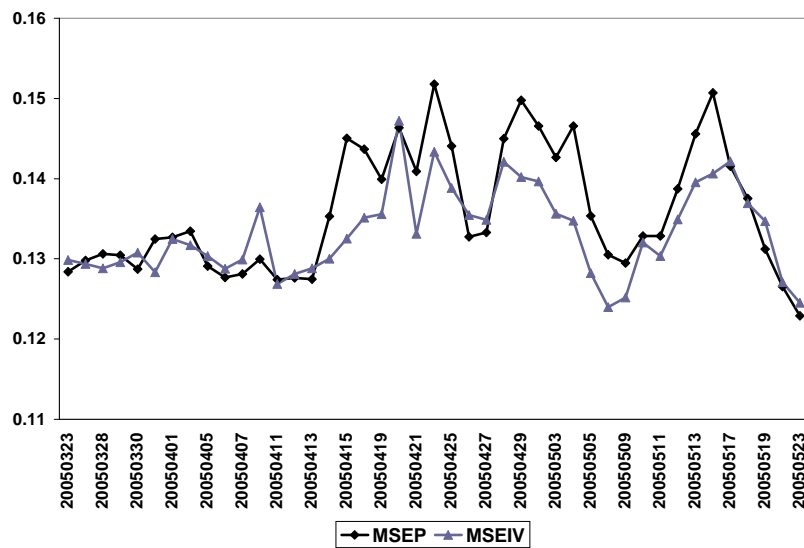
Appendix C: NLLS Estimates



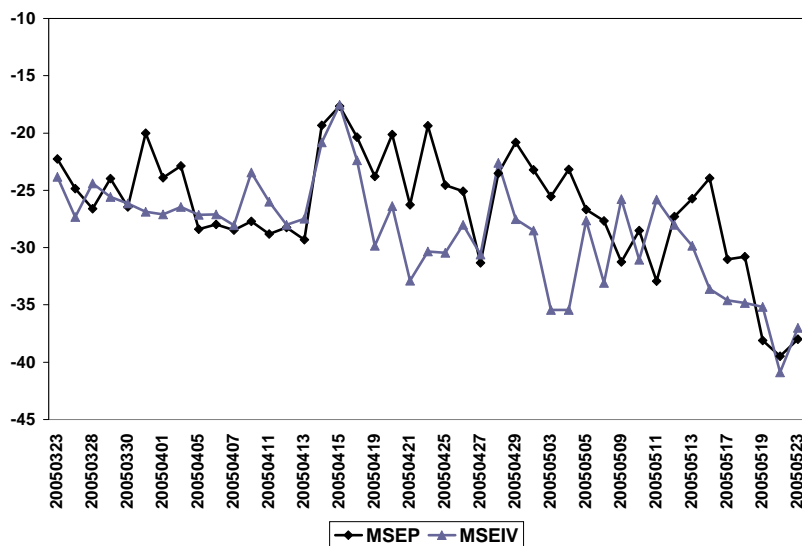
The Heston Model: γ




The Heston Model: ρ




The Hobson-Rogers Model: η_1



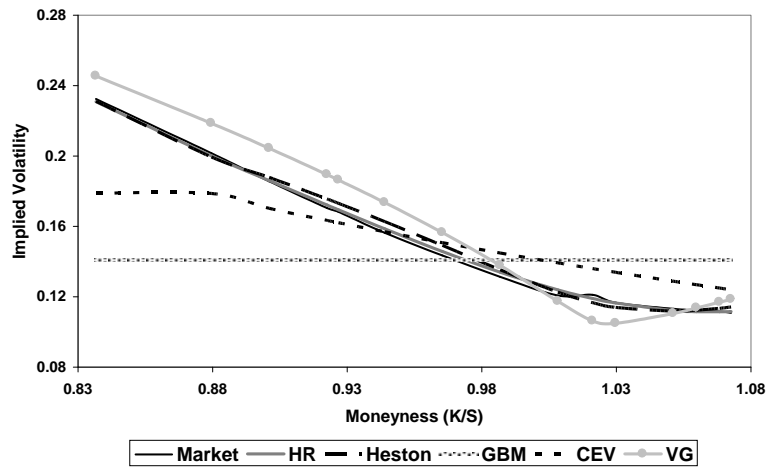
The Hobson-Rogers Model: η_2




The Hobson-Rogers Model: η_3

The Hobson-Rogers Model: λ


Appendix D: Implied Volatilities



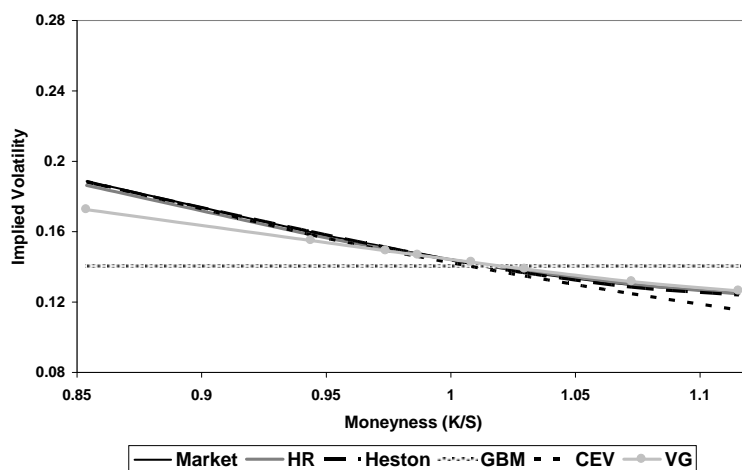
Implied volatilities of contracts traded on March 29, 2005
with maturity 38 days



Implied volatilities of contracts traded on March 29, 2005
with maturity 57 days



Implied volatilities of contracts traded on March 29, 2005
with maturity 120 days



Implied volatilities of contracts traded on March 29, 2005
with maturity 180 days

Appendix E: Out-Of-Sample Fit to Option Prices

The MSE column reports the average (AVG), median (MED) and standard deviation (SD) of daily mean squared errors of forecasts. The RMSE column reports the average (AVG), median (MED) and standard deviation (SD) of daily root mean squared errors of forecasts. The % column reports the number of days on which the MSE of the GBM/CEV/VG/Heston model was lower than than of the CEV/VG/Heston/HR model.

	days	MSE			RMSE			%
		AVG	MED	SD	AVG	MED	SD	
GBM	1	1.21e-03	1.21e-03	3.3e-04	3.54e-02	3.48e-02	4.7e-03	0.00
	2	1.27e-03	1.22e-03	3.7e-04	3.50e-02	3.49e-02	5.1e-03	0.00
	3	1.31e-03	1.23e-03	4.3e-04	3.57e-02	3.50e-02	5.7e-03	0.00
	4	1.34e-03	1.24e-03	4.3e-04	3.61e-02	3.53e-02	5.6e-03	0.00
	5	1.38e-03	1.31e-03	4.3e-04	3.67e-02	3.62e-02	5.6e-03	0.00
CEV	1	2.08e-04	1.82e-04	8.4e-05	1.42e-02	1.35e-02	2.7e-03	0.26
	2	2.31e-04	1.97e-04	1.1e-04	1.48e-02	1.40e-02	3.3e-03	0.30
	3	2.50e-04	1.92e-04	1.5e-04	1.53e-02	1.39e-02	4.1e-03	0.28
	4	2.78e-04	2.05e-04	1.9e-04	1.60e-02	1.43e-02	4.8e-03	0.35
	5	2.95e-04	2.10e-04	2.1e-04	1.63e-02	1.45e-02	5.4e-03	0.42
VG	1	1.60e-04	1.29e-04	8.6e-05	1.23e-02	1.14e-02	3.0e-03	0.07
	2	1.93e-04	1.48e-04	1.4e-04	1.32e-02	1.22e-02	4.2e-03	0.02
	3	2.27e-04	1.63e-04	1.8e-04	1.42e-02	1.28e-02	5.1e-03	0.09
	4	2.62e-04	1.77e-04	2.2e-04	1.52e-02	1.33e-02	5.6e-03	0.05
	5	2.76e-04	2.01e-04	2.1e-04	1.56e-02	1.42e-02	5.7e-03	0.02
Heston	1	6.79e-05	4.69e-05	7.0e-05	7.50e-03	6.84e-03	3.4e-03	0.16
	2	6.74e-05	4.92e-05	6.8e-05	7.61e-03	7.01e-03	3.1e-03	0.26
	3	8.04e-05	4.31e-05	8.8e-05	8.08e-03	6.56e-03	3.9e-03	0.40
	4	8.38e-05	6.27e-05	8.0e-05	8.43e-03	7.92e-03	3.6e-03	0.42
	5	7.81e-05	5.97e-05	6.4e-05	8.27e-03	7.73e-03	3.1e-03	0.49
HR	1	6.65e-05	2.80e-05	1.1e-04	6.51e-03	5.29e-03	4.9e-03	–
	2	1.08e-04	3.33e-05	2.2e-04	7.98e-03	5.77e-03	6.7e-03	–
	3	1.57e-04	5.55e-05	2.7e-04	9.94e-03	7.45e-03	7.7e-03	–
	4	1.45e-04	5.17e-05	2.4e-04	9.68e-03	7.19e-03	7.2e-03	–
	5	1.41e-04	6.15e-05	2.3e-04	9.89e-03	7.84e-03	6.6e-03	–