Derivative Pricing under Multivariate Stochastic Volatility Models with Application to Equity Options

Ph.D. Thesis Proposal

Mihaela Şerban
Department of Statistics,
Carnegie Mellon University

October 11, 2006

Abstract

We examine the joint time series of option prices and returns on the S&P 500 index and a set of stocks drawn from the index with a new arbitrage-free multivariate stochastic volatility model that captures a market effect. The preliminary results show that the new model fits well for all the marginal time series for different periods of time. The price of volatility risk is estimated from option prices simultaneously for stocks and the index. We find that the volatility risk premia is not constant over time and that it differs across stocks. However, our model yields very small option pricing errors both in-sample and out-of-sample. Our future work will focus on (1) improving the estimation from the joint data of returns and option prices, and test the statistical significance of model parameters, (2) compare our modeling approach with existing univariate models from a trading perspective, and (3) extend our model with a latent process to describe the price of volatility risk.

1 Introduction

Today options play an increasingly important role in financial markets. Within ten years of the establishment of the first options organized exchange in 1973, the volume of trading in stock options grew to a level often exceeding, in terms of share equivalents, that of the New York Stock Exchange. Huge volumes are now traded on many exchanges throughout the world and also over the counter by financial institutions. The underlying assets include stocks, market indices, foreign currencies, futures and debt instruments.

Option valuation was first satisfactorily developed by Black and Scholes (1973) and Merton (1973). The authors derived the arbitrage-free price of a European option under the assumption that the stock price follows a geometric Brownian motion with constant drift (i.e. constant risk-free rate) and constant volatility. The main problem with their derived option price is that it is often significantly different from the actual trading price that one can observe in the market. Such systematic valuation errors are often attributed to the fact that volatility does not remain constant over time. In addition to temporal variation, it is also well known that the Black, Scholes, and Merton (BSM) implied volatility (i.e. the resulting volatility parameter from matching the observed option price with the BSM price) exhibits cross-sectional variation, known as the "smile" effect. The name came from the fact that, prior to the October 1987 market crash, the BSM implied volatility function of the S&P 500 index options had the shape of a smile across strike prices. However, after October 1987, some studies suggested that the index implied volatility function decreases as the strike price increases.

The implied volatility describes a process which is often different from the realized volatility (see, Figure 1). To hedge the risk of the volatility process, buyers of options pay a premium, called *volatility risk premium* or the *price of volatility risk*. Empirically, it was found that the pricing of the volatility risk in options is a possible explanation for the difference between implied and realized volatility (see Bakshi and Kapadia (2003)). In the context of option valuation, three important questions arise. Is the volatility risk premium statistically significant in options? Does the price of volatility risk vary across time? Do different assets have different prices of volatility risk? If so, is there any common component across assets, such as market or industry? This thesis tries to address these questions on the price of volatility risk by developing a multivariate model that allows one to examine and capture features of the price of volatility risk for stocks and the S&P 500 index simultaneously.

The volatility risk premium was examined with univariate volatility models, and for index options mainly (see Bates(2000), Pan(2002), Eraker (2004) among others). In the context of index options, there is no common conclusion on whether the volatility risk premium is significantly different from zero or whether it is constant over time and strike prices. Also, literature gives much more attention to index options than to equity options. Literature investigating the properties of volatility risk premium from equity options is almost nonexistent. However, Bakshi and Kapadia (2003) studied equity options and found that individual equity options embedded a negative market volatility risk premium, in general, assuming that idiosyncratic volatility was not priced.

The models developed to study the price of volatility risk belong to a large class of volatility models, known as stochastic volatility (SV) models. This class of models is popular for two main reasons. First, it captures the return volatility term structure while explaining the smile effect. Secondly, options can be valued under some stochastic volatility models analytically (see Heston model, 1993) or numerically (for example, by Monte Carlo simulations or by finite difference methods). However, to study the price of volatility risk across stocks, we need to capture both the term structure of return volatilities for different assets and the dynamics of the cross-correlations.

Straightforward extensions of popular univariate stochastic volatility models (Heston (1993) or its generalization developed by Bates (2000)) may be computational infeasible to estimate. The calibration for univariate SV models is already very complex due to the latent stochastic volatility, an extension to a multivariate framework involving a large number of parameters would make the model difficult to fit even a 3-dimensional return series. In this thesis we propose a novel multivariate stochastic volatility model that describes the evolution of the covariance matrix through an observed common factor and latent idiosyncratic stochastic volatilities. We fit our model to a 13-dimensional time series of returns using the maximum likelihood criterion. Preliminary results show that our multivariate model fits all marginal return series well. However, returns alone cannot give information about the price of volatility risk. We need option price data (or implied volatilities computed from options) to estimate this premium. Therefore, we add to our model a parameter for each asset, to represent the volatility risk premium and we estimate it from option prices, for all assets simultaneously. The estimation of the volatility risk premia from option price data uses the parameter estimates obtained from returns data. In other words, we first use the returns data to estimate one set of parameters, and then conditioned on these estimates, we calibrate the price of volatility risk from option prices for each asset. Our results show that the price of volatility risk is different across assets and varies over time when estimated for different time windows. However, we did not obtain standard errors for our parameter estimates. We are particularly interested in estimating the standard errors for the price of volatility risk to test whether it is statistically significant. Most empirical studies don't report standard errors for the price of volatility risk since they fit their model to returns and option prices separately. Nevertheless, there are few authors who estimated all model parameters from the S&P 500 index returns and option prices jointly. These authors used methods such as Generalized Method of Moments (GMM) (see Pan (2002)), Efficient Methods of Moments (EMM) (see Chernov and Ghysels (2000)), and Bayesian methods (see Eraker (2004)). However, given the difficulty in estimating SV models, all previous empirical studies on the joint time series of returns and options have been for a single underlying (i.e, an index or an interest rate). Our work is the first estimation of a multivariate stochastic volatility model from the joint time series of returns and derivatives on different underlying assets. Moreover, while most researchers worked only with European options (i.e., index options), we study also American-style options for individual stocks, options that are more difficult to price than European options. Therefore, methods previously used to fit univariate returns and option prices jointly, may be inefficient or difficult to implement for our multivariate model and data. As a consequence, we develop a stochastic EM-type algorithm to calibrate our model to returns and option price data jointly.

The next section discusses existing modeling approaches to value options, and methods to estimate model parameters from returns and option prices. Section 3 presents preliminary results from fitting our proposed model to returns and option price data. Section 4 discusses work that remains to be done, including an estimation procedure to calibrate our model to returns and option price data jointly, a comparison of our modeling approach with existing univariate models from an application perspective, and an extension of our model to describe the dynamics of the volatility risk premium.

2 Literature Overview

Most attempts to explain the BSM implied volatility smile concentrate on relaxing the BSM assumption of constant volatility by allowing the return volatility to evolve through time deterministically or stochastically. Two deterministic time-varying volatility models have gained popularity in practice, the Cox and Ross (1976) constant elasticity of variance (CEV) model, and the implied binomial tree models of Dupire (1994), Derman and Kani (1994), and Rubinstein (1994). Emanuel and MacBeth (1982) found that in sample, the CEV model fits option price data better than the BSM constant volatility model, but the CEV model does no better than the BSM model out-of-sample. On the other hand, the problem with the implied binomial tree framework is that it provides flexible models that overfit the option prices at any point in time. However, empirically it was found that the parameters of simple deterministic volatility models are highly unstable through time (see Dumas, Fleming, and Whaley, 1998). All these observations suggest that deterministic volatility models may not be able to explain the shape of the implied volatility curve.

2.1 Stochastic Volatility Models

Empirical studies on option pricing with stochastic volatility (SV) models showed that these models have the potential to explain the shape of the implied volatility curve. Moreover, these models seem to allow for large market movements or fat tails in the returns distribution by assuming a stochastic element in the time evolution of the conditional variance process. Two popular examples of option pricing models with SV are Hull and White (1987), and Heston (1993). Let $S(t)$ and $v(t)$ denote the spot price and the return variance at time t, respectively. The model introduced by Hull and White assumes the following stochastic processes:

$$dS(t) = \mu S dt + \sigma(t) S dW_1(t), \tag{1}$$

$$v(t) = \sigma^2(t)$$

$$dv(t) = \phi v dt + \xi v dW_2(t), \tag{2}$$

where W_1 and W_2 are two correlated Wiener processes, μ, ϕ and ξ are unknown parameters, and $\sigma(t)$ is the return volatility process.

The Heston model assumes that the spot asset $S(t)$ follows the diffusion process given by (1) and that the volatility $\sigma(t)$ follows an Ornstein-Uhlenbeck process (OU), or that the variance process $v(t)$ follows the square-root process (Cox, Ingersoll, and Ross)

$$dv(t) = k(\theta - v)dt + \xi\sqrt{v} dW_2(t), \tag{3}$$

where W_1 and W_2 are two correlated Wiener processes, and k (the speed of mean reversion), θ (long-run mean) and ξ are unknown parameters. Heston provides a closed-form solution for European options using Fourier transform methods, while Hull and White value European options in a series form for the case in which W_1 and W_2 are independent.

Another example of diffusion models is the class of exponential linear models (Tauchen, 2004). In this class, the basic model is given by two equations, equation (1) for the asset price evolution, and equation (4) for the log-volatility dynamics:

$$\sigma_t = \exp(\xi_0 + \xi_1 v(t))$$

$$dv(t) = (\alpha_0 + \alpha_1 v)dt + \alpha_2 dW_2(t), \tag{4}$$

where ξ_1, ξ_2, α_0, α_1, and α_2 are unknown parameters.

There have been numerous theoretical extensions to models (1)-(4). However, few of them have been tested against their simpler form (for example, the Heston model) empirically. One extension was the introduction of jumps in the underlying and/or volatility processes. A number of empirical studies investigated whether the stochastic volatility jump-diffusion specifications could explain the pricing biases in the BSM model. Most papers invoked the October 1987 market crash for the deficiency of an SV model to fit the return series and the option prices. In particular, recent work by Bates (2000), Pan (2002) and Eraker (2000, 2004) focused on whether incorporating discontinuous jumps in the asset price and/or volatility improved the fit over simpler SV models to S&P 500 return data and index option prices. Bates (2000) found, for example, that stochastic volatility significantly reduced option pricing errors compared to the BSM model. His general model was an extension of the Heston model allowing for Poisson jumps in the asset price with constant or time-varying intensity. This jump diffusion model seemed to bring a slight improvement over the Heston model in terms of option errors. However, his approach used option prices only, and consequently, parameters were estimated only under the risk neutral measure without being able to identify volatility and jump risk premia. Pan (2002) adopted the Bates (2000) model and used the joint data of returns and option prices to estimate model parameters. She concluded that jump diffusions were important in capturing variations in BSM volatilities, and that, the Heston model was strongly rejected for the joint data, as well as for the univariate data. Contrary to the results in Bates(2000) and Pan(2002), Eraker (2004) found the Bates model with jump components didn't seem to improve significantly upon the simpler stochastic volatility (Heston) model in-sample and out-of-sample option pricing. However, his findings showed that jump diffusion models fitted the time series dimension of the data reasonably well. Overall, researchers found that the Heston model calibrated with moment-based methods did not fit the returns and option prices on the S&P 500 well.

Another extension to models (1)-(4) was the introduction of a second unobserved volatility factor. Several other authors argued that the Heston model did not fit well long series of daily stock index returns (Anderson *et al* (2002), Chernov *et al* (2003), Eraker *et al* (2003)) because there was only one volatility process, which had to accomplish two purposes simultaneously. Specifically, it needed to capture the persistent stochastic volatility (with very small values of the mean reversion speed parameter k), and it needed to accommodate the fat tails (with very large values of k). Therefore, some researchers introduced a second volatility factor with a different mean reversion speed. Two-factor volatility structures have been developed empirically by Gallant, Hsu, and Tauchen (1999), Barndorff-Nielsen and Shephard (2001a, 2002a), Alizadeh, Brandt, and, Diebold (2002), among others. Also, Bakshi and Kapadia (2003) build a two-factor stochastic volatility model based on the assumption of a common latent factor in the marginal volatilities (see model given by (5)). However, they do not fit their model to returns and/or options, instead they are interested in deriving delta-hedged gains as a function of market volatility risk premium.

$$\begin{aligned}
\frac{dS_i(t)}{S_i(t)} &= \mu_i[S_i, V_i]dt + \sqrt{V_i(t)}dW^1(t), \\
V_i(t) &= \beta_i V_m(t) + Z_i(t), \\
dV_m(t) &= (\theta - kV_m)dt + \eta\sqrt{V_m}dW^2(t),
\end{aligned} \tag{5}$$

where $S_i(t)$ and $V_i(t)$ denote the stock price and the variance of firm i, $V_m(t)$ denotes the market index variance, and β, θ, k, η are unknown parameters. The overall conclusion in the above studies was that the two-factor volatility model did a better job than the one-factor model, but it still didn't fit the daily returns data well. Tauchen *et al* (1999) extended the two-factor model to a three-factor volatility model and found that this latter model did better than the former, arguing that the three-factor model mimicked the long-memory behavior of financial volatility.

Empirically, it was found that the dependence in the volatility structure initially decays quickly at short lags, and then the decay slows at longer lags, when volatility is measured using high frequency data over a few years or using daily data over decades. To capture this feature, researchers have been working with both discrete and continuous time long memory SV. Recent work focused on the square root model driven by fractionally integrated Brownian Motion (Comte, Coutin, and Renault, 2003) and on the infinite superposition of non-negative OU processes (Barndorff-Nielsen, 2001a). These two models have the advantage that option valuations can be performed without high computational cost. We will not investigate further these models since we use daily data for a time period of only a few years. The overall conclusion is that stochastic volatility models are suitable for returns data, and for option price data if volatility/jump risk premia is parameterized. Given this fact, we build our multivariate model with stochastic volatility models for the marginal returns series.

2.1.1 Multivariate Stochastic Volatility Models

Numerous empirical studies exist for pricing options under univariate stochastic volatility models, but there are none that investigate multivariate volatility models using the data of returns and option prices jointly. Moreover, the literature on multivariate volatility models is not rich. Various discrete time multivariate models have been proposed in the literature to parameterize the conditional covariance (see, for example the multivariate GARCH model introduced by Engle and Kroner (1995), the dynamic conditional correlation model of Engle (2002), and models based on copula distributions of Şerban, Brockwell, Lehoczky and Srivastava (2006)). An alternative approach for achieving a more manageable and parsimonious multivariate volatility model makes use of factor structures (see Diebold and Nerlove (1989) for example). Each factor is allowed to have its own latent dynamic structure, which is parameterized as an ARCH process. An empirical drawback of this approach is that the covariance matrix of the idiosyncratic innovations is assumed constant over time. A continuous-time multivariate stochastic volatility model is introduced by Gourieroux and Sufana (2004) in the context of derivative pricing. They propose a multivariate extension of the Heston model, modeling the dynamics of the covariance matrix as a Wishart autoregressive (WAR) process. However, the authors provide no empirical application of their model. Our study is the first to develop and fit a cross-sectional multivariate model to returns and option prices jointly.

2.2 Inference

A major difficulty with the use of discrete and, in particular, continuous time SV models is that they are hard to estimate because they don't have readily available likelihood functions. Formally, the SV likelihood is given as follows. Let $\mathbf{r} = (r_1, ..., r_T)$ and $\mathbf{v} = (v_1, ..., v_T)$ denote the vector of returns and volatilities, respectively. If Θ is the set of model parameters, then the probability density function of the data given Θ may be written as

$$f(\mathbf{r}; \Theta) = \int f(\mathbf{r}, \mathbf{v}; \Theta) d\mathbf{v} = \prod_{t=1}^{T} \int f(r_t | v_t; \Theta) f(v_t | \mathcal{F}_{t-1}; \Theta) dv_t. \tag{6}$$

For parametric SV models such as (2)-(4), the density $f(r_t|v_t; \Theta)$ is known in closed form, but $f(v_t|\mathcal{F}_{t-1}; \Theta)$ is not available. This led to two approaches in the literature to estimating SV models. Numerous authors focus on simple moments-based estimators, whereas others focus on computationally intensive methods such as simulated maximum likelihood methods and MCMC techniques.

In the first class of estimation methods, the Generalized Method of Moments (GMM) and the Efficient Method of Moments (EMM) have been used predominantly in the econometrics literature. The GMM was developed by Hansen (1982) from the method of moments procedure and is closely related to the classical theory of minimum chi-square estimation. However, this method is not particularly efficient, since the unconditional moments available in closed-form are different from the (efficient) score moments associated with the likelihood function. The EMM procedure (Gallant and Tauchen, 1996) has the potential to deliver efficient inference, since it is based on approximating the score (moment) vector. The main idea in the EMM is to use a tractable semi-parametric (SNP) density estimate to approximate the log-likelihood, and then approximate the (pseudo) score moments (given by SNP) using Monte Carlo simulations from the true SV model. An important drawback for GMM and EMM is the lack of any estimates of the underlying latent state variable. Some authors used the Kalman filter to estimate the latent variable even though the simplest SV model assumes some non-Gaussian components. However, the Kalman filter has been generally found to perform poorly when the noise is non-gaussian. Gallant and Tauchen (1998) developed a technique within EMM, called re-projection, in which they used the same idea of a tractable SNP density estimate in order to evaluate the one-step-ahead conditional mean and variance analytically. The EMM technique is not necessarily very efficient, and the re-projection method for estimating the latent variable is numerically intensive.

In the second class of calibration methods, the MCMC approach worked well for discrete SV models by Shephard (1993), and Jacquier, Polson, and Rossi (1994). A key advantage of the MCMC algorithm is that the distribution of the latent state vector is obtained as an implicit part of the estimation. The procedure solves the smoothing problem of determining $f(v_{t+j}|\mathbf{r_T})$, but it doesn't yield the filtering distribution $f(v_t|\mathbf{r_t})$ which is important in many applications. However the filtering problem can be tackled via so-called particle filters (see, for example, Gordon, Salmond, and Smith (1993), Doucet, deFreitas, and Gordon (2001), and Pitt and Shephard (1999), for more details). The extension of the MCMC technique to a continuous-time setting is discussed in Elerian, Chib, and Shephard (2001) and Eraker(2001). In summary, the MCMC approach works well for many practical problems, but in more complex settings there might be serious issues with its implementation.

In general, most estimation procedures developed for non-linear state-space models can be adopted for continuous-time models. Some of the SV models have closed-form conditional distributions readily available in continuous time, and the above mentioned techniques can be applied. Some, however, don't have closed-form distributions, but it is possible to discretize the model by using first order (Euler scheme) and second order (Milestein scheme) approximations. Moreover, Monte Carlo methods (such as the bootstrap filter of Gordon *et al* (1993), or the method introduced in Brockwell (2005)) can be used to approximate the likelihood function (6) for SV models in (1)-(4).

Inference from returns only is not sufficient to predict option prices. According to modern asset pricing theory, the value of an option can be computed as the expectation under the risk-neutral measure of the discounted future payoff. Therefore, options are priced not using the true process, but using a corresponding risk-neutral process that incorporates the appropriate compensation for volatility risk and/or jump risk. SV models under a risk-neutral measure usually have the same parametric structure as the ones under the true measure. They also have common parameters. The difference between the objective and risk-neutral measures defines the risk premia. It has been documented in the literature (Doran and Ronn (2004), Eraker (2004) for example) that the price of volatility risk, for example, is of key importance for pricing options. However, information on the volatility risk premium (which is usually the risk-neutral specific parameter) cannot be obtained only from the return data. Therefore, there is a growing body of research which advocates the use of the joint data of returns and option prices to fit SV models. Papers by Chernov and Ghysels (2000), Bates (2000), Pan (2002), Jiang (2002), Jones (2003), Eraker (2004) all use some combination of a time series of the underlying price and one or more time series of liquid, short-term, near-the-money option prices to estimate the model parameters under both the objective and risk-neutral measures. Most of the above authors use moments-based inference methods, except for Eraker(2004) and Jones(2003) who used Bayesian analysis to arrive at their estimates.

2.2.1 Volatility Risk Premium

There is no common conclusion on the properties of the volatility risk premium. Some authors argued that this premium was statistically significantly different from zero (Melenberg and Werker (2001), for EOE index options; Bakshi and Kapadia (2001), and Coval and Shumway (2001) for options on S&P 500 and S&P 100), and some argued that while negative, the volatility risk premium was insignificant for S&P 500 options (Pan (2002), Eraker (2004)). Doran and Ronn (2004) documented the impact of the market price of volatility risk on volatility bias through simulations. Their findings suggest that across several models the market price of volatility risk is a critical parameter that generates the differences in risk-neutral (BSM implied volatility) and objective volatility (realized volatility). Moreover, Melenberg and Werker (2001) studied the time-series properties of the volatility risk premiums (estimated at daily frequency) and concluded that they were non-constant and exhibit significant autocorrelation.

Literature investigating the properties of the volatility risk premium from the equity option market is almost nonexistent. However, Bakshi and Kapadia (2003) investigated the market volatility risk premium for the equity option market with the model given by (5). They found that individual equity option prices embedded a negative market volatility risk premium, in general. However, they showed that the volatility risk from equity options was much smaller than that from index options, and idiosyncratic volatility is not priced. Their conclusion is based on regressing stock delta-hedged gains versus market realized volatility and stock realized volatility averaged over all stocks, obtaining a significant coefficient for the market volatility and an insignificant one for the stock return volatility. However, their empirical results are questionable, because, for example, they used realized market and stock volatilities, rather than the market and idiosyncratic volatilities filtered from their model.

3 Preliminary Results

In this section we propose a new continuous-time multivariate stochastic volatility model and demonstrate that it fits well the return series data from 12 stocks representing three industrial sectors and the S&P 500 index. We also evaluate this model from an option pricing perspective.

3.1 The model

Our multivariate diffusion model builds on the idea that stock returns have a market component, which we assume is observed through the S&P 500 index. Other factors can be considered as well, such as the ones from the discrete-time Fama-French model. However, Figure 2 shows, for the example of 12 stocks, that one eigenvalue dominates the other ones over time. Therefore, for our initial study we restrict ourself to considering only one factor in predicting the stock returns. We also specify a diffusion model for the market factor, and model both the idiosyncratic stock return volatility and the market (index) return volatility as autonomous stochastic processes. This model introduces four sources of uncertainty to the underlying stock price dynamics: (1) stock return shocks, (2) market return shocks, (3) stock volatility shocks, and (4) market volatility shocks. The prices of these risk factors are parameterized and estimated in Section 3.4.

3.2 The data-generating process

Let $(\Omega, \mathcal{F}, P)$ be a probability space and $(\mathcal{F}_t)$ an information filtration satisfying the usual conditions (see, e.g. Protter, 1990). We denote by $\mathbf{S} = (S_1, ..., S_N)$ and by I an N-dimensional vector of stock price processes and the index process, respectively. Our model assumes the following data-generating process for the stock price vector $\mathbf{S}$ and for the index I

$$
\begin{aligned}
\frac{dS_k(t)}{S_k(t)} &= (\mu_k - \beta_k \mu_I)dt + \beta_k \frac{dI(t)}{I(t)} + \sqrt{v_{S_k}(t)} dW_k^{(S)}(t), \\
\frac{dI(t)}{I(t)} &= \mu_I dt + \sqrt{v_I(t)} dW_I^{(S)}(t), \\
v_{S_k}(t) &= \exp(x_{S_k}(t)), \\
dx_{S_k}(t) &= \alpha_k(b_k - x_{S_k}(t))dt + \sigma_k d\epsilon_k^{(X)}(t), \\
d\epsilon_k^{(X)}(t) &= \gamma_k d\overline{W}^{(X)}(t) + dW_k^{(X)}(t), \\
v_I(t) &= \exp(x_I(t)), \\
dx_I(t) &= \alpha_I(b_I - x_I(t))dt + \sigma_I dW_I^{(X)}(t),
\end{aligned} \tag{7}
$$

where, for $k = 1, .., N$, $W_k^{(S)}(t)$, $W_I^{(S)}(t)$, $\overline{W}^{(X)}(t)$, $W_k^{(X)}(t)$ and $W_I^{(X)}(t)$ are standard Brownian motions with the following correlation parameters

$$
\begin{aligned}
0 &= dW_k^{(S)}(t) d\epsilon_k^{(X)}(t), \\
0 &= dW_I^{(S)}(t) dW_I^{(X)}(t), \\
\rho_k^{SI} &= dW_k^{(X)}(t) dW_I^{(X)}(t), \\
0 &= d\overline{W}^{(X)}(t) dW_k^{(X)}(t), \\
\rho &= d\overline{W}^{(X)}(t) dW_I^{(X)}(t) \\
0 &= dW_k^{(S)}(t) dW_I^{(S)}(t).
\end{aligned}
$$

Applying Îto's Lemma to the first two equations in (7), we obtain the following model for stock and index log-returns

$$
\begin{aligned}
d\log(S_k(t)) &= (\mu_k - \beta_k \mu_I - \tfrac{v_{S_k}(t)}{2} - \beta_k(\beta_k - 1)\tfrac{v_I(t)}{2})dt + \beta_k d\log(I(t)) \\
&+ \sqrt{v_{S_k}(t)} dW_k^{(S)}(t), \\
d\log(I(t)) &= (\mu_I - \tfrac{v_I(t)}{2})dt + \sqrt{v_I(t)} dW_I^{(S)}(t), \\
v_{S_k}(t) &= \exp(x_{S_k}(t)), \\
dx_{S_k}(t) &= \alpha_k(b_k - x_{S_k}(t))dt + \sigma_k d\epsilon_k^{(X)}(t), \\
d\epsilon_k^{(X)}(t) &= \gamma_k d\overline{W}^{(X)}(t) + dW_k^{(X)}(t), \\
v_I(t) &= \exp(x_I(t)), \\
dx_I(t) &= \alpha_I(b_I - x_I(t))dt + \sigma_I dW_I^{(X)}(t).
\end{aligned} \tag{8}
$$

This model captures important features of the stock returns dynamics, namely, the market effect and the stochastic return volatility. First, the market effect is modeled by the process I that introduces two sources of uncertainty to the stock return processes, the index return shock $W_I^{(S)}$ and the index volatility shock $W_I^{(X)}$, which initially are assumed to be independent. The index return log-volatility x_I follows an OU process with constant long-run mean level b_I, speed of mean-reversion α_I, and volatility coefficient σ_I. Higher values of σ_I imply that the asset price distribution will have fatter tails. Second, the idiosyncratic stock return volatility is modeled by the process v_{S_k} which introduces the shock $\epsilon_k^{(X)}$. The stock return log-volatility x_{S_k} follows an OU process with constant long-run mean b_k, speed of mean-reversion α_k, and volatility coefficient σ_k, for $k = 1, 2, ..., N$. We further expand on the covariance structure between stock return volatilities by assuming a common shock $\overline{W}$ and an idiosyncratic shock $W_k^{(X)}$, both correlated with the index log-volatility shock $W_I^{(X)}$ with coefficients ρ and ρ_k^{SI}, respectively. The intuition behind these correlations is that the index is nothing other than a weighted average of stocks (we fit only stocks from S&P 500) and thus its variance process will be a function of the noise processes from the stock variance process.

There is evidence in the literature (e.g. Pan(2002)) that the shock in the index return, $W_I^{(S)}$, and the shock in the index volatility, $W_I^{(X)}$ are correlated with a significant constant coefficient. This correlation parameter is typically found to be negative and implies that a fall in prices usually will be accompanied by an increase in volatility, which is sometimes known as the "leverage effect". For model simplification, we assume that the return shocks and the log-volatility shocks are not correlated for both stock and index returns. Investigations on whether these correlations are significantly non-zero are left for future work.

3.3 Inference from the returns data

Since we sample the returns at discrete times, we need a discrete version of the continuous-time model (8). The relatively simple form of the differential equations in (8) allows us to obtain analytical solutions for the return conditional distribution and the log-volatility transition distribution. The resulting discrete version of model (7) is given by (14) in the appendix.

We fit the discrete model (14) to daily returns on 12 stocks (from Table 5) and the S&P 500 index (source: CRSP from Wharton Research Data Services). The stock data we analyzed consist on returns from daily closing prices from 01/2003 to 12/2004. We chose these particular stocks because they belong to three different industrial sectors (IT, Finance and Materials) and because the options on these stocks are more liquid than options on other stocks belonging to the same industry. We estimated the parameters from model (14) by using the maximum likelihood criterion. The likelihood function (6) was approximated by using the bootstrap filter (Gordon *et al* 1993). Goodness-of-fit was tested using the same particle filter used for estimation. We obtained a set of residuals (see Brockwell, 2005 for more details) by first obtaining one-step predictive marginal cumulative distributions and then converting them into normally distributed "residuals". Then, we applied standard methods to test the normality of the residuals. In particular, we used the Kolmogorov-Smirnov (KS) test. Parameter estimates and p-values for goodness-of-fit tests, using data from year 2003 for example, are presented in Table 2. The results in Table 2 show that model (14) fits well both the stock and index marginal distributions. Moreover, we performed likelihood ratio tests to examine whether the parameters β_k, ρ_k^{IS}, and γ_k are significantly different from zero. We can see that all the β parameters, except for one, are significantly non-zero at level 1%, while most of the parameters ρ_k^{IS} and γ_k are significantly non-zero at level 5%. We found that the correlation coefficient ρ is not significantly different from zero, therefore we'll exclude it from the model in our further analysis. All these results are consistent over other periods of time.

In the next section we consider statistical inference from the option data. In particular, we want to investigate the dynamics of risk-neutral specific parameters (the volatility risk premia). Before that, we investigate the stability of the model parameters under the objective measure. We fit model (14) to an overlapping moving window of 251 observations (one year) beginning with February 2003 and ending with August 2004. For example, we obtain one set of parameters $\Theta_{obj} = (\{\mu_k\}, \mu_I, \{\beta_k\}, \{b_k\}, \{\alpha_k\}, \{\sigma_k\}, b_I, \alpha_I, \sigma_I, \{\rho_k^{IS}\})$ for the time period February 2003 - January 2004, then another set for March 2003 - February 2004, and so on up to August 2004, in total there are 7 sets of parameters estimates. All these estimates are plotted in Figures 3-9. The estimates for ϕ, b, σ, and β don't seem to vary over time for each asset, instead the estimates for ρ^{IS} and γ do exhibit time variation for some of the stocks. The time series properties of these particular parameters will be investigated in future work.

3.4 The price of volatility risk

We mentioned in Section 2.2 that options are priced under a risk-neutral measure. Therefore, to use the time series of option prices, we need to specify a model under this measure besides the model under the objective measure. Following the common practice in econometrics and because our new model does a good job in fitting the returns under the true measure, we assume that the continuous-time model under the risk-neutral measure is given by

$$\begin{aligned}
\frac{dS_k(t)}{S_k(t)} &= r(1-\beta_k)dt + \beta_k\frac{dI(t)}{I(t)} + \sqrt{v_{S_k}(t)}dW_k^{(S)}(t), \\
\frac{dI(t)}{I(t)} &= rdt + \sqrt{v_I(t)}dW_I^{(S)}(t), \\
v_{S_k}(t) &= \exp(x_{S_k}(t)), \\
dx_{S_k}(t) &= \alpha_k(b_k - x_{S_k}(t) - \lambda_k\sigma_k)dt + \sigma_k d\epsilon_k^{(X)}(t), \\
d\epsilon_k^{(X)}(t) &= \gamma_k d\overline{W}^{(X)}(t) + dW_k^{(X)}(t), \\
v_I(t) &= \exp(x_I(t)), \\
dx_I(t) &= \alpha_I(b_I - x_I(t) - \lambda_I\sigma_I)dt + \sigma_I dW_I^{(X)}(t),
\end{aligned} \tag{9}$$

where, for $k = 1,..,4$, $W_k^{(S)}(t)$, $W_I^{(S)}(t)$, $\overline{W}^{(X)}(t)$, $W_k^{(X)}(t)$ and $W_I^{(X)}(t)$ are standard Brownian motions with the following correlation parameters

$$\left\{\begin{aligned}
0 &= dW_k^{(S)}(t)d\epsilon_k^{(X)}(t), \\
0 &= dW_I^{(S)}(t)dW_I^{(X)}(t), \\
\rho_k^{SI} &= dW_k^{(X)}(t)dW_I^{(X)}(t), \\
0 &= d\overline{W}^{(X)}(t)dW_k^{(X)}(t), \\
0 &= d\overline{W}^{(X)}(t)dW_I^{(X)}(t) \\
0 &= dW_k^{(S)}(t)dW_I^{(S)}(t).
\end{aligned}\right.$$

The difference between equations (7), the model under the objective measure, and equations (9), the model under a risk-neutral measure, is two-fold. First, under the risk-neutral measure, the stock and index mean return rate is r, the risk free interest rate which we assume to be constant. Second, in (9), we introduce new parameters $\Theta_{neutral} = (\{\lambda_k\}, \lambda_I)$

which measure the price of volatility risk for each stock and for the index. The rest of the parameters are common to both measures. In our current study, we plugged in the estimates obtained from the returns data (see Section 3.3) and estimated $\Theta_{neutral}$ from option prices. Some authors parameterized the volatility premium as being proportional to the level of volatility (Heston, 1993, Pan, 2002 among others). As a first step in our analysis, we assume $\Theta_{neutral} = (\{\lambda_k\}, \lambda_I)$ to be constant over time, and constant over option moneyness (strike/spot price) and maturity. We will then expand our model if these assumptions prove to be inappropriate.

The option data are from OptionMetrics of Wharton Research Data Services. OptionMetrics contains data on all US exchange-listed index and equity options. For each observation day (from 01/2003 to 12/2004), we collected all bid-ask closing quotes for put options only. We removed all option records with open interest zero, options with maturity longer than 6 months, and options with moneyness less than 0.9 (deep out-of-the money options). The reason we eliminated options deep out-of-the-money is that our model doesn't account for default risk, which might be taken into account for deep out-of-the-money options.

We computed the theoretical option prices (under model (9)) using Monte Carlo methods. The value of a put on the S&P 500 index, which is an European option, at time t is

$$p_t = \widetilde{E_t}((K - S_T)^+),$$

where $\widetilde{E_t}$ is the conditional expectation under a risk-neutral measure, K is the option strike price, and T is the option maturity. Therefore, simulations from model (9) together with variance reduction methods can be applied directly to price European options. However, equity options are American options, and their valuation is more complex than pricing European options. This is because an American option can be exercised at any point in time until expiration. Thus, besides approximating the expected value of the discounted payoff, one needs to approximate the early exercise premium. We used the least-square method (LSM) introduced by Longstaff and Schwartz (2001) in order to value American equity options.

Using the parameter estimates under the objective measure, we estimated the volatility risk premia from option price data for each stock and for the index at monthly and weekly frequency. In other words first we fitted model (7) to returns from the period of time February 2003-January 2004 and obtained estimates for Θ_{obj}. Next, we used a non-linear regression to obtain estimates for $\Theta_{neutral}$ using option prices (average of bid and ask quotes) observed in every month/week beginning with January 2004 and up to July 2004. Plots of these estimates at a monthly frequency are in Figures 10-11 and at a weekly frequency in Figures 12-13. Three things are worth mentioning. First, the parameter estimates seem to vary over time for almost all stocks. Second, the prices of volatility risk are different across stocks and from the index. Finally, stocks have both positive and negative volatility premia, while the index volatility risk price is always negative. A negative volatility risk premium means that whenever volatility is high, options are more expensive than what is implied by the objective measure as the investor asks for a higher premium. Conversely, in low-volatility periods, the options are less expensive. Since the idiosyncratic volatility risk premium can be positive for some of the stocks, it seems that buyers of individual equity options leave less money on the table than buyers of index options. However, in our model, the stock volatility risk price is a combination of the index volatility risk and the idiosyncratic stock volatility risk premium. We will further examine the time series, cross-sectional and cross-stock properties of this premium in future work. We also would like to investigate the economic implications of changes in the volatility risk premia.

3.5 Pricing Errors

We investigate the fit of model (9) for the panel data of option prices by looking at the pricing errors (predicted price versus observed price) at monthly frequency in-sample and out-of-sample. Figures 14 and 15 show the in-sample and out-of-sample pricing errors for CSCO. Residuals and forecasting errors are small for the all stocks, suggesting that our model fits the option price data reasonably well.

4 Proposed Research

The next steps for this research may be considered in three areas:

1. Estimate model parameters under the objective and the risk-neutral measure jointly. Test the statistical significance of model parameters.
2. Compare our modeling approach with existing univariate models from a trading perspective.
3. Extend our model with a process to describe the dynamics of the volatility risk premium.

4.1 Estimate model parameters under the objective and the risk-neutral measure jointly. Test the statistical significance of model parameters

In future work, one of the problems we would like to address is to jointly estimate the model parameters under the objective measure and the risk-neutral measure from the multivariate data (stocks returns, index returns, equity option prices, index option prices). We plan to develop a stochastic EM-type algorithm to maximize the joint log-likelihood of returns and options. The EM-type algorithm will allow us to obtain standard errors for parameter estimates using the EM iterates. In particular, once we have the standard errors for the price of volatility risk we can test whether this parameter is statistically significant for each stock and the index, respectively. Also, we can test the statistical significance of new model parameters such as the correlation between returns and volatilities.

4.1.1 A stochastic EM algorithm

Let $\Theta = (\Theta_{obj}, \Theta_{neutral})$ be the model parameters, where Θ_{obj} represents the parameters under the objective measure, and $\Theta_{neutral}$ represents the parameters under the risk-neutral measure (i.e. the price of volatility risk). Also, we denote by $\mathbf{Y_T}$ the multivariate data of observed stock and index returns, by $\widehat{\mathbf{O}}_\mathbf{T}$ the observed option prices, and by $\mathbf{X_T}$ the latent log-volatility, from time 1 to T. Then we can write the *complete* joint log-likelihood as:

$$\begin{aligned} \log f((\widehat{\mathbf{O}}_\mathbf{T}, \mathbf{Y_T}), X_0, \mathbf{X_T}|\Theta) &= \log f(\widehat{\mathbf{O}}_\mathbf{T}|\mathbf{Y_T}, X_0, \mathbf{X_T}, \Theta) + \log f(\mathbf{Y_T}, X_0, \mathbf{X_T}|\Theta) \\ &= \log f(\widehat{\mathbf{O}}_\mathbf{T}|\mathbf{Y_T}, X_0, \mathbf{X_T}, \Theta_{obj}, \Theta_{neutral}) + \log f(\mathbf{Y_T}, X_0, \mathbf{X_T}|\Theta_{neutral}) \end{aligned} \tag{10}$$

The stochastic EM algorithm allows us to maximize the joint log-likelihood $\log f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}})|\Theta)$ using an expectation-maximization iterative procedure. This algorithm is particular useful in our case since we cannot directly maximize the log-likelihood $\log f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}})|\Theta)$, but we can maximize a lower bound, which iteratively gets closer to the log-likelihood:

$$\begin{aligned}\log f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}})|\Theta) &= \log \int f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}})|X_0, \mathbf{X}_{\mathbf{T}}, \Theta) f(X_0, \mathbf{X}_{\mathbf{T}}|\Theta) dX_0 d\mathbf{X}_{\mathbf{T}} \\ &\geq \int \log f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}})|X_0, \mathbf{X}_{\mathbf{T}}, \Theta) f(X_0, \mathbf{X}_{\mathbf{T}}|\Theta) dX_0 d\mathbf{X}_{\mathbf{T}}.\end{aligned}$$

By using the log-likelihood function we can actually split the estimation procedure into two steps: first update the objective parameters and then update the risk-neutral parameters.

Specifically, the steps in a stochastic EM will be:

1. Start with $\Theta^{(0)}$.

   For $i \geq 0$ (until convergence), given $\Theta^{(i)}$ perform step 2 and step 3 and obtain $\Theta^{(i+1)}$,

2. (**E-step**) Compute the expectation

$$\begin{aligned}Q(\Theta|\Theta^{(i)}) &= \int \log f((\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}}), \mathbf{X}_{\mathbf{T}}|\Theta) f(\mathbf{X}_{\mathbf{T}}|(\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}}), \Theta^{(i)}) d\mathbf{X}_{\mathbf{T}} \\ &= \underbrace{\int \log f(\widehat{\mathbf{O}}_{\mathbf{T}}|\mathbf{Y}_{\mathbf{T}}, \mathbf{X}_{\mathbf{T}}, \Theta) f(\mathbf{X}_{\mathbf{T}}|(\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}}), \Theta^{(i)}) d\mathbf{X}_{\mathbf{T}}}_{Q_2(\Theta|\Theta^{(i)})} \\ &+ \underbrace{\int \log f(\mathbf{Y}_{\mathbf{T}}, \mathbf{X}_{\mathbf{T}}|\Theta) f(\mathbf{X}_{\mathbf{T}}|(\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}}), \Theta^{(i)}) d\mathbf{X}_{\mathbf{T}}}_{Q_1(\Theta_{obj}|\Theta^{(i)})}\end{aligned} \tag{11}$$

   We can approximate $Q_1(\Theta_{obj}|\Theta^{(i)})$ and $Q_2(\Theta|\Theta^{(i)})$ using MC simulations by

$$Q_1(\Theta_{obj}|\Theta^{(i)}) \approx \frac{1}{n_i} \sum_{j=1}^{n_i} \sum_{t=1}^{T} (\log f(y_t|X_t^{(j)}, \Theta_{obj}) + \log f(X_t^{(j)}|X_{t-1}^{(j)}, \Theta_{obj})) \tag{12}$$

   and

$$Q_2(\Theta|\Theta^{(i)}) \approx \frac{1}{n_i} \sum_{j=1}^{j=n_i} \sum_{t=1}^{T} \log f(\widehat{O}_t|X_t^{(j)}, S_t, \Theta), \tag{13}$$

   where $\{X_t^{(j)}\}$ are simulated from $f(X_t|(\widehat{\mathbf{O}}_{\mathbf{T}}, \mathbf{Y}_{\mathbf{T}}), \Theta^{(i)})$.

3. (**M-step**) Obtain $\Theta^{(i+1)}$ by maximizing $Q_1(\Theta_{obj}|\Theta^{(i)})$ over Θ_{obj}. Using $\Theta_{obj}^{(i+1)}$, update to $\Theta_{neutral}^{(i+1)}$ by maximizing $Q_2((\theta_{obj}^{(i+1)}, \Theta_{neutral})|\Theta^{(i)})$ over $\Theta_{neutral}$.

4.1.2 Test the statistical significance of model parameters

The EM algorithm in Section 4.1.1 allows us to maximize the log-likelihood and thus, to test the statistical significance of model parameters under the objective and the risk neutral measure. Also, we plan to find a procedure that uses the EM iterates in order to estimate the standard errors of parameter estimates, and thus test their statistical significance. In particular, we are interested to test the statistical significance of the price of volatility risk.

4.1.3 Investigate model assumptions

First, we plan to investigate whether the price of volatility risk is constant over groups of options with different strikes and maturities. Our model assumes that the risk-neutral parameters $\{\lambda_k\}$ and λ_I don't depend on the options parameters. However, Jiang (2002) found that the risk-neutral process (represented by the λ parameters in our case) estimated from different segments of the options market is not unique, suggesting a segmented market. Therefore, we would like to estimate the λs from different groups of options and see whether they are significantly different.

Second, we would like to test whether the correlation between returns and volatility is zero. Researchers have found that the correlation between the asset returns and volatility is significant for the S&P 500 index returns (see e.g. Pan(2002), Eraker(2004)). The correlation parameter between returns and volatility was typically found to be negative. This means that a fall in prices usually will be accompanied by an increase in volatility, which is sometimes known as the "leverage effect". Therefore, we plan to add new correlation parameters to our model and test whether they are statistically significant.

4.2 Compare our modeling approach with existing univariate models from a trading perspective

Univariate models for spot price and volatility are currently used to price options in almost all trading applications. Therefore, we would like to examine whether our multivariate model performs better than independent univariate models, for stocks and the S&P500 index, in a trading strategy. Specifically, first, we calibrate our model and separate univariate models for stocks and the index using observations from a training period of a few years. Then, we will compute realized returns from the trading strategy using model price predictions, for the remaining years in the sample. The level of realized returns will be considered as an objective criterion to compare modeling approaches.

4.3 Extend our model with a process to describe the dynamics of the volatility risk premium

Figures 10-11 and 12-13 seem to indicate that the price of volatility risk varies from week to week and from month to month in a mean-reverting way, for each asset. Therefore, we plan to extend our model with a new specification for the volatility risk premium. Specifically, we will add new equations to model (9) to describe the dynamics of $\{\lambda_k\}$ and λ_I as latent mean-reverting processes.

Recent papers capture the term structure of the volatility risk premium by assuming that it is proportional to the volatility level (Chernov and Ghysels (2000), Pan (2002) and Eraker (2004) among others). Given the popularity of this specification for the volatility risk premium, we would like to compare it with a more general specification such as an autonomous mean-reverting process.

5 Figures and tables

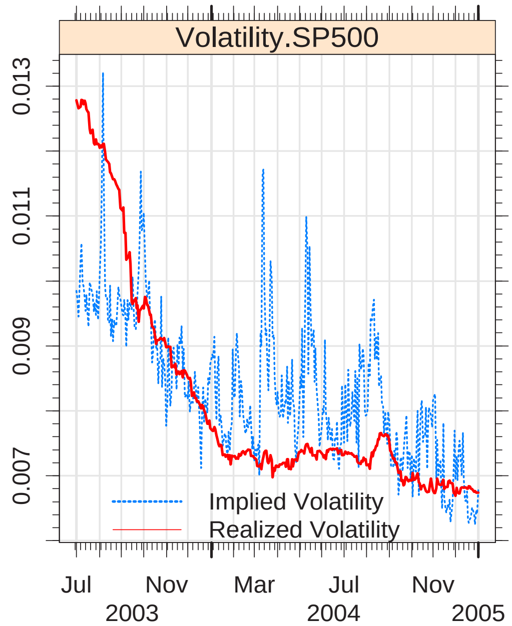

Figure 1: BSM implied volatility and realized volatility for options and returns on S&P500. The realized volatility is computed as the sample variance over an overlapping window of 120 observations.

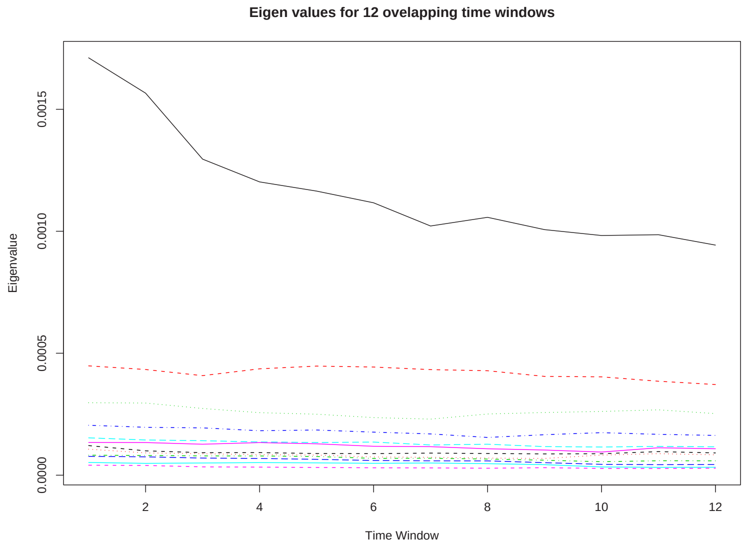

Figure 2: Eigenvalues of the sample covariance matrix of 12 stock returns. Each set of eigenvalues was computed for an overlapping window of 251 observations using data from 2003 and 2004. The stocks are Cisco Systems Inc., DELL, Microsoft, IBM, Citi, JP Morgan, AIG, Bank of America, Newmont Mining, Alcoa, Dow Chemical, and EI DuPont de Nemours.

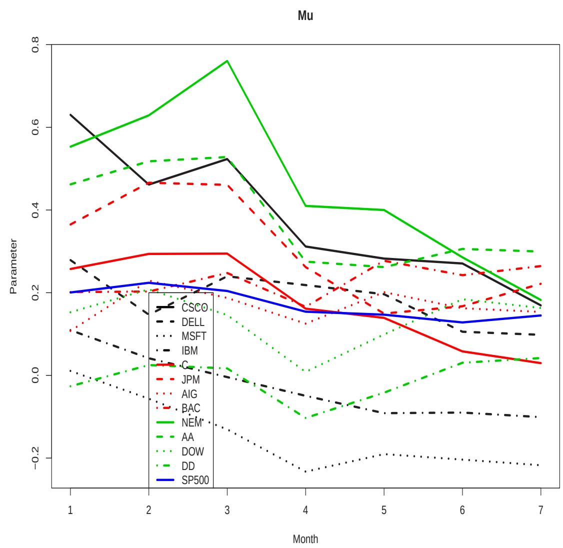

Figure 3: Estimated μ using observations from seven overlapping windows of 252 observations

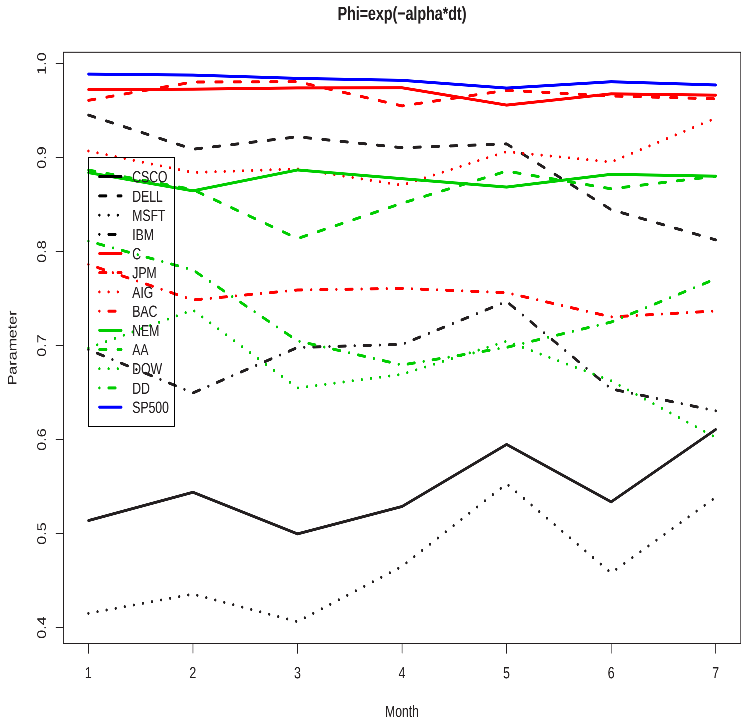

Figure 4: Estimated ϕ using observations from seven overlapping windows of 252 observations

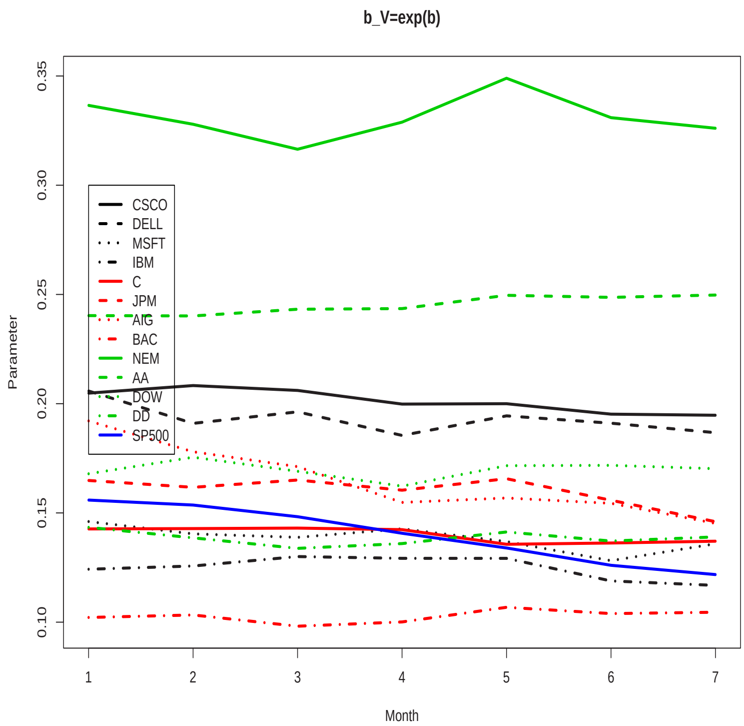

Figure 5: Estimated b using observations from seven overlapping windows of 252 observations

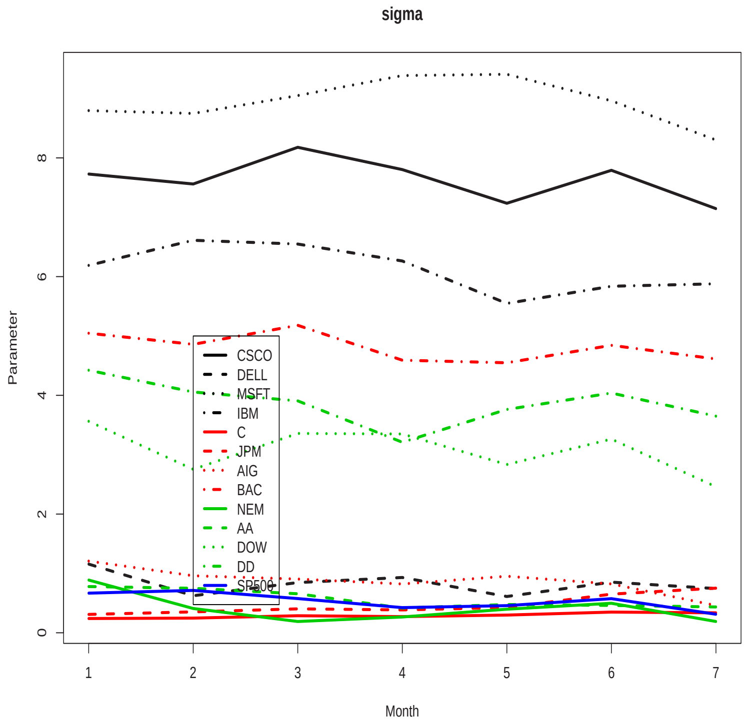

Figure 6: Estimated σ using observations from seven overlapping windows of 252 observations

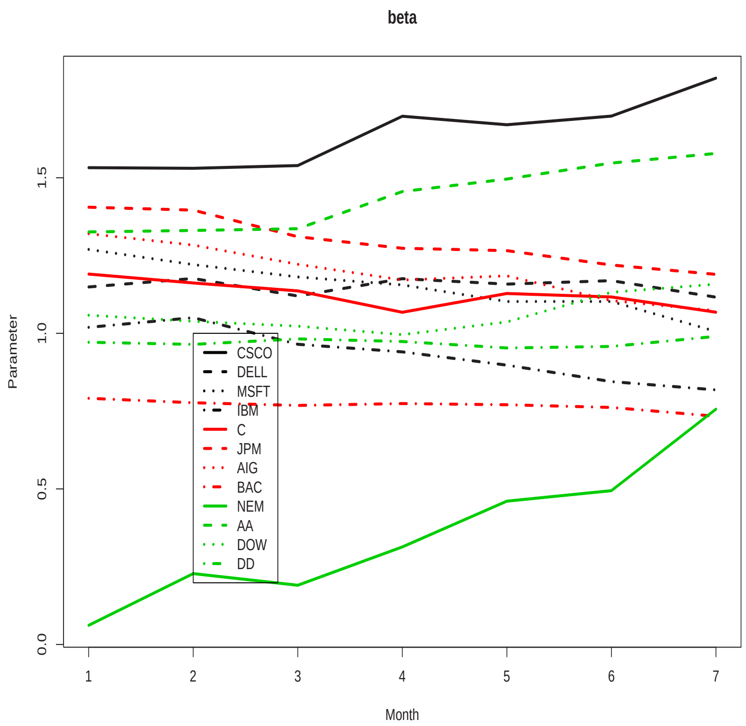

Figure 7: Estimated β using observations from seven overlapping windows of 252 observations

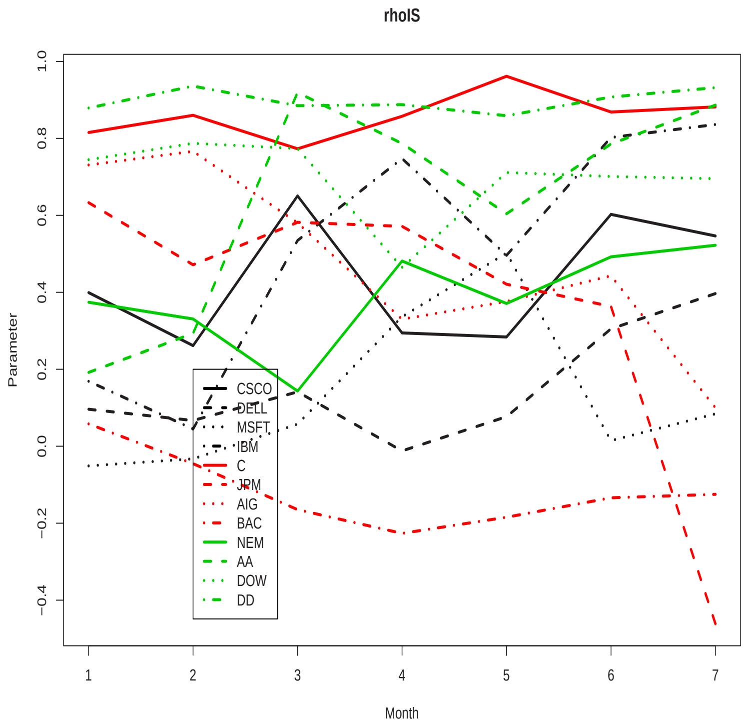

Figure 8: Estimated ρ^{IS} using observations from seven overlapping windows of 252 observations

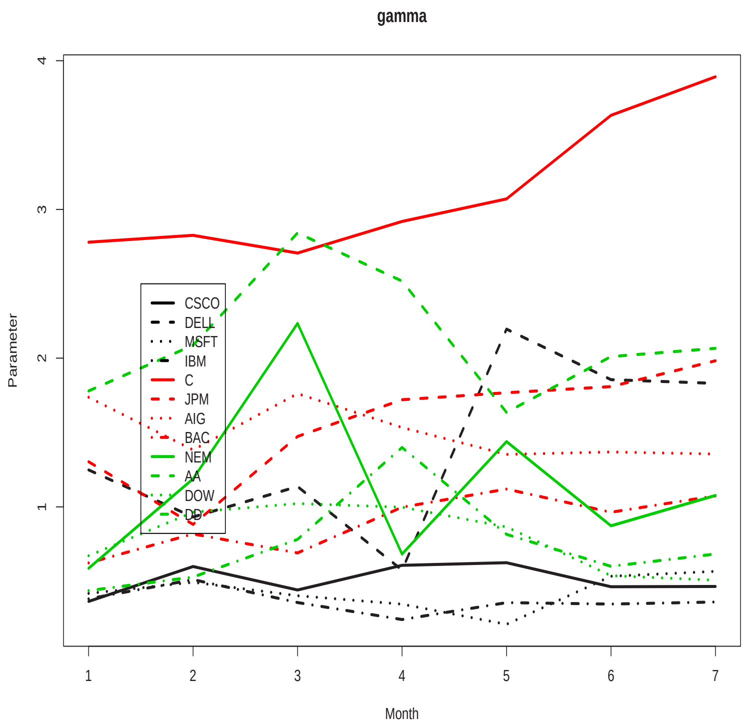

Figure 9: Estimated γ using observations from seven overlapping windows of 252 observations

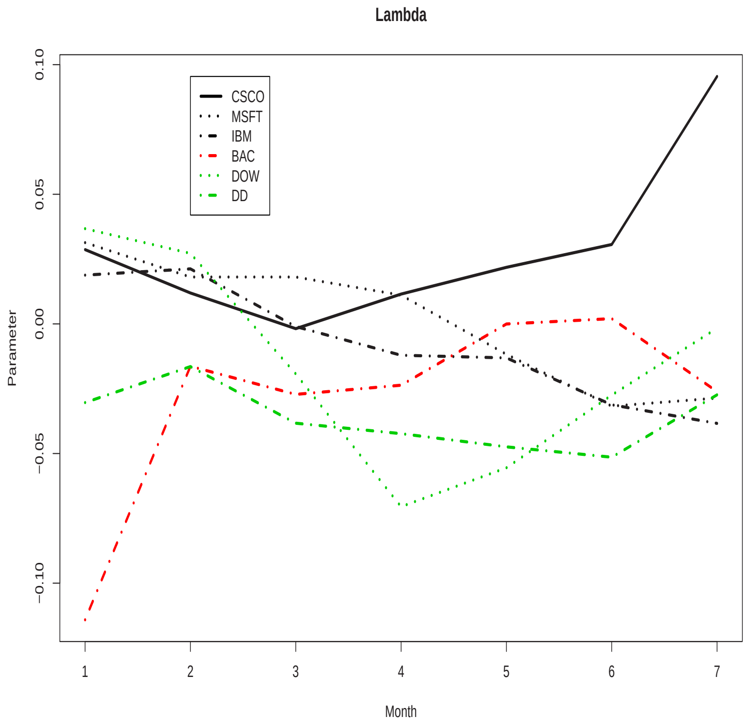

Figure 10: Estimated λ using observations from 7 non-overlapping windows of 20 observations for the assets in the legend

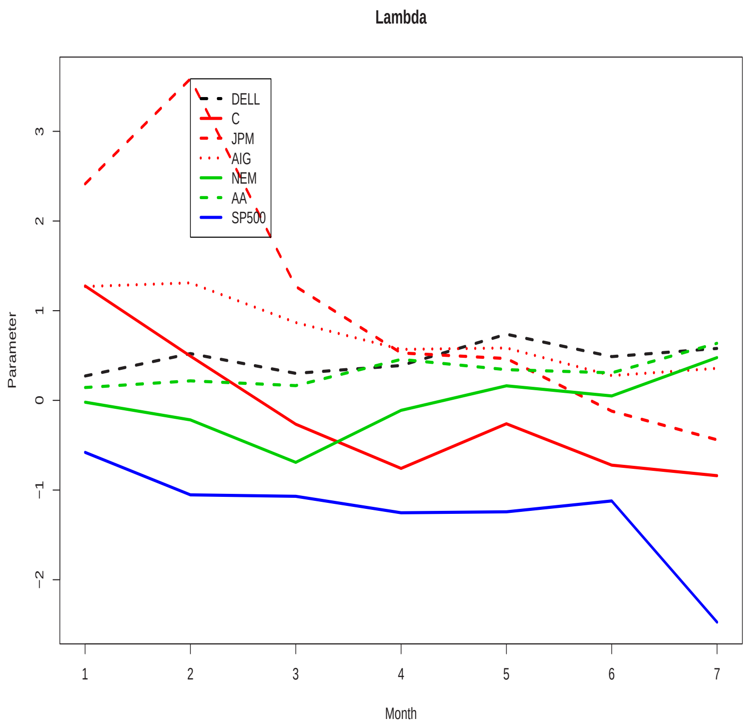

Figure 11: Estimated λ using observations from 7 non-overlapping windows of 20 observations for the stocks in the legend

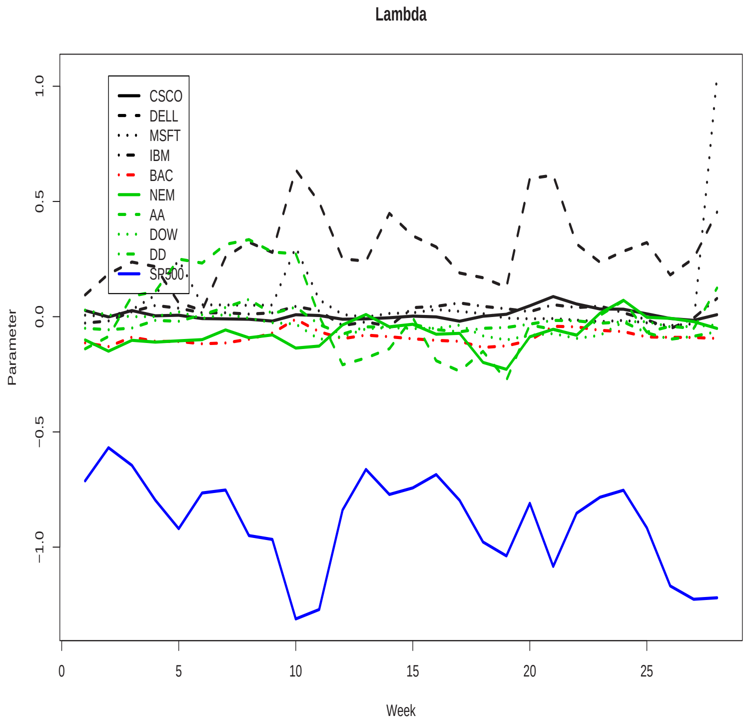

Figure 12: Estimated λ using observations from 28 non-overlapping windows of 5 obs.

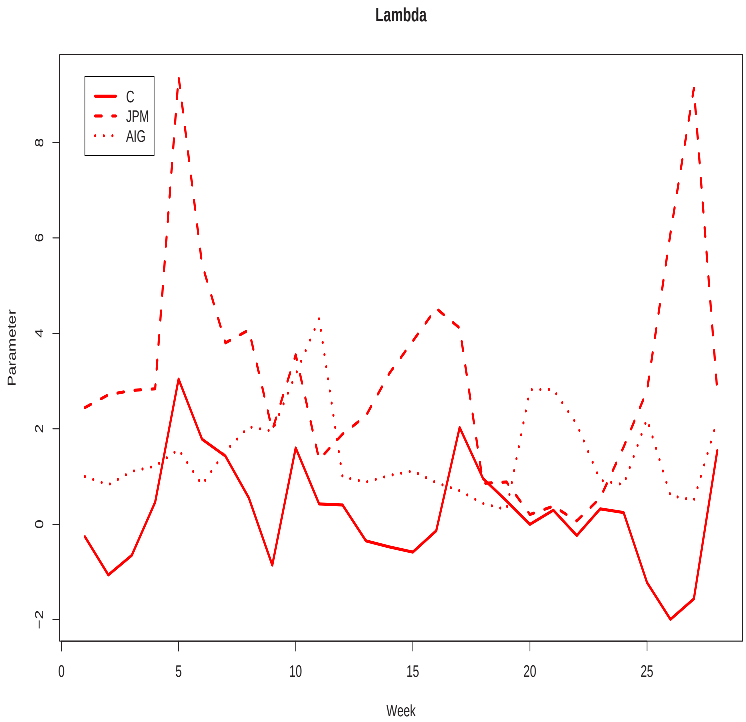

Figure 13: Estimated λ using observations from 28 non-overlapping windows of 5 obs.

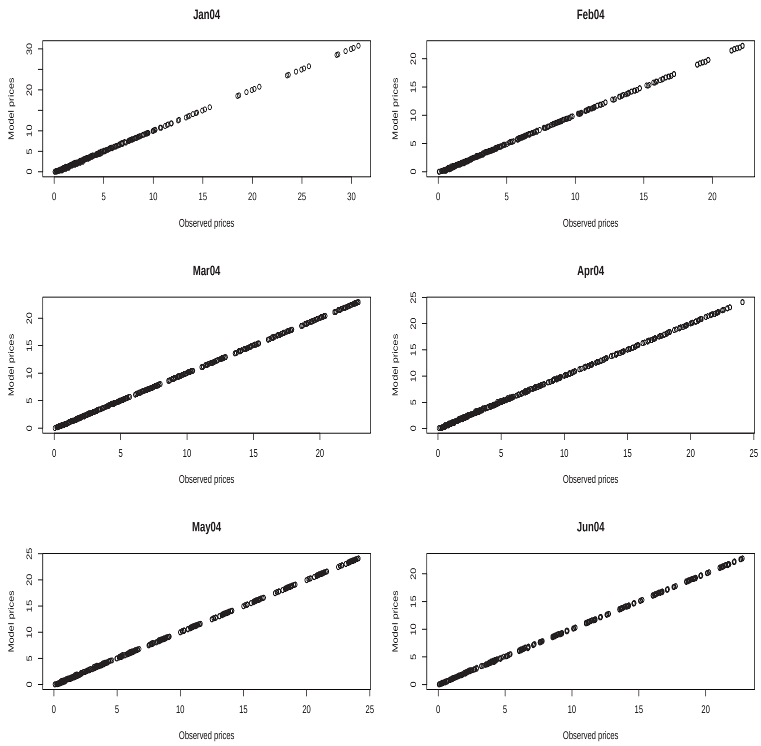

Figure 14: Residuals from option prices for CSCO

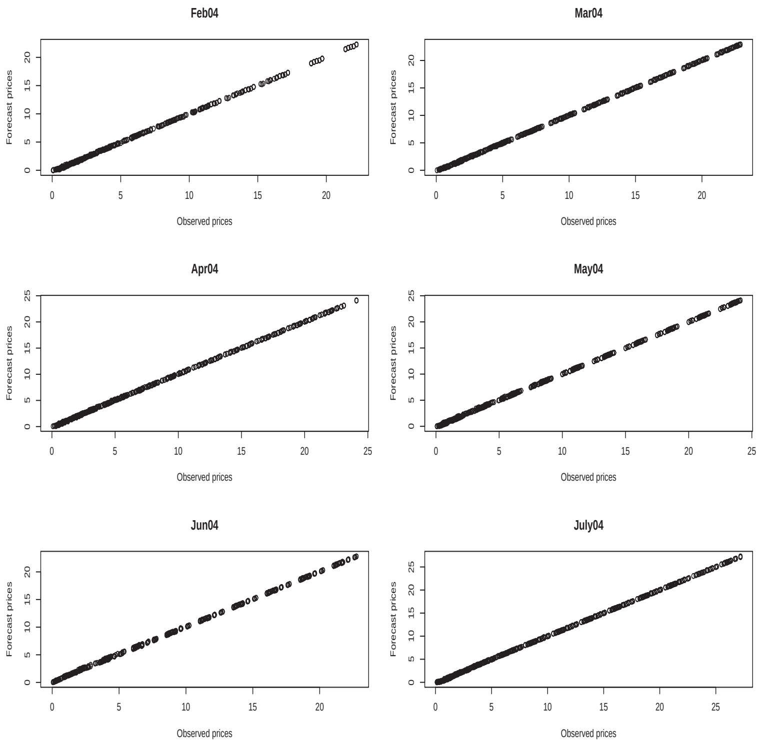

Figure 15: One-month ahead forecasting errors from option prices for CSCO

Ticker	Company	Industry
CSCO	Cisco Systems Inc.	IT
DELL	Dell Inc	IT
MSFT	Microsoft Corp.	IT
IBM	International Business Machines Corp.	IT
C	Citigroup Inc.	Finance
JPM	JP Morgan	Finance
AIG	American International Group	Finance
BAC	Bank of America	Finance
NEM	Newmont Mining	Materials
AA	Alcoa Inc.	Materials
DOW	Dow Chemical	Materials
DD	EI DuPont de Nemours	

Table 1: The 12 analyzed stocks belong to three different industries

	p-value	$\widehat{\mu}$	$\widehat{\phi}$	$\widehat{b_V}$	$\widehat{\sigma}$	$\widehat{\beta}$		$\widehat{\rho}^{IS}$		$\widehat{\gamma}$	
CSCO	0.40	0.56	0.51	0.20	7.65	1.54	**	0.41	*	0.35	*
DELL	0.98	0.28	0.95	0.21	1.16	1.14	**	0.11	*	1.24	*
MSFT	0.78	0.02	0.42	0.15	8.91	1.27	**	-0.04		0.41	*
IBM	0.56	0.10	0.69	0.12	6.17	1.02	**	0.18	*	0.37	*
C	0.89	0.26	0.97	0.14	0.24	1.19	**	0.82	*	2.75	*
JPM	0.49	0.35	0.96	0.16	0.31	1.4	**	0.64	*	1.16	*
AIG	0.92	0.11	0.91	0.19	1.22	1.32	**	0.74	*	1.71	*
BAC	0.70	0.20	0.79	0.10	5.08	0.79	**	0.07	*	0.61	*
NEM	0.85	0.56	0.89	0.34	0.89	0.07		0.39	*	0.57	*
AA	0.88	0.47	0.89	0.24	0.79	1.33	**	0.20	*	1.75	*
DOW	0.96	0.16	0.70	0.17	3.57	1.06	**	0.75	**	0.65	*
DD	0.47	-0.02	0.81	0.14	4.44	0.97	**	0.88	**	0.44	*
SP500	0.98	0.20	0.99	0.16	0.67						

Table 2: Goodness-of-fit tests and estimates from the multivariate stochastic volatility model(year 2003). The "**" and "*" stand for significantly different from zero at level 1% and 5%, respectively. We display parameter estimates in a more meaningful form than in their initial form in (14). That is, we report $\phi = \exp(-\alpha * dt)$ instead of α, which measures the persistence in log-volatility, and we report $b_V = \exp(b)$ instead of b to have a better sense of the long-run mean of the volatility.

APPENDIX

A Discrete-time version of model (8)

$$
\begin{aligned}
Y_k(n) &= \log(S_k(n+1)/S_k(n)) \\
&= (\mu_k - \beta_k\mu_I - V_k(n)/2 - \beta_k(\beta_k - 1)V_I(n)/2)\Delta t + \beta_k M(n) \\
&+ \sqrt{V_k(n)\Delta t} Z_k(n), \\
\\
M(n) &= \log(I(n+1)/I(n)) \\
&= (\mu_I - V_I(n)/2)\Delta t + \sqrt{V_I(n)\Delta t} Z_I(n) \\
\\
V_k(n) &= exp(X_k(n)), \\
X_k(n) &= b_k + \exp(-\alpha_k\Delta t)(X_k(n-1) - b_k) + \sigma_k c_k \epsilon_k(n), \\
\epsilon_k(n) &= \gamma_k \bar{\epsilon}(n) + \eta_k(n), \\
\\
V_I(n) &= exp(X_I(n)), \\
X_I(n) &= b_I + \exp(-\alpha_I\Delta t)(X_I(n-1) - b_I) + \sigma_I c_I \epsilon_I(n),
\end{aligned} \tag{14}
$$

where, for $k = 1, .., N$, $c_k^2 = \frac{1-\exp(-2\alpha_k\Delta t)}{2\alpha_k}$, $c_I^2 = \frac{1-\exp(-2\alpha_I\Delta t)}{2\alpha_I}$, $Z_k(n)$, $Z_I(n)$, $\bar{\epsilon}(n)$, $\eta_k(n)$

and $\epsilon_I(n)$ are standard normal r.v. with the following correlation coefficients

$$\begin{aligned} 0 &= corr(Z_k(n), \epsilon_k(n)), \\ 0 &= corr(Z_I(n), \epsilon_I(n)), \\ \rho_k^{IS} &= corr(\eta_k(n), \epsilon_I(n)), \\ 0 &= corr(\bar{\epsilon}(n), \eta_k(n)), \\ \rho &= corr(\bar{\epsilon}(n), \epsilon_I(n)), \\ 0 &= corr(Z_i(n), Z_I(n)). \end{aligned}$$

References

[1] Alizadeh, S., Brandt, M. and Diebold, F.X. (2002), "Range-Based Estimation of Stochastic Volatility Models.", *Journal of Finance*, Vol 57, pp. 1047-1092.

[2] Bakshi, G. and Kapadia, N. (2002), "Delta Hedged Gains and the Pricing of Volatility Risk", *The Review of Financial Studies*, Vol 16, pp.527-566.

[3] Bakshi, G. and Kapadia, N. (2003), "Volatility Risk Premiums Embedded in Individual Equity Options: Some New Insights", *Journal of Derivatives*, Vol 11, pp.45-54.

[4] Barndorff-Nielsen, O.E., and Shephard, N. (2001a), "Non-Gaussian Ornstein-Uhlenbeck models and some of their uses in financial economics", *Journal of the Royal Statistitcal Society, Series B, Methodological*, Vol 63, pp. 167-241.

[5] Barndorff-Nielsen, O.E., and Shephard, N. (2002a), "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models", *Journal of the Royal Statistitcal Society, Series B, Methodological*, Vol 64, pp. 253-280.

[6] Bates, D.S. (2000), "Post-'87 crash fears in the S&P 500 futures option market, *Journal of Econometrics*, Vol 94, pp. 181-238.

[7] Black, F. and Scholes, M. (1973), "The pricing of options and corporate liabilities", *Journal of Political Economy*, Vol 81, pp. 637-659.

[8] Bollen, N.P.B. and Whaley, R.E. (2004), "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?", *Journal of Finance*, Vol. 59, pp. 711- 753.

[9] Brockwell, A.E. (2005), "Recursive Kernel Density Estimation of the Likelihood for Generalized State-Space Models", Technical Report, Department of Statistics, Carnegie Mellon University.

[10] Chernov, M. and Ghysels, E. (2000), "A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of option valuation", *Journal of Financial Economics*, Vol 56, pp. 407-458.

[11] Comte, F., Coutin, L., and Renault, E. (2003), "Affine fractional stochastic volatility models", Unpublished paper: University of Montreal.

[12] Coval, J. and Shumway, T. (2001), "Expected Option Returns", *Journal of Finance*, Vol 56, pp. 983-1009.

[13] Cox, J.C. and Ross, S.A. (1976), "The valuation of options for alternative stochastic processes", *Journal of Financial Economics*, Vol. 3, pp. 145-166.

[14] Derman, E. and Kani, I. (1994), "Riding on the smile", *Risk*, Vol 7, pp. 32-39.

[15] Diebold, F.X. and Nerlove, M. (1989), "The dynamics of exchange rate volatility: a multivariate latent factor ARCH model", *Journal of Applied Econometrics*, Vol 4, pp. 1-21.

[16] Doran, J.S. and Ronn, E.I. (2004), "On the Market Price of Volatility Risk", Working paper, University of Texas at Austin.

[17] Doucet, A., N. de Freitas, and Gordon, N.J. (2001), "Sequential Monte Carlo Methods in Practice", New York:Springer-Verlag.

[18] Dumas, B., Fleming, J. and Whaley, R.E. (1998), "Implied Volatility Functions: Empirical tests", *Journal of Finance*, Vol 53, pp. 2059-2106.

[19] Dupire, B. (1994), "Pricing with a smile", *Risk*, Vol. 7, pp. 18-20.

[20] Elerian, O., Chib, S., and Shephard, N. (2001), "Likelihood inference for discretely observed non-linear diffusions", *Econometrica*, Vol 69, pp. 959-993.

[21] Engle, R.F. (2002), "Dynamic Conditional Correlation - A simple class of multivariate GARCH", *Journal of Business and Economics Statistics*, Vol 17(5), pp. 425-446.

[22] Emanuel, D.C. and MacBeth, J. (1982), "Further results on the constant elasticity of variance call option pricing model",*Journal of Financial and Quantitative Analysis*, Vol. 4, pp. 533-554.

[23] Eraker, B. (2001), "Markov chain Monte Carlo analysis of diffusion models with application to finance", *Journal of Business and Economic Statistics*, Vol 19, pp. 177-191.

[24] Eraker, B. (2004), "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices", *Journal of Finance*, Vol 59, pp. 1367-1403.

[25] Gallant, A.R. and Tauchen, G. (1996), "Which Moments to Match", *Econometric Theory*, Vol. 12, pp. 657-681.

[26] Gallant, A.R. and Tauchen, G. (1998), "Reprojecting partially observed systems with application to interest rate diffusions", *Journal of the American Statistical Association*, Vol. 93, pp. 10-24.

[27] Gallant, A.R., Hsu, C.T. and Tauchen, G. (1999), "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance", *The Review of Economics and Statistics*, Vol. 81 (4), pp. 617-631.

[28] Gordon, N.J., Salmond, D.J., and Smith, A.F.M. (1993), "A novel approach to nonlinear and non-Gaussian Bayesian state estimation", *IEE-Procedings F*, 140, pp. 107-113.

[29] Gourieroux, C. and Sufana, R. (2004), "Derivative Pricing with Multivariate Stochastic Volatility: Application to Credit Risk", Working paper, University of Toronto.

[30] Hansen, L.P. (1982), “Large Sample Properties of Generalized Method of Moments Estimators”, *Econometrica*, Vol 50, pp. 1029-1054.

[31] Hull, J. and White, A. (1987), “The Pricing of Options on Assets with Stochastic Volatility”, *Journal of Finance*, Vol 42, pp. 281-300.

[32] Heston, S.L. (1993), “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Optitons”, *The Review of Financial Studies*, Vol 6, number 2, pp. 327-343.

[33] Jacquier, E., Polson, N.G., and Rossi, P.E. (1994), “Bayesian analysis of stochastic volatility models (with discussion)”, *Journal of Business and Economic Statistics*, Vol 12, pp. 371-417.

[34] Jiang, G.J. (2002), “Testing Option Pricing Models with SV, Random Jumpa and Stochastic Interest rate”, *International Review of Finance*, pp. 233-272.

[35] Jones, C.S. (2003), “The dynamics of stochastic volatility: evidence from underlying and option markets”, *Journal of Econometrics*, Vol 116, pp. 181-224.

[36] Longstaff, F.A. and Schwartz, E.S. (2001), “Valuing American Options by Simulation: A Simple Least-Squares Approach”, *The Review of Financial Studies*, Vol. 14, number 3, pp. 113-147.

[37] Melenberg, B. and Werker, J.M. (2001), “The implicit Price of Volatility Risk: An Empirical Analysis”, Working paper, Tilburg University.

[38] Merton, R.C. (1973), “The Theory of Rational Option Pricing”, *Bell Journal of Economics and Management Science*, 4, pp. 141-183.

[39] Pan, J. (2002), “The jump-risk premia implicit in options: evidence from an integrated time-series study”, *Journal of Financial Economics*, Vol 63, pp. 3-50.

[40] Pitt, M.K. and Shephard, N. (1999), “Filtering via simulation: auxiliary particle filter”, *Journal of the American Statistical Association*, Vol 94, pp. 590-599.

[41] Rambharat, B.R. (2005), “Valuation Methods for American Derivatives in a Stochastic Volatiltity Framework”, Ph.D. Thesis, Department of Statistics, Carnegie Mellon University.

[42] Rubinstein, M. (1994), “Implied binomial trees”, *Journal of Finance*, Vol 49, pp. 771-818.

[43] Şerban, M., Brockwell, A., Lehoczky, J. and Srivastava, S. (2006), “Modeling the Dynamic Dependence Structure in Multivariate Financial Time Series”, submitted to the *Journal of Time Series Analysis*.

[44] Tauchen, G. (2004), “Remarks on Recent Developments in Stochastic Volatility: Statistical Modelling and General Equilibrium Analysis”, Working paper, Duke University.