Pricing Options under Stochastic Volatility: An Empirical Investigation*

Luca Benzoni†

December 2001

*I am grateful for the comments and suggestions of Marco Cagetti, Tim Conley, Kent Daniel, Eric Ghysels, Jack Kareken, Bob Korajczyk, Debbie Lucas, Jesper Lund, Guido Rossi, Costis Skiadas, Harold Zhang, seminar participants at the University of British Columbia, Columbia University, INSEAD, Loyola University, McGill University, the University of Minnesota, MIT, Northwestern University, Stanford University, the University of Utah, Waterloo University, the Econometric Society Winter Meeting 1999 and, especially, Torben Andersen, John Heaton, Erzo Luttmer, Bob McDonald and Ravi Jagannathan. Of course, all errors remain my sole responsibility. I am indebted to the Chicago Board Option Exchange for providing S&P 500 option price data. Previous versions of this paper were circulated under the title "Pricing Options under Stochastic Volatility: An Econometric Analysis," 1998.

†Finance Department, 3-122 Carlson School of Management, 321 19th Ave. South, Minneapolis, MN 55455, Tel. 612-624-1075, Fax 612-626-1335, e-mail lbenzoni@umn.edu.

Pricing Options under Stochastic Volatility: An Empirical Investigation

Abstract

In this paper, I investigate the empirical properties of two stochastic volatility diffusion models using data on both S&P 500 returns and option prices. I find that a negative relationship between returns and volatility innovations is critical for fitting the pronounced asymmetry in the patterns of options' implied volatilities and the skewness in index returns. Variance risk is priced in the S&P500 option market. I find a risk premium that is not only statistically significant, but also has an economic impact on option prices. A non-zero premium reduces pricing errors considerably. Also, it fattens the tails of the state-price densities (SPDs) and changes the implied volatilities obtained from option prices using either of the two stochastic volatility models. The performance of the option pricing model is very sensitive to the underlying returns variance. I compare return- and option-based estimates of volatility and find that the latter provide the best results in reducing option pricing errors. Moreover, considerable discrepancies are found between return- and option-based variance estimates. This result is indicative of the presence of inconsistencies in the joint model for option prices and index returns. Finally, the analysis of option pricing errors, SPDs and other model diagnostics indicate that the two stochastic volatlity diffusions considered in this paper perform similarly.

1 Introduction

Even though routinely used to price European options, the Black-Scholes model is well known for its pricing biases. Practitioners are familiar with the "volatility smile." Before the 1987 crash, equity options that were deeply in- and out-of-the-money had higher implied volatilities than at-the-money contracts. This phenomenon became accentuated after 1987. Since the crash, out-of-the-money puts and in-the-money calls have had higher implied volatilities than have other contracts, and the "smile" has assumed the asymmetric shape of a "smirk"; see, e.g., Rubinstein (1994).

The qualitative pattern observed in implied volatilities can be reconciled with the empirical evidence on stock returns. Equity returns exhibit excess skewness, leptokurtosis and pronounced conditional heteroskedasticity, all characteristics at odds with the assumptions underlying the Black-Scholes model; see, e.g., Bollerslev, Engle and Nelson (1994) and references therein. The presence of outliers, captured by the excess of kurtosis and conditional heteroskedasticity, and the asymmetry in the returns distribution is qualitatively consistent with the higher prices of out-of-the-money options and the pronounced asymmetry of the volatility smirk. Consequently, different extensions of the Black-Scholes model have been suggested to account for the salient features of the data. In particular, a number of studies have investigated the continuous-time stochastic volatility model. This specification avoids many of the shortcomings of the constant variance diffusion assumed by Black and Scholes, and may still be cast within the class of representative agent models which allow for derivative pricing via equilibrium arguments; see, e.g., Bakshi, Cao and Chen (1997, 2000), Bates (1996a,b, 2000), Chernov and Ghysels (2000), Chernov et al. (1999, 2000), Eraker (2000), Eraker, Johannes and Polson (2001), Pan (2001), and Jones (1999).

In this paper, I present estimates of two common stochastic volatility diffusion models, one with a log-variance specification (see, e.g., Melino and Turnbull (1990)) and the other with a square-root specification (see, e.g., Heston (1993)), and compare them based on goodness-of-fit criteria for both S&P 500 returns and derivative prices. In my application, I use a two-stage estimation procedure. In the first stage, I use a time series of daily S&P 500 returns and a simulated method of moments (SMM) procedure (see, e.g., Duffie and Singleton (1993)) to estimate the structural parameters of the model. More specifically, I obtain moment conditions from an implementation of the efficient method of moments (EMM) procedure of Gallant and Tauchen (1996), which provides a convenient setting for consistent continuous-time estimation and analysis of the model specifications. In the second stage, I use a simulation methodology and a sample of S&P 500 option prices to estimate the risk adjustment which is necessary for derivative pricing. This procedure relies on the first-stage EMM estimates of the stochastic volatility parameters.

The presence of stochastic volatility makes it natural to ask whether there is a premium for volatility risk embedded in observed option prices. There is at present no consensus about the magnitude of the premium, even though there have been several attempts to measure it. One of the contribution of this paper is to develop an estimation procedure, also described in detail below, for estimating that premium using S&P 500 option prices and returns. Another is the development of the asymptotic properties of the two-stage estimator, which are used to assess the statistical significance of the volatility risk premium. Moreover, I develop diagnostics which make possible an understanding of the economic importance of the estimated premium. Finally, I provide a number of goodness-of-fit criteria that allow me to compare and assess the (mis)pricing of the two diffusion-based specifications.

In my analysis, I investigate continuous-time stochastic volatility models within and outside the affine class. However, I do not incorporate jumps in my specifications. There is a mounting evidence in support of the presence of jumps in S&P 500 returns and possibly in their volatility; see, e.g., Andersen, Benzoni and Lund (2001) and Eraker, Johannes and Polson (2001). However, a number of contributions have pointed out that jumps are only of second-order importance in fitting option prices, in comparison with stochastic volatility, except possibly for the shortest lived options; see, e.g., Bakshi, Cao and Chen (1997) and Eraker (2000). Therefore, since the focus of this paper is on option pricing, I decide to concentrate on pure stochastic volatility models. More specifically, I investigate the economic and statistical implications of such specifications. In the present context, adding jumps would make the empirical analysis considerably more difficult and possibly complicate the interpretation of results.

I deliberately choose a two-stage estimation procedure. Even though it is not fully efficient, there are, however, several advantages in using it. First, my two-step approach allows me to combine a very large sample of index returns with a relatively small sample of option prices. Trading of derivative contracts started only relatively recently, and an estimation procedure which uses simultaneously the two data sets would limit the length of the underlying returns sample used in the estimation. This would be a considerable limitation, because a large sample of returns is crucial for precise estimation of the models. In particular, it is important for capturing the dynamics of the strongly persistent volatility process and for pinning down the coefficient of correlation between shocks to volatility and shocks to returns. As is discussed below, this model parameter turns out to be critical for fitting important characteristics of the S&P 500 returns and option prices. Second, a joint estimation using two data sets together would involve additional computational problems and, as a practical matter, limit analysis to affine model specifications, which deliver closed-form option pricing formulas. In contrast, my two-step estimation methodology allows me to also investigate the option pricing properties of the non-affine log-variance model and compare them to those of the more commonly used square-root specification. The procedure provides a natural setting for the application of the EMM technique, which delivers efficient and consistent estimates for the continuous-time stochastic volatility models, and offers powerful model diagnostics and specification tests. Third, my two-stage approach delivers diagnostics for the models estimated directly under the "physical" probability measure, in addition to characterizing the "risk neutral" dynamics. It is therefore possible to identify the presence of potential specification problems in the underlying return dynamics and disentangle them from other issues arising from the characterization of the model risk premia.

I find that the log-variance stochastic volatility model provides a much better fit for S&P 500 returns and option prices than that produced by standard one-factor diffusion models. More specifically, the correlation between shocks to volatility and shocks to the underlying returns captures nicely the negative skewness observed in returns. The results of the EMM estimation are indicative of the presence of strong asymmetries, and the analysis of option prices confirms that evidence.

Variance risk is priced in the S&P 500 option market. I find a risk premium that is not only statistically significant, but also has an economic impact on option prices. A non-zero premium reduces pricing errors considerably. Such a premium also fattens the tails of the state-price densities (SPDs) and changes the implied volatilities obtained from option prices using the stochastic volatility model ("SV implied volatilities"). More specifically, when the premium is set equal to zero, SV implied volatilities increase to compensate for the unaccounted volatility risk.

I find that the performance of the option pricing model is very sensitive to the underlying returns variance. I compute estimates of volatility from index returns using the Kalman filter and the reprojection method of Gallant and Tauchen (1998), and from derivative prices by minimizing the deviations between market and stochastic volatility option prices. In-sample, return-based estimates of variance do not provide nearly as good results as the SV implied volatilities evaluated from same-day option prices, which in comparison with Black-Scholes reduce the squared pricing errors by a factor of 2-3. Similar conclusions are obtained from the out-of-sample analysis. In this case, day-before SV implied volatilities are used to evaluate derivative prices; this reduces the squared pricing errors by a factor of 2 in comparison with Black-Scholes. The Heston model performs similarly; the square-root specification delivers EMM diagnostics which are in line with the log-variance representation, and produces comparable in- and out-of-sample option pricing errors.

A direct examination of the SV implied volatilities across different levels of option "moneyness" suggests that the incorporation of a stochastic volatility factor lessens but does not eliminate the volatility smile. Other diagnostics reveal the existence of considerable discrepancies between the time series of return-based estimates of variance and SV implied volatilities, and those discrepancies are indicative of inconsistencies in the stochastic volatility model for both option prices and stock returns.

Finally, the analysis of SPDs at different points in time is used to highlight the main characteristics of stochastic volatility option prices. SPDs exhibit a fatter left tail at longer maturities, compared to Black-Scholes. On the other hand, they do not show signs of asymmetry and leptokurtosis close to expiration. This evidence suggests that further extensions of the model, e.g., allowing for jumps, may improve the pricing of very the shortest term options.

The remainder of the paper is organized as follows. In Section 2, I set out the stochastic volatility specifications to be analyzed, and in Section 3 I discuss the estimation methodology, stressing how it contrasts with other procedures in the relevant literature. In Section 4, I report the results of the EMM estimation of the stochastic volatility models and the associated specification tests. Then, in Section 5, I set up the derivative pricing model and discuss the strategy for estimating option prices and the premium for variance risk. The asymptotic properties of the estimator are developed in Section 6, and in Section 7 I report the empirical option pricing results. In Section 8, I offer some concluding observations.

2 Model Specification

A common extension of the Black-Scholes model incorporates a stochastic volatility factor:

$$\frac{dS_t}{S_t} = (\mu + c\,V_t\,)\,dt + \sqrt{V}_t\,dW_{1,t}\,, \tag{1}$$

where the (log-)variance process is assumed to exhibit mean-reversion, as in

$$d\ln V_t = (\alpha - \beta \ln V_t\,)dt + \eta\,dW_{2,t}\,, \tag{2}$$

or

$$dV_t = (\alpha - \beta V_t\,)dt + \eta\,\sqrt{V}_t\,dW_{2,t}\,, \tag{3}$$

with W_1 and W_2 standard Brownian motions. Stochastic volatility induces excess kurtosis in the return process. Its level is largely determined by the volatility parameters α, β and η. Moreover, the model allows for correlation between shocks to volatility and shocks to index returns induced by the coefficient $\rho = \mathrm{corr}(dW_{1,t}, dW_{2,t})$. This makes it possible to capture the negative skewness (asymmetry) observed in equity returns. In the rest of the paper, I refer to the special case of $\rho = 0$ as the "symmetric" stochastic volatility model and to the general case of $\rho \neq 0$ as the "asymmetric" model. Finally, I allow the volatility factor to enter the mean coefficient, and therefore rule out arbitrage opportunities when the variance takes a (near) zero value.[1]

The square-root specification (3) provides a tractable setting for derivative pricing: Heston (1993) obtains a closed-form solution for the option premium when the underlying returns obey (1) and (3). The log-variance specification (2) is inspired by the EGARCH model of Nelson (1991), which has been used successfully for fitting equity returns and is more in line with discrete-time stochastic volatility models that have been studied. Also, the model (1) and (2) is easily converted to state-space form, which, as illustrated in Appendix A, delivers a simple estimate of the unobservable variance process. (Alternative estimates of the variance process based on both equity returns and option prices are discussed below.) Nevertheless, no closed-form solution is available for model (1) and (2), so that numerical methods must be used to obtain option prices; see, e.g., Melino and Turnbull (1990). It is strictly an empirical issue to determine which one of the two models, (1) and (2) or (1) and (3), provides the better fit for index returns and option prices. It is therefore important to rely on an estimation procedure which can be used for both specifications and delivers the diagnostics necessary to assess the performances of the two models.

3 Estimation Methodology

When equity returns are described by a continuous-time model with latent variables, a closed-form expression for the discrete-time transition density of the process is generally not available, and standard estimation techniques, such as maximum likelihood (ML), cannot be used. In some cases, an approximate likelihood function can be computed by first discretizing the model. But there is a risk of inconsistent estimates; see, e.g., Lo (1988). Even when ML is in principle feasible, empirical applications are computationally challenging if the latent variables have to be integrated out of the likelihood function—as is typically necessary in the presence of stochastic volatility.

Alternative techniques for the estimation of continuous-time models have been developed in recent years: among others, those of Aït-Sahalia (1996), Bandi and Phillips (1998),

[1] This specification prevents the occurrence of a positive risk premium when the variance is zero and, with $\mu = r - d$, yields the risk-neutral dynamics in (22). The c coefficient captures the volatility-in-mean effect. But the existence of such effect is problematic (see Glosten, Jagannathan and Runkle (1993), Nelson (1991) and references therein), so that it should not be surprising if this coefficient is found to be statistically insignificant.

Bandi and Nguyen (2000), Conley at. al. (1997), Hansen (1995), Jiang and Knight (1997), Johannes (1999), and Stanton (1997), who suggest a number of (semi-)nonparametric procedures. Unfortunately, they are all difficult to apply in the presence of unobservable factors such as, for example, stochastic volatility. Simulation-based procedures are computationally more intensive, but offer more flexibility. As an example, there is the Monte Carlo Markov Chain method used by Elerian, Chib and Shephard (1998), Eraker (2001), Jacquier, Polson and Rossi (1994), Jones (1998) and Kim, Shephard and Chib (1998). More recent work along these lines is also found in Johannes, Kumar and Polson (1998) and Eraker, Johannes and Polson (2001), who explore a pure jump model and a square-root stochastic volatility jump-diffusion. The advances in Duffie, Pan and Singleton (2000) have inspired new methods based on empirical characteristic functions; see, e.g., Singleton (2001), Chacko and Viceira (1999), Jiang and Knight (1999) and Carrasco et al. (2001). However, these new approaches cannot be easily extended to the log-variance specification.

In my application, I use the EMM procedure of Gallant and Tauchen (1996), a simulated method of moments technique. This approach offers the flexibility necessary for estimating non-affine models with latent variable, of which model (1) and (2) is an example. The SMM procedure of Duffie and Singleton (1993) matches sample population moments with simulated moments. The EMM technique refines this approach by using different moment conditions, obtained from the expectation of the score of an auxiliary model which closely approximates the distribution of the data. One of the primary advantages of the technique is that EMM estimates achieve the same degree of efficiency as the ML procedure when the score of the auxiliary model asymptotically spans the score of the true model. It also delivers powerful specification diagnostics that provide guidance in the model selection. Andersen, Benzoni and Lund (2001) also conduct an EMM application that relies on a sample of S&P 500 index returns. However, this paper differs significantly from that of Andersen et al.. In Andersen et al. the emphasis is on fitting the underlying return dynamics, while in this paper the focus is on the option pricing implications of the model. Other EMM applications can be found in, e.g., Chernov and Ghysels (2000), and Chernov et. al. (1999, 2000).[2]

In most of the studies mentioned so far, only samples of returns are used, and models are estimated under the "physical" probability measure. (A notable exception is Chernov

[2]More work on continuous- and discrete-time estimation of stochastic volatility models using asset returns and method-of-moments type procedures can be found in, among other papers, Andersen and Lund (1996, 1997), Bollerslev and Zhou (2001), Gallant, Hsu and Tauchen (1999), Ho, Perraudin and Sørensen (1996), Jiang and van der Sluis (1999), Liu and Zhang (1997), Meddhai (2001), Melino and Turnbull (1990), Pastorello, Renault and Touzi (1994), and van der Sluis (1997). Other estimation techniques for stochastic volatility models are surveyed by Ghysels, Harvey, and Renault (1998).

and Ghysels (2000), which is discussed below in more detail.) But there are a number of authors who advocated the use of option prices, and as a result estimated the model under the "risk-neutral" probability measure. Bates (1996a, 2000) extends the model (1) and (3) by allowing for (multiple) stochastic volatility factors and jumps in the return process. He uses a sample of option prices only to simultaneously estimate the model parameters and the variance process V by minimizing the sum of squared residuals (SSR) between market and stochastic volatility prices. Bakshi, Cao and Chen (1997) use a square-root specification with stochastic volatility, jumps and stochastic interest rates and provide extensive comparisons of the full model, simpler special cases and the benchmark Black-Scholes model. The estimation is based on the minimization of the SSR and makes use of only a sample of option prices.

The mounting evidence in support of stochastic volatility has generated interest in the possible existence of a volatility risk premium. Buraschi and Jackwert (2001) reject the pricing restrictions of a deterministic volatility specification and conclude that models need to incorporate priced risk factors to price and hedge derivatives. Bakshi and Kapadia (2001) and Coval and Schumway (2001) investigate the returns on a number of derivative investment strategies and provide indirect evidence in support of a negative volatility risk premium. Jones (2001) conducts a factor analysis of option expected returns, and concludes that factors other than the return on the underlying security contribute to these expected returns, although factor-based models appear to be insufficient to explain their magnitude. The implications for optimal investment decisions in the presence of volatility and jump risks are studied in, e.g., Liu and Pan (2001).

There are several recent contributions in which both primitives and derivatives data are used to make inference about option prices. This approach has the advantage of delivering direct estimates of the model risk premia, the premium for variance risk included. Chernov and Ghysels (2000) fit the Heston model using EMM and both equity-index returns and Black-Scholes implied volatilities. They find that the best pricing and hedging results are achieved when the model is estimated using a univariate approach and just S&P 500 option prices, rather than a multivariate approach and both primitives and derivatives data. Interestingly, they do not find the correlation between shocks to volatility and equity returns to be important in describing S&P 500 returns and option prices. Pan (2001) estimates a square-root stochastic volatility jump-diffusion process similar to the one found in Bates (2000) using GMM and both S&P 500 returns and option prices. She emphasizes the importance of jump risk premia in reconciling the dynamics implied by the two data samples and in explaining the asymmetry in option prices. Eraker (2000) uses the MCMC approach and both S&P 500 returns and option prices. He finds that jumps to returns and volatility are an important element of the underlying dynamics. Jones

(2000) exploits the VIX implied volatility index and daily S&P 100 returns to estimate a stochastic volatility diffusion with constant elasticity of variance (CEV) that extends the Heston (1993) representation. He concludes that the CEV extension generates more realistic crash probabilities and values of skewness and kurtosis much more consistent with their sample values than the square-root specification does. Jiang and van der Sluis (1999) fit a discrete-time log-variance model with a stochastic interest rate using EMM and a sample of interest rates and equity returns data. They perform an out-of-sample analysis estimating the premium for variance risk with day-before option prices.[3]

4 EMM Estimation of the Stochastic Volatility Model for S&P 500 Returns

In this section, I report on the estimation of the parameters in (1) and (2) and (1) and (3), and provide a comparison of the two models using a goodness-of-fit criterion based on S&P 500 daily returns. In Section 4.1, I discuss the selection and the quasi-maximum likelihood estimation of a semi-nonparametric (SNP) model for the S&P 500 daily returns. Moment conditions are then obtained from the expectation of the SNP scores and are used for the continuous-time EMM estimation, as explained in Section 4.2.

4.1 The SNP Auxiliary Model

The key to a successful application of EMM is the choice of an auxiliary model which closely approximates the conditional returns distribution. Gallant and Long (1997) show that when the score function of the auxiliary model asymptotically spans the score of the true model, EMM is asymptotically efficient. They also show that SNP densities are a good choice for this task.

Quasi-maximum likelihood estimation is performed on the SNP family of conditional densities

$$f_K(r_t|x_t;\xi) = \left(\nu + (1-\nu\,) \times \frac{[P_K(z_t,x_t)]^2}{\int_{\mathrm{R}}[P_K(z_t,x_t)]^2\phi(u)du}\right)\frac{\phi(z_t)}{\sqrt{h_t}}\,,$$

where ν is a small constant (fixed at 0.01), $\phi(.)$ is the standard normal density, r_t is the time-t index return, x_t is a vector of lagged return observations,

$$z_t \;=\; \frac{r_t-\mu_t}{\sqrt{h_t}}\,,$$

[3] Investigations of the volatility risk premium can also be found in, e.g., Guo (1998), Melino and Turnbull (1990), Poteshman (1998), and Kapadia (1998). Related work can be found in the empirical literature on the bias in volatility forecasts computed using option prices; see, e.g., Chernov (2000), Christensen and Prabhala (1998), Day and Lewis (1992), Lamoureux and Lastrapes (1993), Poteshman (2000) and Santa-Clara and Yan (2001).

$$
\begin{aligned}
\mu_t &= \phi_0 + c\,h_t + \sum_{i=1}^{s} \phi_i r_{t-i} + \sum_{i=1}^{u} \delta_i \varepsilon_{t-i}\,, \\
\ln h_t &= \omega + \sum_{i=1}^{p} \beta_i \ln h_{t-i} + (1 + \alpha_1 L + \ldots + \alpha_q L^q)\,[\,\theta_1 z_{t-1} + \theta_2\,(b(z_{t-1}) - \sqrt{2/\pi})\,]\,, \\
b(z) &= |z| \;\text{ for } |z| \geq \pi/2K, \quad b(z) = (\pi/2 - \cos(Kz))/K \;\text{ for } |z| < \pi/2K\,, \\
P_K(z,x) &= \sum_{i=0}^{K_z} a_i(x) z^i = \sum_{i=0}^{K_z} \left(\sum_{|j|=0}^{K_x} a_{ij} x^j \right) z^i\,, \qquad a_{00} = 1\,,
\end{aligned}
$$

j is a multi-index vector, $x^j \equiv (x_1^{j_1}, \ldots, x_M^{j_M})$ and $|j\,| \equiv \sum_{m=1}^{M} j_m$. As in Andersen and Lund (1997), $b(z)$ is a smooth (twice-differentiable) function, with $K = 100$, that closely approximates the absolute value operator in the EGARCH variance equation.

An ARMA term, extended with volatility-in-mean effects, is the natural specification for the conditional mean. Gallant and Long (1997) show that some non-Markovian score generators are valid auxiliary models. This makes the EGARCH specification of Nelson (1991) an ideal choice for the conditional volatility. Besides being more parsimonious than ARCH forms, it readily accommodates an asymmetric response of returns innovations to volatility shocks and makes non-negativity constraints on the volatility parameters unnecessary. The mixture appearing in the conditional density $f_K(r_t|x_t;\xi)$ is introduced to avoid stability problems in the EMM estimation.[4] Within this SNP family, the main task of the semi-nonparametric polynomial expansion in the conditional density is to capture any excess kurtosis in the return process and, to a lesser extent, any asymmetry not accommodated by the EGARCH leading term. It therefore provides a parsimonious and yet accurate representation of the return process. In practice, the polynomial representation P_K is given by Hermite orthogonal polynomials. I allow for heterogeneity in the polynomial expansion ($K_x > 0$); but when the model is confronted with the data, these terms are insignificant, which suggests that the EGARCH leading term suffices to capture this source of heterogeneity.

My application makes use of daily S&P 500 returns from 1/3/1953 to 12/31/1996, a sample of 11,076 observations. Descriptive statistics for the data are provided in Table 1. According to the results of a Dickey-Fuller non-stationarity test, also provided in Table 1, the unit-root hypothesis is convincingly rejected for the returns sequence. A time series plot of the price and return process is provided in Figure 1.

S&P 500 returns exhibit autocorrelation that cannot be readily incorporated in the standard diffusions of this paper. Also, it is likely that a significant portion of this autocorrelation may be spurious, induced by non-synchronous trading in the stocks of the

[4] $P_K(z, x_t)$ may equal zero for a given simulated trajectory, and a zero value causes numerical problems when evaluating the score function. This was pointed out by Qiang Dai of NYU.

index; see, e.g., Lo and MacKinlay (1990). Such autocorrelation is thus of questionable economic significance. Finally, and most importantly, it is of much less importance for option pricing than are volatility fluctuations. The inference for the variance process is largely unaffected by the specification of the mean dynamics—a result documented in Andersen, Benzoni and Lund (2001). For the forgoing reasons, filtering the data seems reasonable. A relatively high-order AR form is necessary to model the predictability in the sample of returns considered in this application. But a parsimonious MA(1) does the same more effectively. Hence, I estimate an MA(1) model for the S&P 500 daily returns and rescale the residuals to match the sample mean and variance in the original data set. The rescaled residuals are then used in the estimation of the SNP model as the original data set. On the other hand, I do not model day-of-the-week, week, month and year effects explicitly. My approach falls somewhere in between that found in Gallant, Rossi and Tauchen (1992), who prefilter the data extensively, and that found in other studies, e.g., Chernov and Ghysels (2000), Chernov et al. (2000), and Pan (2001), in which raw returns are used.

Within the above SNP family, the Bayesian (BIC) and Hannan and Quinn (H-Q) information criteria are used for model selection. (Actual values of the statistics are not reported.) This selection strategy yields an ARMA(0,0)-EGARCH-M(1,1)-Kz(8)-Kx(0), the same specification as that used in Andersen, Benzoni and Lund (2001). Parameter estimates and the associated standard errors are given in Table 2. Ljung-Box tests (not reported) for the autocorrelation in the residuals confirm that the selected specification successfully removes the systematic first- and second-order dependencies in the data.

4.2 The Continuous-Time Stochastic Volatility Model

In this section, I outline the EMM estimation of the continuous-time stochastic volatility models (1) and (2) and (1) and (3), and discuss the results of the associated specification tests.

Let $\{r_t(\psi)\}_{t=1}^{\mathcal{T}(N)}$ denote a sample of S&P 500 returns simulated from the stochastic volatility model using the parameter vector $\psi = (\mu,\ \alpha,\ \beta,\ \eta,\ \rho)$; let $\{x_t(\psi)\}_{t=1}^{\mathcal{T}(N)}$ be a sequence of variables containing lagged S&P 500 returns from the identical simulation. The EMM estimator of ψ is then defined by

$$\hat{\psi}_N = \arg\min_{\psi}\ m_{\mathcal{T}(N)}(\psi,\hat{\xi})'\, W_N\, m_{\mathcal{T}(N)}(\psi,\hat{\xi})\,,$$

where $m_{\mathcal{T}(N)}(\psi,\hat{\xi})$ is the expectation of the score function, evaluated by Monte Carlo integration at the quasi-maximum likelihood estimate of the auxiliary model parameter

$\hat{\xi}$, i.e.

$$m_{\mathcal{T}(N)}(\psi, \hat{\xi}) = \frac{1}{\mathcal{T}(N)} \sum_{t=1}^{\mathcal{T}(N)} \frac{\partial \ln f_K(r_t(\psi)|x_t(\psi); \hat{\xi})}{\partial \xi},$$

and the weighting matrix W_N is, as in Gallant and Tauchen (1996), a consistent estimate of the inverse asymptotic variance matrix of the score function. In simulating the return sequence $\{r_t(\psi), x_t(\psi)\}_{t=1}^{\mathcal{T}(N)}$, two antithetic samples of 75,000 × 10 + 5,000 rates of return are generated from the stochastic volatility model at time intervals of one-tenth of a day.[5] The first 5,000 observations are discarded to eliminate the effect of initial conditions. Finally, a sequence of $\mathcal{T}(N) = 75,000$ daily returns is obtained by summing the elements of the simulated sample in groups of 10.

The results of the estimation of the model (1) and (2) are reported in Table 3. Parameter estimates are expressed in percentage form on a daily basis ($dt = 1$). First, the simulation is performed imposing the restriction $\rho = 0$. That constraint is subsequently relaxed, so that the model can capture the asymmetries in the stock returns and variance. The resulting estimate of ρ, -0.5778, is significant and, as expected, negative. Moreover, an unconstrained ρ dramatically improves the fit of the model; the value of the chi-square statistic drops from 116.98 to 31.53. This is indicative of the presence of negative skewness in stock returns, which is readily captured by the asymmetric stochastic volatility factor, and is qualitatively consistent with the asymmetry observed in option prices. A negative correlation of this magnitude has also been documented in studies of stock returns and associated implied volatilities. Among others, Dumas, Fleming and Whaley (1998) find that the correlation in the first differences of prices and Black-Scholes implied volatilities is -0.57, a figure almost identical to the estimate of ρ reported in Table 3. Similarly, Bakshi, Cao and Chen (1997) estimate ρ from a set of S&P 500 option prices, obtaining a comparable result. Still, the model is rejected.[6] Andersen, Benzoni and Lund (2001) provide further model diagnostics; as they report, the stochastic volatility model is unable to accommodate the excess kurtosis observed in the S&P 500 returns.

For an interpretation of the remaining parameter estimates, notice first that the positive β rules out non-stationarity in the variance process. Also, β provides an indirect measure of the persistence in the (log-)variance process. Based on the estimates for the model (1) and (2), $\exp\{-\beta\} = 0.985$, which indicates a strong daily persistence in (log-)variance. This finding is in line with the results reported in the discrete-time literature; see, e.g., Bollerslev, Engle and Nelson (1994) and references therein.

[5] The Euler scheme is used in the simulation; see, e.g., Kloeden and Platen (1992). EMM estimation with two antithetic simulated samples of 150,000 × 10 + 10,000 returns produce the same results.

[6] This finding is in line with those in other empirical studies on stochastic volatility models in discrete and continuous time; see, e.g., Chernov and Ghysels (2000), Gallant, Hsieh and Tauchen (1997), Gallant and Tauchen (1997), Liu and Zhang (1997) and van der Sluis (1997).

Next, I turn to the model (1) and (3), which is estimated with and without the restriction $\rho = 0$. The results, reported in Table 3, are qualitatively similar to those obtained for the log-variance specification (1) and (2).[7] The value of the chi-square statistic is equal to 30.74 in the case of $\rho \neq 0$. The p-value is comparable to the one obtained for the log-variance model. Other model diagnostics (not reported) show that this candidate model also fails to accommodate the excess of kurtosis observed in the S&P 500 returns.

Based on the analysis conducted above, the specifications (1) and (2) and (1) and (3) seem to perform about the same in fitting stock returns. Neither model passes the specification tests: Andersen, Benzoni and Lund (2001) conclude that continuous-time models must incorporate both stochastic volatility and discrete jumps in order to provide an adequate characterization of equity returns.[8] However, a number of contributions have pointed out that jumps are only of second-order importance in fitting option prices, in comparison with stochastic volatility, except possibly for the shortest lived options; see, e.g., Bakshi, Cao and Chen (1997) and Eraker (2000). Therefore, the stochastic volatility models (1) and (2) and (1) and (3) may serve as a good starting point for explaining option prices. In the present context, adding jumps would make the empirical analysis considerably more difficult and possibly complicate the interpretation of results.

5 Option Pricing Methodology

In this section, I develop the empirical methodology used to estimate the premium for variance risk and compute derivative prices under stochastic volatility. I focus on the log-variance specification (1) and (2), which is more challenging than the square-root representation (1) and (3) because no closed-form solution for the derivative price is available. The methodology developed below is then extended to the Heston model in Section 7.

5.1 The Pricing Model

Suppose that the underlying stock price S obeys the SDEs (1) and (2) and generates a constant dividend yield d. Consider a derivative on the security S with expiration date T, the price f of which depends only on S and the underlying returns variance V:

[7] To facilitate a comparison with the empirical option pricing literature it may be helpful to convert parameter estimates for the square-root model to decimal form on a yearly basis. They become $\alpha = 0.0514$, $\beta = 3.9312$, $\eta = 0.1971$ and $\rho = -0.5973$.

[8] See also Chernov et al. (1999), Eraker (2000), Eraker, Johannes and Polson (2001), Pan (2001). Alternatively, more general volatility specification have been advocated for fitting the index returns' leptokurtosis. See, e.g., Chernov et al. (2000), Jones (1999), and Meddahi (2001).

$f = \{f(S_t, V_t, \tau), t \in [0, T]\}$, where $\tau \equiv T - t$ is the time to maturity. Also, let $h(S_T)$ denote the derivative's payoff at expiration, and $g(S_t, \tau)$ its instantaneous pay-out rate.

An application of Itô calculus and standard equilibrium arguments for a representative agent economy (see, e.g., Cox, Ingersoll and Ross (1985)) yield:

$$Lf(S, V, \tau) = -G(S, \tau)\,, \tag{4}$$

where

$$\begin{aligned} Lf = -\frac{\partial f}{\partial \tau} \;&+\; \frac{1}{2}\Big(VS^2\frac{\partial^2 f}{\partial S^2} + \eta^2 V^2 \frac{\partial^2 f}{\partial V^2} + 2\,\rho\,\eta\, S\, V^{3/2}\,\frac{\partial^2 f}{\partial S \partial V}\Big) - rf + \\ +\;\; & (r-d)S\frac{\partial f}{\partial S} - V\Big(\beta \ln V + \lambda(V)/V - \alpha - \frac{1}{2}\eta^2\Big)\frac{\partial f}{\partial V}\,, \end{aligned} \tag{5}$$

the function $\lambda(V_t)$ is the premium for variance risk and $G(S, \tau) \equiv g(S, \tau) f(S, V, \tau)$. The initial condition

$$f(S, V, 0) = h(S)\,, \tag{6}$$

and (4) characterize the price f.

Additional assumptions are required for computational tractability. Bates (1996a) points out that a no-arbitrage condition is $\lambda(V_t = 0) = 0$ and so suggests assuming that the premium for variance risk is proportional to V_t. Here, as in Melino and Turnbull (1990), $\lambda(V) = \lambda \eta V$.

5.2 Estimation Strategy

The derivative pricing problem (4) and (6) admits a unique solution; see, e.g., Benzoni (2001). In particular, the price of a European call option solves (4) with $G(S, \tau) = 0$ and the initial condition

$$h(S_T) = \max\{0, S_T - K\}. \tag{7}$$

The objective is to compute the time-t option price which solves (4) and (7). Note that the solution to this problem is a function

$$f(S_t, V_t, \tau\,; \alpha, \beta, \eta, \rho, \lambda, K)\;:\; \mathbf{R}^+ \times \mathbf{R}^+ \times [\,0, \mathrm{T}\,] \;\to \mathbf{R}^+\,,$$

The stock price S_t and the strike price K are observable at time t, but the true values of the parameters $(\alpha,\ \beta,\ \eta,\ \rho)$ and λ are unknown; also unknown is the time-t variance V_t. Therefore the pricing of an option contract necessarily involves the estimation of these inputs.

To obtain those estimates, I exploit the information contained in equity returns and option prices in a multi-stage procedure. In a first stage, a sample of daily stock returns is used to estimate the parameter vector ψ, as explained in Section 4. In the second step, the premium for variance risk is estimated using a sample of option prices and the EMM estimate of ψ.[9] I develop the asymptotic properties of the multiple-step estimator and test the statistical significance of the risk premium coefficient λ.

In the following sections, I provide the details of the methodology for estimating λ. Two strategies are presented. The first follows a simulated method of moments approach; the second is based on the largely diffused practice of minimizing the SSR. Both methods make use of an econometric model for option pricing errors, which is described in the next section.

5.3 An Econometric Model of Option Pricing Errors

No-arbitrage and equilibrium arguments are central to the theoretical option pricing literature and are commonly used to value derivative securities. The option pricing implications of these models are, however, often inconsistent with the data, and model prices do not match derivative prices exactly. Market microstructure effects, e.g. non-synchronous trading and price discreteness, measurement errors and other market frictions, in addition to possible model specification errors, play an important role in explaining such mispricing.

Of course, by arbitrarily increasing the number of factors one can always duplicate perfectly the observed paths of option prices. Alternatively, Rubinstein (1994) suggests the use of implied binomial trees to obtain a perfect in-sample fit for derivative prices. Dumas, Fleming and Whaley (1998) develop a deterministic volatility function option valuation model that has the potential for fitting any observed cross-section of option prices exactly. These approaches, though, are sensitive to over-fitting problem. The same specification which obtains a perfect in-sample fit can perform poorly out-of-sample. This is documented by, e.g., Dumas et al. (1998), who find that an ad-hoc Black-Scholes implementation consistently outperforms the deterministic no-arbitrage model, even if the in-sample fit of the latter is quasi perfect. (See also Jacquier and Jarrow (2000) and Renault (1996).)

To avoid these problems, I assume

$$\frac{\tilde{f}(\tau, K)}{S_t} = \frac{f(S_t, V_t, \tau\,; \psi, \lambda, K\,)}{S_t} + \varepsilon_t\,, \tag{8}$$

where $\tilde{f}(\tau, K)$ is the observed option price for a given strike price K and τ days to maturity, $f(S_t, V_t, \tau\,; \psi, \lambda, K\,)$ is the price predicted by the stochastic volatility model,

[9] Aït-Sahalia (1996) adopts a similar approach in estimating the term structure of interest rates and in pricing interest rates derivatives. A nonparametric estimation of the continuous-time model is performed first, and an interest rate risk premium is computed in a second stage, conditional on the estimates obtained in the first stage.

and the error process $\varepsilon = \{\varepsilon_t, t \in [0, T]\}$ is assumed to be stationary and ergodic with zero mean. Note that, to simplify notation, I suppress the arguments K and τ of the error term ε.

Given the model for the option pricing errors, in the following sections I illustrate the estimation of the market price for variance risk.

5.4 Estimating the risk premium by SMM

The assumption that the error term in (8) has zero mean yields the moment condition[10]

$$E\left[\frac{\tilde{f}(\tau, K)}{S_t}\right] - E\left[\frac{f(S_t, V_t, \tau\,; \psi, \lambda, K)}{S_t}\right] = 0\,. \tag{9}$$

For any given strike price K, only a small sample of option prices is available. Indeed, since the exercise price is set by the Chicago Board Option Exchange (CBOE) as a function of the underlying index level, options with a given strike K are only traded over a relatively brief period. Of course, these moments are matched for different strike prices K. However, if options with different exercise prices are used in estimating $E[\tilde{f}(\tau, K)/S_t]$, the same strike price pattern in market prices should be replicated in the simulated sample used to approximate $E[f(S_t, V_t, \tau\,; \psi, \lambda, K)/S_t]$. To circumvent these problems, I exploit the homogeneity property of option prices; see, e.g., Merton (1973). Together with (9), homogeneity implies

$$E\left[\frac{\tilde{f}(\tau, K)}{S_t}\right] - E\Big[f(1, V_t, \tau\,; \psi, \lambda, K/S_t)\Big] = 0.$$

If, at any maturity, options with a strike $K = \gamma S_t$, where γ is a constant, are traded, then the previous equation simplifies to

$$E\left[\frac{\tilde{f}(\tau, \gamma S_t)}{S_t}\right] - E\Big[f(1, V_t, \tau\,; \psi, \lambda, \gamma)\Big] = 0. \tag{10}$$

It is easy to show that the normalized option price process $\{\tilde{f}(\tau, \gamma S_t)/S_t, t \in [0, T]\}$, is stationary. Hence, one can estimate the population moment in (10) using a set of observations $\tilde{f}(\tau, \gamma S_{t_n})/S_{t_n}$, $n = 1, \ldots, N$, on option contracts sampled at times t_n, $n = 1, \ldots, N$, having time to maturity τ and strike price γS_{t_n}, where γ is fixed throughout the sample.

An obvious criticism of this approach is that exercise prices take only discrete values. For S&P 500 contracts, the strike price interval is 5 points, and thus, since $\tilde{f}(\tau, K)/S_t \neq \tilde{f}(\tau, \gamma S_t)/S_t$, assuming $K = \gamma S_t$ introduces a measurement error.

[10] Bakshi, Cao and Chen (2000) exploit similar moment conditions to estimate a square-root stochastic volatility jump-diffusion, using only a sample of option prices.

To avoid this problem, I select, for given values of γ and S_t, the traded contract having exercise price K_t^* nearest to γS_t. Setting $\gamma_t^* = K_t^*/S_t$, I consider the econometric model

$$\frac{\tilde{f}(\tau, K_t^*)}{S_t} = f(1, V_t, \tau\,; \psi, \gamma_t^*\,) + \varepsilon_t\,, \tag{11}$$

where the error process $\varepsilon = \{\varepsilon_t\}$ satisfies the same regularity properties as in (8). From (11), I obtain a moment condition which can be used in the SMM estimation of λ:

$$E\left[\frac{\tilde{f}(\tau, K_t^*)}{S_t}\right] - E\left[f(1, V_t, \tau\,; \psi, \lambda, \gamma_t^*\,)\right] \;\; = \;\; 0\,. \tag{12}$$

To exploit condition (12), $E[\tilde{f}(\tau, K_t^*)/S_t]$ is approximated by the average of a sequence $\{\tilde{f}(\tau, K_{t_n}^*)/S_{t_n}\}_{n=1}^{N}$ of option contracts sampled at times t_n, $n = 1, \ldots, N$, each having time to maturity τ, with $K_{t_n}^*$ chosen to be the strike price closest to γS_{t_n}, where γ is fixed throughout the sample. The approach is operational, and uses the entire span of option prices, going back to the year of the introduction of the S&P 500 contract.

The expectation $E[f(1, V_t, \tau\,; \psi, \lambda, \gamma_t^*\,)]$ is computed by simulation. I use numerical methods to calculate the stochastic volatility price which appears in the expectation. The parameter vector ψ is replaced by the EMM estimate $\hat{\psi}$. For $S_t = 1$ and a fixed time to maturity τ, I create a fine grid of call option prices, each a function of λ, K, and the variance process V_t, by numerically solving the partial differential equation (4) with initial condition (7). Using the finite differencing method discussed in Appendix B, I compute the option price on the grid of possible strike price and variance values for a given value of λ.

Once the entire grid of option prices is available, it is straightforward to compute the expectation $E[f(1, V_t, \tau\,; \psi, \lambda, \gamma_t^*\,)]$ by simulation. First, I simulate a sample $\{S_{t_n}, V_{t_n}\}_{n=1}^{\mathcal{T}(N)}$ using model (1) and (2). Then, I construct a sequence of strike prices $\{\gamma_{t_n}^*\}_{n=1}^{\mathcal{T}(N)}$, where $\gamma_{t_n}^* = K_{t_n}^*/S_{t_n}$, and $K_{t_n}^*$ is the strike price nearest to γS_{t_n}. Finally, I extract a sequence $\{f(1, V_{t_n}, \tau\,; \hat{\psi}, \lambda, \gamma_{t_n}^*\,)\}_{n=1}^{\mathcal{T}(N)}$ from the grid of option prices. Using this sample, the expected value of the stochastic volatility price appearing in the moment conditions (12) is approximated by $1/\mathcal{T}(N) \sum_{n=1}^{\mathcal{T}(N)} f(1, V_{t_n}, \tau\,; \hat{\psi}, \lambda, \gamma_{t_n}^*\,)$.

In this approach, (12) is approximated by

$$\frac{1}{N} \sum_{n=1}^{N} \frac{\tilde{f}(\tau, K_{t_n}^*)}{S_{t_n}} - \frac{1}{\mathcal{T}(N)} \sum_{n=1}^{\mathcal{T}(N)} f(1, V_{t_n}, \tau\,; \hat{\psi}, \lambda, \gamma_{t_n}^*\,) \;\; \cong \;\; 0\,. \tag{13}$$

Selecting different values of γ and τ, I obtain a set of sample moment conditions which is then used to compute a consistent SMM estimate $\hat{\lambda}$.

Finally, the time-t option price is computed by solving numerically the pricing partial differential equation (4) with its initial condition (7), using $\hat{\lambda}$ and the EMM estimate $\hat{\psi} = (\hat{\mu}\; \hat{\alpha}\; \hat{\beta}\; \hat{\eta}\; \hat{\rho}\,)$.

5.5 Estimating the risk premium by minimizing the SSR

The SMM method outlined above delivers consistent estimates of λ. Also, it provides a setting for assessing the statistical significance of the volatility risk premium and constructing a Chi-square statistic for testing the over-identifying restrictions. However, it is silent about the economic properties of the stochastic volatility option pricing model.

To gain some knowledge of these properties, I also consider an estimate of λ that minimizes the SSR for the market and stochastic volatility prices. Given a set of $n = 1 \dots N$ contracts having $\tau_n = T - t_n$ days to maturity and strike price K_i, $i = 1, \dots, I$, the estimated market price of variance risk λ is

$$\hat{\lambda} = \arg\min_{\lambda} \sum_{n=1}^{N} \sum_{i=1}^{I} \left\{ \frac{f(S_{t_n}, \hat{V}_{t_n}, \tau_n; \hat{\psi}, \lambda, K_i)}{S_{t_n}} - \frac{\tilde{f}(\tau_n, K_i)}{S_{t_n}} \right\}^2 , \tag{14}$$

where $\hat{V}_t$ is an estimate of the time-t variance. The option price $f(S_{t_n}, \hat{V}_{t_n}, \tau_n; \hat{\psi}, \lambda, K_i)$, as a function of λ, is computed solving (4) and (7) numerically. The specific details of the numerical algorithm are given in Appendix B.

The choice of the conditional variance estimate $\hat{V}_t$ is critical for the performance of the option pricing model. In the following application, three estimates are considered. Two are computed using stock returns only. One is obtained applying the Kalman filter to the discrete version of model (1) and (2) (for details, see Appendix A); the other is obtained using the reprojection method (see, e.g., Gallant and Tauchen (1998), Chernov and Ghysels (2000), Chernov et al. (2000), and Jiang and van der Sluis (1999)). The third estimate of V_t makes use of the information contained in option contracts: it is given by the sequence $\{\hat{V}_{t_n}\}_{n=1}^{N}$ that, simultaneously with λ, minimizes the SSR between market and stochastic volatility prices.[11] In what follows, the option-based estimates are called "SV implied volatilities" to distinguish them from the more commonly used Black-Scholes implied volatilities. A comparison of the conditional variance estimates, based on the performance of the option pricing model, addresses the extensively debated issue of which source of information should be used to price options, and provides a basis for testing the internal consistency of the model.

[11] This approach is also followed by Bates (1996a) and Bakshi, Cao and Chen (1997). SV implied volatilities will produce a smaller SSR between market and stochastic volatility prices than Kalman filter estimates. Actually, the use of implied volatilities could bias the results, since by design this approach directly minimizes the SSR over V_t. This and related issues are addressed in Section 7.

6 Asymptotic Properties of the Estimators

The solution to the option pricing problem involves a sequential estimation procedure, which makes a direct derivation of the estimator's asymptotic properties cumbersome. The problem is better handled by writing the multiple-step estimation in a "method of moments" form, so that standard asymptotics from GMM and SMM theory can be applied.[12]

The following proposition summarizes the asymptotic properties of the estimator; additional details are provided in Appendix C. Under the assumptions and the regularity conditions discussed in Duffie and Singleton (1993) and Gallant and Tauchen (1996), we have the following

Proposition 1 *Suppose $N/\mathcal{T}(N) \to \xi$ as $N \to \infty$. The asymptotic distribution of the sequential estimator $(\hat{\psi}\ \hat{\lambda})$ is*

$$\begin{bmatrix} \sqrt{N}(\hat{\psi} - \psi) \\ \sqrt{N}(\hat{\lambda} - \lambda) \end{bmatrix} \stackrel{D}{\to} N(0, H)\,. \tag{15}$$

The expression for H depends on the estimation method of λ that is adopted (see Appendix C).

The price of an option with a given strike price K satisfies

$$\sqrt{N}\,[\,f(S,V,\tau;\hat{\psi},\hat{\lambda}) - f(S,V,\tau;\psi,\lambda)] \stackrel{D}{\to} N(0,\Lambda' H \Lambda) \tag{16}$$

where $\Lambda \equiv D_{\psi,\lambda} f(S,V,\tau\,;\psi,\lambda)$, i.e. Λ is the vector of the derivatives of the option price with respect to the parameter vector $(\psi\ \lambda)$.

Proof: See Appendix C. □

7 Empirical Option Pricing Results

In this Section, I apply the procedure discussed above to estimate λ, the market price for variance risk, assess its statistical significance and discuss its economic relevance for the pricing of European call options on the S&P 500 index (SPX). I also compare the performances of the log-variance model (1) and (2), the square-root specification (1) and (3), and the Black-Scholes benchmark. In- and out-of-sample diagnostics are used to highlight the main features of the different specifications. Additional model diagnostics based on SV implied volatilities are then provided. First, I check for the presence of possible

[12] See Newey (1984) for a discussion of a method of moments interpretation of multiple-step estimators, and Heaton (1995) for an application of this methodology.

patterns, one being the Black-Scholes smile. Then I compare SV implied volatilities and return-based estimates of the variance process, which have been obtained using either the Kalman filter or the reprojection method of Gallant and Tauchen (1998). Finally, stochastic volatility SPDs are used to compare and interpret the (mis)pricing produced by the different specifications.

7.1 Option pricing when the risk premium is estimated by SMM

In this section, I report on the SMM estimation of λ, based on the methodology discussed in Section 5.4. The choice of γ and τ in the moment condition (12) is guided by the characteristics of the option contracts traded in the market. It is well known that call options deep in-the-money are often very illiquid, as compared to at- and out-of-the-money contracts; hence I set γ equal to one. With regard to time-to-maturity, there are two issues. The presence of stochastic volatility is generally more important for options that are far from maturity than for contracts close to expiration. On the other hand, at longer maturities the option contracts are relatively illiquid. Because of this trade-off, I set τ equal to 3, 4 and 5 weeks, thus obtaining 3 moment conditions. The nominal risk-free interest rate is fixed at 6%; also, a constant 2% dividend yield is assumed in this application and hereafter. Option prices are computed numerically from the pricing equation (4).

The results are provided in Table 4. Standard errors are consistently estimated from the asymptotic variance matrix (27). Note that the risk premium is statistically significant (the coefficient's t-ratio is -3.8), which confirms the conjecture that variance risk is priced by the market. This conclusion is supported by the results of a likelihood ratio test, which rejects the null $\lambda = 0$. Further model diagnostics are provided by a test for over-identifying restrictions: the p-value for the Chi-square statistics is 5%.

7.2 Option pricing when the risk premium is estimated minimizing the SSR

In order to provide better economic insights into the performances of the stochastic volatility models, I also estimate the market price of variance risk λ by minimizing the SSR of market and stochastic volatility prices, as discussed in Section 5.5. I use a sample of S&P 500 options expiring in March 1997; daily observations from December 30, 1996 to March 14, 1997 are from the CBOE tapes. From here on, the nominal risk-free interest rate is fixed at 5.1%.

A number of alternative estimates for the variance process are investigated here. First, I consider return-based methods and value options using the Kalman filter estimate of the daily conditional variance V_t. Initially, I set $\lambda = 0$ and compute the corresponding value of the SSR, reported in Table 5. I then minimize the SSR over the coefficient λ. The squared option pricing errors drop significantly in comparison with the $\lambda = 0$ case (Table 5). This also confirms the existence of a non-zero premium for variance risk. This finding is consistent with those in Melino and Turnbull (1990) and Jiang and van der Sluis (1999). However, the estimate of λ obtained with this method is relatively imprecise; the SMM approach discussed in Section 7.1 (Table 6) provides a better estimate.[13] Standard errors are computed by plugging the estimates $(\hat{\psi}\ \hat{\lambda})$ and $\hat{V}$ into the asymptotic variance matrix (28).

This analysis provides a first diagnostic of in-sample model performance based on the comparison of the SSRs produced by the stochastic volatility model and the Black-Scholes benchmark. In drawing this comparison, two estimates of Black-Scholes volatility are considered. The first is obtained estimating the diffusion coefficient in the Black-Scholes model by EMM, as in Andersen, Benzoni and Lund (2001). The second is the value of the implied volatility which minimizes the SSR of Black-Scholes and market prices. The results are reported in Table 5: for both estimates of the Black-Scholes volatility, the log-variance model delivers a smaller SSR. However, the improvement is not as large as one might have expected. Little difference in the performances of the Black-Scholes and stochastic volatility models is also suggested by the comparison of the average absolute pricing error produced by the two, which is approximately \$1 in both cases (Table 5).

To shed some light on the last result, I investigate the effect of different variance estimates on option pricing errors. As an alternative to the Kalman filter method, I consider now the reprojection technique of Gallant and Tauchen (1998) (See also Chernov and Ghysels (2000), Chernov et al. (2000) and Jiang and van der Sluis (1999)). I simulate a sample of returns from the stochastic volatility model using the parameter estimates in Table 3, and fit an SNP conditional density $f_K(r_t|x_t;\xi)$ on the simulated sequence of returns. The family of SNP densities considered for this application is the one used in Section 4.1. The simulated sample consists of 75,000 daily returns. Based on standard specification tests I select an ARMA(0,0)-EGARCH-M(1,1)-Kz(4)-Kx(0) representation. This model is readily used in the estimation of the conditional variances of S&P 500 returns:

$$E(r_t|x_t;\xi) \;\; = \;\; \int y\, f_K(y|x_t;\xi)\, dy\,, \tag{17}$$

[13]The formal inference about λ is based on the results reported in Section 7.1, since estimates of λ obtained from the minimization of the sum of squared residuals are not consistent. Nevertheless, the analysis based on option pricing errors illustrated in this section provides additional economic insights into the performance of the two stochastic volatility models.

$$Var(r_t|x_t;\xi) \quad = \quad \int (y - E(y|x_t;\xi))^2 \, f_K(y|x_t;\xi)\, dy\,, \tag{18}$$

where r_t denotes the actual S&P 500 return and x_t is a vector of lagged actual returns. I use these variance estimates for computing option prices. First, I evaluate the SSR at $\lambda = 0$. Then, I minimize the SSR over λ. The results, reported in Table 5 under the heading "reprojection method" are very similar to those obtained using the Kalman filter method. So also are the average absolute price deviations.

The two return-based estimates of V_t do not do well in eliminating the mispricing of options. Hence, I turn now to option-based measures of variance, and use SV implied volatilities, obtained using the procedure described in Section 5.5, to assess the model performance. As indicated in Table 5, the change to SV implied volatilities results in quite a sharp decrease in the SSR, to a value which is dramatically smaller than that of the Black-Scholes benchmark. This improvement is also reflected in the average absolute pricing deviation, which is only \$0.59 or alternatively 6.51%. One possible interpretation of the changes in pricing errors produced by the use of SV implied volatilities is that a reliable estimate of the conditional variance is critical for getting at correct option prices. Same-day SV implied volatilities obtained from option prices may contain more information than do return-based estimates obtained from stock returns only, and thus be a more reliable measure of volatility for the purpose of derivative pricing. It is also true, though, that using implied volatilities may bias results, since by design this approach directly minimizes the SSR over V_t. It is also worth noting that the SSR computed at the optimal value of λ is only marginally smaller than when λ is set equal to zero (Table 5). The market premium for variance risk affects option prices by changing the unconditional mean of the variance process. When λ is set equal to zero, the value of the implied volatility obtained from option prices is significantly higher than that found when λ is fixed at its estimated value. This offsets the difference in the values of λ, and reduces option pricing errors. To resolve the issue of whether SV implied volatilities bias results, it is advisable to investigate further the time series properties of the sequence of implied volatilities, in order to uncover possible inconsistencies between the time series of returns and option prices. Moreover, to assess whether a non-zero λ is necessary to fit option prices, it is also important to evaluate the out-of-sample performance of the model and to examine the impact of λ on SPDs. These issues are explored in the following sections.

To conclude this subsection, I consider the importance of the asymmetry parameter ρ for option pricing. More specifically, I compute the SSR corresponding to the EMM estimates for the symmetric model ($\rho = 0$), and compare the pricing errors to those obtained using the preferred asymmetric specification ($\rho \neq 0$). The results for the log-variance representation are reported in Table 5; no matter what variance estimate is used in computing option prices, the asymmetric model performs remarkably better than the symmetric model. This confirms the findings reported in Section 4.2: a strongly negative asymmetry coefficient ρ is crucial not only to fitting the underlying returns dynamics, but also the characteristics of option prices.

7.3 The Heston model

In this section, I compare the log-variance option pricing model (1) and (2) to the alternative specification (1) and (3) considered by Heston (1993). As in the log-variance case, appropriate adjustments are made to account for the presence of non-diversifiable volatility risk. For a representative agent economy, this is achieved by writing the model in "risk-neutral" form:

$$\frac{dS_t}{S_t} = (r-d)dt + \sqrt{V}_t \, d\hat{W}_{1,t} \,, \tag{19}$$

$$dV_t = (\alpha - \beta V_t - \lambda(V_t)\,)\, dt + \eta \sqrt{V}_t \, d\hat{W}_{2,t} \,, \tag{20}$$

where $\hat{W}_1$ and $\hat{W}_2$ are standard Brownian Motions under a "risk-neutral" probability measure and $\text{corr}(d\hat{W}_{1,t}, d\hat{W}_{2,t}) = \rho$. For tractability, the premium for variance risk is assumed to be proportional to the conditional variance, i.e. $\lambda(V_t) = \lambda V_t$. With this specification, a closed-form formula, provided in Appendix D, is available for computing option prices.

For a comparison with the log-variance model, I estimate λ by minimizing the SSR between Heston and market prices, as in Section 7.2. As in the log-variance case, the EMM estimate $\hat{\psi}$ is used in (14). In evaluating option prices, I approximate the conditional variance with those SV implied volatilities that minimize the SSR for market and stochastic volatility prices, computed at the optimal value of λ. The optimal value of λ is -0.0278 in this case. Under the risk-neutral distribution, the persistence parameter of the variance process is $\beta + \lambda$. The estimated value of λ implies that the variance process is non-stationary under the risk-neutral distribution, a result consistent with what Pan (2001) finds using a different estimation method. Interestingly, this does not happen in the case of the log-variance specification. The value of the SSR appears in Table 5. Also in this respect the Heston model performs similarly to the log-variance specification, and dramatically out-performs the Black-Scholes benchmark.

7.4 Out-of-sample results

Model diagnostics are obtained evaluating the out-of-sample SSR for market and stochastic volatility prices. I use a data set of European call options on the S&P 500 index expiring in June 1997; closing prices are sampled daily from March 31, 1996, to June 13, 1997. The sum of squared residuals is computed as in (14), using the in-sample estimates of $\hat{\psi}$, $\hat{\lambda}$, and $\hat{\sigma}$ given in Section 7.2.

I start with the log-variance representation and use the Kalman filter to estimate the variance process. The results are reported in Table 7. Interestingly, the model does not perform as well as the Black-Scholes benchmark. Next, I reproject V_t using (17)-(18), but without getting any improvement. The results appear in Table 7 under the heading "reprojection method." Much better performance is obtained instead when option prices are evaluated at day-before SV implied volatilities. The SSR decreases by a factor of approximately 2 in comparison with Black-Scholes. In Table 7 I also report the SSR obtained using same-day implied volatilities, which, given the model parameters estimates, provide a lower bound for pricing errors.

The results for the Heston representation are similar; the value of the SSR evaluated at day-before SV implied volatilities is only slightly bigger than that obtained in the log-variance case (Table 7).

In summary, the analysis in this subsection confirms that the choice of the return volatility estimate is critical for option pricing. Specifically, variance estimates obtained from derivative prices dominate those computed by filtering stock returns, and they reduce considerably Black-Scholes pricing errors.

7.5 SV Implied Volatilities

Since the choice of the conditional variance estimate is critical for the purpose of derivative pricing, it is important to investigate further the properties of SV implied volatilities obtained from option prices and those of return-based estimates of variance.

First, I use the stochastic volatility model to compute a volatility estimate based on option prices for contracts with different degrees of moneyness. In doing this, I fix the error term in (8) at zero, and solve for the value of the variance which makes market and model prices equal. I use options expiring in April 1997 with up to two months to maturity. The SV implied volatilities are computed solving (4) numerically and are given in Figure 2.

Allowing for the presence of stochastic volatility flattens but does not eliminate the volatility smile. For example, on March 21, 1997, the values of the standard deviations implied by the stochastic volatility model are concentrated in a 1.2% range for at- and out-of-the-money calls (2.3% for Black-Scholes). The pattern of Black-Scholes implied volatilities, also in Figure 2, is a strong asymmetric smile. The pattern of the estimates implied by the log-variance model is less strong; it appears to be more a scatter around some central value. However, in-the-money calls exhibit significantly higher implied volatilities than do other contracts. This indicates that stochastic volatility is unable to explain fully in-the-money call and out-of-the-money put prices.

Next, I investigate the time series properties of volatilities. Specifically, I regress SV implied volatility over return-based variance estimates, as in

$$\text{Imp. Vol.} = b_1 + b_2 \times \text{Vol. Est.} + \varepsilon,$$

where the independent variable is initially given by Kalman filter estimates, and then by reprojected volatilities. This application uses data from the sample period from December 30, 1996 to March 21, 1997. The results are reported in Table 8. Overall, there are considerable discrepancies between return and option based estimates of volatility. This finding is consistent with the difference in option pricing errors produced by the two volatility estimates, and is indicative of inconsistencies in the joint model for option prices and stock returns.[14]

7.6 State-price densities

It is well known that in a stochastic volatility setting the price of a derivative with a single liquidating payoff $h(S_T)$ is given by

$$f(S_t, V_t, \tau; \psi, \lambda) = \mathrm{e}^{-r\tau} \int_{\mathbf{R}^+ \times \mathbf{R}^+} h(S_T) \pi(S_T, V_T, \tau; S_t, V_t, \psi, \lambda) dS_T dV_T \,, \tag{21}$$

where $\pi(S_T, V_T, \tau; S_t, V_t, \psi, \lambda)$ denotes the SPD conditional on the time-t stock price S_t and the variance V_t.

The SPD π summarizes the characteristics of derivative prices and can therefore be used to compare alternative pricing models. To facilitate the contrast with Black-Scholes, it is convenient to integrate out the terminal variance V_T from π and introduce

$$\pi^*(S_T, \tau; S_t, V_t, \psi, \lambda) = \int_{\mathbf{R}^+} \pi(S_T, V_T, \tau; S_t, V_t, \psi, \lambda) \, dV_T \,.$$

Different methods deliver an estimate of $\pi^*(S_T, \tau; S_t, V_t, \psi, \lambda)$.[15] In this application, I rely on Monte Carlo integration. First, I simulate a sequence of terminal stock prices using

$$\frac{dS_t}{S_t} = (r - d)dt + \sqrt{V}_t \, d\tilde{W}_{1,t} \,, \tag{22}$$

$$d \ln V_t = (\tilde{\alpha} - \beta \ln V_t \,) dt + \eta \, d\tilde{W}_{2,t} \,, \tag{23}$$

[14] Bates (1996a) and Bakshi, Cao and Chen (1997) also find internal inconsistencies in the joint model of option prices and underlying returns.

[15] See, e.g., Aït-Sahalia and Lo (1998), Aït-Sahalia et al. (2001) and Breeden and Litzenberger (1978). Closed-form solutions are available for the Black-Scholes and Heston models (see, e.g., Bakshi, Cao and Chen (2000)) but not for the log-variance representation.

where $\tilde{\alpha} \equiv \alpha - \lambda\eta$, and $\tilde{W}_1$ and $\tilde{W}_2$ are Standard Brownian Motions under a risk neutral probability measure. Then I obtain a nonparametric estimate of the SPD by fitting a kernel to the simulated sample of prices.

To develop some initial feel for the behavior of the stochastic volatility SPDs, I set $\lambda = 0$. I use SV implied volatilities obtained from same-day option prices as a proxy for the time-t variance V_t, and substitute its EMM estimate for ψ. I compute stochastic volatility SPDs for both the log-variance and Heston models and compare them to the Black-Scholes benchmark. Figure 3, top panel, depicts the March 21, 1997, SPD estimates relative to option contracts expiring in April 1997. At first glance, the plots reveal that, as compared to Black-Scholes, the stochastic volatility SPDs exhibit excess skewness and kurtosis. This is consistent with the empirical evidence about option prices after the 1987 crash, which is summarized by the asymmetric smile in the Black-Scholes implied volatilities; see, e.g., Rubinstein (1994). It is also of interest that there is no apparent difference between the SPDs implied by the log-variance and Heston models. This confirms that the derivative pricing implications of the two representations are very similar.

Next, I compute SPDs assuming a non-zero market price of variance risk: λ is set equal to the estimates obtained in the previous subsections. SPDs are evaluated at a new value of the implied volatility, calculated using $\lambda = -0.1335$ in the log-variance model and $\lambda = -0.0278$ in the square-root model. The results are reported in Figure 3, bottom panel, and do not differ significantly from those obtained with $\lambda = 0$. This finding is consistent with the results discussed in Section 7.2. The difference in the value of the λ coefficient is offset by the new implied volatility estimate, so that in both cases the estimated SPDs fit very closely the risk-neutral density implicit in option prices. Hence the similarity in the patterns of the two panels.

In order to investigate the impact of λ on SPDs, it is therefore important to isolate the effect of λ on the volatility estimate. To this end, I evaluate SPDs using return-based estimates of the variance obtained using both the Kalman filter and the reprojection methods, which are independent of the market price of variance risk. Figure 4 illustrates the impact of a non-zero λ on the log-variance model SPDs, and compares them to those of Black-Scholes. This analysis reinforces the conclusion that the premium for variance risk is not only statistically but also economically significant. That is, the presence of a non-zero λ fattens considerably the tails of the risk-neutral distribution of stock prices.

The final question is how SPDs change over time compared to the Black-Scholes benchmark. The answer is given in Figure 5, in which are depicted the SPD estimates based on option contracts expiring in April 1997, evaluated at different maturities. Differences are more remarkable the longer the time to expiration. At longer maturities, stochastic volatility risk-neutral densities become more leptokurtic and have a fatter left tail. This is not surprising; the stochastic volatility model generates outliers mainly over longer maturities because of the continuity of the return and volatility sample paths. The patterns in the Figure suggest a potential limitation of the model. To generate excess kurtosis at the shortest maturities, it is necessary to further generalize the model, for example, by incorporating a jump component.[16]

8 Conclusions

In this paper the Black-Scholes model for equity returns is generalized to a continuous-time stochastic volatility setting. Specifically, two common stochastic volatility diffusion models, the log-variance (Melino and Turnbull (1990)) and the square-root (Heston (1993)) specifications, are compared using as a criterion the fit for S&P 500 index returns and option prices.

Stochastic volatility improves the fit of S&P 500 returns considerably. However, both models fail the EMM specification tests, which suggests that further extensions are necessary for modeling equity index returns, and, indeed, Andersen et al. (2001) conclude that continuous-time models must incorporate both stochastic volatility and discrete jumps in order to provide an adequate characterization of S&P 500 returns. Similar conclusions are reached by Chernov et al. (1999), Eraker (2000), Eraker, Johannes and Polson (2001), and Pan (2001).

Stochastic volatility also greatly enhances the performance of the option pricing model by flattening the volatility smile and reducing considerably option pricing errors. The analysis demonstrates the high sensitivity of stochastic volatility prices to the level of the return volatility. Return-based estimates of the variance process, obtained using either the Kalman filter or the reprojection method of Gallant and Tauchen (1998), produce large pricing errors. The mispricing is considerably reduced when option prices are evaluated using SV implied volatilities obtained from a cross-section of option prices. Also, large discrepancies are observed across return- and option-based measures of volatilities, suggesting the existence of inconsistencies in the joint model for equity returns and derivative prices.

Volatility risk is priced in the S&P 500 option market; the risk premium coefficient is statistically significant. Analysis of SPDs, options implied volatilities and pricing errors confirms this finding and shows the relevant economic impact that the risk adjustment has on derivative prices.

Finally, the examination of SPDs suggests that the stochastic volatility model is able

[16]This conclusion is in line with those in, e.g., Bakshi and Cao and Chen (1997), Eraker (2000), and Das and Sundaram (1998).

to generate skewness and leptokurtosis in the return process over relatively long time horizons, but not so much over the shortest maturities. This evidence indicates that further extensions of the model, e.g., incorporating a jump factor, may improve the model pricing of the shortest term options.

Appendices

A Kalman Filter Estimate of the Variance Process

In order to compute the Kalman Filter estimates of the variance process, I write the log-variance stochastic volatility model in state-space form. First, time is made discrete. So the model is

$$\begin{aligned} r_t &= \mu\Delta t + \sqrt{V}_t \Delta W_{1,t}\,, \\ \ln V_t - \ln V_{t-1} &= (\alpha - \beta \ln V_{t-1})\Delta t + \eta \Delta W_{2,t}\,, \end{aligned}$$

where $r_t \equiv \frac{\Delta S_t}{S_t}$. Note that when the joint errors distribution is symmetric, the state-space form has uncorrelated disturbances; see Harvey and Shephard (1996). Second, the observation equation is linearized:

$$\begin{aligned} L_t &= \ln V_t + \ln \Delta t + E[\ln(Z_{1,t}^2)] + \xi_t\,, \\ \ln V_t &= \alpha\Delta t + (1-\beta\Delta t)\ln V_{t-1} + \eta\sqrt{\Delta t}\, Z_{2,t}\,, \end{aligned}$$

where $L_t \equiv \ln(r_t - \mu\Delta t)^2$, $\xi_t \equiv \ln(Z_{1,t}^2) - E[\ln(Z_{1,t}^2)]$, $Z_{1,t}$ and $Z_{2,t}$ are distributed N(0,1).

The Kalman filter then yields

$$\hat{\ln V}_{t|t} = \alpha\Delta t + \Gamma_t\,(L_t - \alpha\Delta t - \ln\Delta t + 1.27) + (1-\beta\Delta t)\,(1-\Gamma_t)\,\hat{\ln V}_{t-1|t-1}\,,$$

where Γ_t is recursively defined by

$$\begin{aligned} \Gamma_t &= \frac{\eta^2\Delta t\, y_t + \pi^2/2\, x_t}{\eta^2\Delta t\, y_t + \pi^2/2\, x_t + \pi^2/2}\,, \\ y_t &= 1 + (1-\beta\Delta t)^2(1-\Gamma_{t-1})^2\, y_{t-1}\,, \\ x_t &= \Gamma_{t-1}^2(1-\beta\Delta t)^2 + (1-\Gamma_{t-1})^2(1-\beta\Delta t)^2\, x_{t-1}\,, \end{aligned}$$

with $y_1 = 1$ e $x_1 = 0$.

The parameter vector $(\mu\ \alpha\ \beta\ \eta)$ is replaced by the EMM estimates $(\hat{\mu}\ \hat{\alpha}\ \hat{\beta}\ \hat{\eta})$ computed in the first stage; the sample variance of continuously compounded rate of returns is used as initial condition for the recursion. The daily variance sequence is constructed using a sample of S&P 500 returns from January 1983 to March 1997; the first part of the sequence is then discarded, to eliminate the effect of the initial condition.

B Numerical Solution to the Option Pricing Partial Differential Equation

In this Appendix, I discuss the application of finite differencing methods to solve the pricing differential equation (4).

In the SMM procedure of Section 5.4 a stochastic volatility option price must be computed for each simulated strike $\gamma^*_{t_n}$, $n = 1, \ldots, \mathcal{T}(N)$. Since the initial condition for (4) depends on the exercise price, this application entails solving numerically a partial differential equation for each simulated strike—a task which is computationally not feasible. In practice, though, this problem is avoided by using the homogeneity property of the option price. At each simulated point, the option's value in (13) can be computed as:

$$f(1, V_{t_n}, \tau\,; \psi, \lambda, \gamma^*_{t_n}) = \gamma^*_{t_n}\; f(1/\gamma^*_{t_n}, V_{t_n}, \tau\,; \psi, \lambda, 1),$$

where $f(1/\gamma^*_{t_n}, V_{t_n}, \tau\,; \psi, \lambda, 1)$ is the solution to (4) for an option having strike equal to unity. Thus, the initial condition to (4) is identical for each simulated point, and I only have to solve the partial differential equation once, with a huge saving in computer time.

As in Aït-Sahalia (1996), the following changes of variables further simplify the problem:

$$s = \frac{C_1 S}{1 + C_1 S} \quad \text{and} \quad v = \frac{C_2 V}{1 + C_2 V},$$

where C_1 and C_2 are positive constants chosen to over-represent the range of stock prices and variance estimates in the sample. I set $g(s, v, \tau) \equiv f(S, V, \tau)$, so that $\frac{\partial f}{\partial \tau} = \frac{\partial g}{\partial \tau}$, $\frac{\partial f}{\partial S} = \frac{\partial g}{\partial s}\frac{ds}{dS}$, $\frac{\partial^2 f}{\partial S^2} = \frac{\partial^2 g}{\partial s^2}\left[\frac{ds}{dS}\right]^2 + \frac{\partial g}{\partial s}\frac{d^2 s}{dS^2}$, and similarly for $\frac{\partial f}{\partial V}$ and $\frac{\partial^2 f}{\partial V^2}$.

It is easy to check that $g(s, v, \tau) : [0,1] \times [0,1] \times [0, T\,] \to \mathbf{R}^+$ solves a parabolic equation similar to (4). To compute its numerical solution, I use two partitions of the interval $[0, 1]$,

$$0 = s_1 < s_2 < \ldots < s_J = 1 \quad \text{and} \quad 0 = v_1 < v_2 < \ldots < v_I = 1,$$

with constant subintervals Δj and Δi, and a partition of the interval $[0, T\,]$,

$$0 = t_1 < t_2 < \ldots < t_M = T,$$

with constant subinterval Δt. For ease of notation, I identify the points of these partitions with the indexes j, i and m, and construct the 3-dimensional grid

$$\{1, 2, \ldots, J\} \times \{1, 2, \ldots, I\} \times \{1, 2, \ldots, M\}. \tag{24}$$

Since diffusive problems are usually best treated with implicit schemes, I use the Crank-Nicolson algorithm. The numerical solution is computed on the 3-dimensional grid from the recursive, implicit equation:

$$\begin{aligned}
-\delta_\tau g^m_{j,i} + D_1(j,i)\,\frac{\delta^2_s g^m_{j,i} + \delta^2_s g^{m-1}_{j,i}}{2} + D_2(j,i)\,\frac{\delta^2_v g^m_{j,i} + \delta^2_v g^{m-1}_{j,i}}{2} + D_3(j,i)\,\frac{\delta_{sv} g^m_{j,i} + \delta_{sv} g^{m-1}_{j,i}}{2} + \\
-r\frac{g^m_{j,i} + g^{m-1}_{j,i}}{2} + D_4(j,i)\,\frac{\delta_s g^m_{j,i} + \delta_s g^{m-1}_{j,i}}{2} + D_5(j,i)\,\frac{\delta_v g^m_{j,i} + \delta_v g^{m-1}_{j,i}}{2} = 0,
\end{aligned} \tag{25}$$

where:

$$\begin{aligned}
\delta_\tau g^m_{j,i} &= \frac{g(j,i,m) - g(j,i,m-1)}{\Delta t}, \\
\delta_s g^m_{j,i} &= \frac{g(j+1,i,m) - g(j-1,i,m)}{2\Delta j}, \\
\delta^2_s g^m_{j,i} &= \frac{g(j+1,i,m) - 2g(j,i,m) + g(j-1,i,m)}{\Delta j^2}, \\
\delta_{sv} g^m_{j,i} &= \frac{g(j+1,i+1,m) - g(j+1,i-1,m) - g(j-1,i+1,m) + g(j-1,i-1,m)}{4\Delta j \Delta i}.
\end{aligned}$$

The terms $D_1(j,i),\ldots,D_5(j,i)$ are the coefficients of the parabolic equation that characterizes g, evaluated at the points of the grid. The numerical derivatives $\delta_v g^m_{j,i}$ and $\delta^2_v g^m_{j,i}$ are computed in a similar fashion.

The option price is computed recursively from (25); for any $m = 1,\ldots,M-1$, the set of prices $\{g(j,i,m+1),\, j=1,\ldots,J,\, i=1,\ldots,I\}$ is obtained as a function of the prices in the previous period, $\{g(j,i,m), j=1,\ldots,J, i=1,\ldots,I\}$. Convergence of the numerical solution to the actual solution of the pde is guaranteed as $\Delta j \to 0$, $\Delta i \to 0$ and $\Delta t \to 0$; see, e.g., Ames (1977).

C Proof of Proposition 1

Asymptotic results are obtained writing the estimation problem of the vector $(\psi\ \lambda)$, $\psi \equiv (\mu\ \alpha\ \beta\ \eta\ \rho)$, in a Method of Moments form, so that standard asymptotics from GMM and SMM theory can be applied. See, e.g., Newey (1984) for a discussion of a method of moments interpretation of multiple-step estimators, and Heaton (1995) for an application of this methodology.

In the first step of the estimation, $\hat{\psi}$ is obtained via EMM from a set of moment conditions $Q_1(\psi) = 0$.

On the other hand, two different methods are used to estimate λ.

In the first approach, $\hat{\lambda}$ is obtained by minimizing the SSR between market and stochastic volatility option prices, as indicated in (14). Hence, $(\psi\ \lambda)$ satisfies the first order condition $Q_2(\psi,\lambda) = 0$, where

$$Q_2(\psi,\lambda) \equiv \sum_{n=1}^{N}\sum_{i=1}^{I} \frac{\partial f(S_{t_n},\hat{V}_{t_n},\tau_n;\psi,\lambda,K_i)}{\partial\lambda}\Big(f(S_{t_n},\hat{V}_{t_n},\tau_n;\psi,\lambda,K_i) - \tilde{f}(\tau_n,K_i)\Big)/S^2_{t_n}\,.$$

In the second approach discussed above, λ is instead estimated via SMM from the moment condition $Q_2(\psi,\lambda) = 0$, where $Q_2(\psi,\lambda)$ is obtained in this case stacking the moment condition (12) for different values of the constant γ and $\tau \in [0, T]$.

Let us then denote with $Q(\psi,\lambda)$ the vector obtained by stacking the moment conditions $Q_1(\psi)$ and $Q_2(\psi,\lambda)$:

$$Q(\psi,\lambda) \equiv \begin{bmatrix} Q_1(\psi) \\ Q_2(\psi,\lambda) \end{bmatrix},$$

and let

$$\Omega \equiv \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{12}' & \Omega_{22} \end{bmatrix}$$

be the variance matrix of $Q(\psi,\lambda)$, partitioned according to the decomposition Q_1–Q_2. Finally, let D be the matrix of the derivatives of Q with respect to ψ and λ:

$$D \equiv \begin{bmatrix} D_{1\psi} & 0 \\ D_{2\psi} & D_{2\lambda} \end{bmatrix},$$

where $D_{1\psi} \equiv D_\psi Q_1(\psi)$, $D_{2\psi} \equiv D_\psi Q_2(\psi,\lambda)$ and $D_{2\lambda} \equiv D_\lambda Q_2(\psi,\lambda)$.

Then, the sequential estimator $(\hat{\psi}\ \hat{\lambda})$ solves:

$$A_N\, Q^N(\psi,\lambda) = 0,$$

where $Q^N(\psi,\lambda)$ is the sample counterpart of the vector of moments $Q(\psi,\lambda)$ and the matrix A_N selects the moment conditions which are set equal to zero. Assume that A_N converges to a matrix A with probability one as $N\to\infty$; given the nature of the multiple-step problem, the matrix A has the structure

$$A=\begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix},$$

where A_1 is the matrix used to select the moment conditions in the EMM estimation of the stochastic volatility model, and A_2 selects the moment conditions used in the second step of the estimation. In particular, A_2 is equal to 1 in the case that we use the approach discussed in Section 5.5, since in this case A_2 coincides with the bottom right element of A and sets the last component of Q equal to zero. Otherwise, A_2 is chosen efficiently, as discussed in Hansen (1982).

Under the assumptions and regularity conditions discussed in Gallant and Tauchen (1996) and Duffie and Singleton (1993),

$$\begin{bmatrix} \sqrt{N}(\hat{\psi}-\psi) \\ \sqrt{N}(\hat{\lambda}-\lambda) \end{bmatrix} \xrightarrow{D} N(0,H),$$

with $H=(1+\xi)\ (AD)^{-1}A\Omega A'(AD)^{-1\prime}$, where the multiplicative factor $(1+\xi)$, $\xi=\lim_{N\to\infty}\frac{N}{T(N)}$, summarizes the effect of the simulation; see, e.g., Gallant and Tauchen (1996) and Gourieroux, Monfort and Renault (1993).

Computing explicitly the elements of the matrix H, we get

$$H_{11}=(1+\xi)\ (A_1D_{1\psi})^{-1}A_1\Omega_{11}A_1'(A_1D_{1\psi})^{-1\prime}.$$

The matrix A_1 was chosen optimally in the EMM estimation of ψ, and, as discussed in Hansen (1982), was set equal to $D_{1\psi}'\Omega_{11}^{-1}$. Substituting into H_{11}, we get

$$H_{11}=(1+\xi)\ (D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1},$$

which is the EMM asymptotic variance matrix of $\hat{\psi}$, as in Gallant and Tauchen (1996).

Consider first the case where λ is estimated by SMM, as discussed in Section 5.4. Then, substituting $A_1=D_{1,\psi}'\Omega_{11}^{-1}$ into H_{22} and H_{21}, we get

$$\begin{aligned} H_{21} &= (1+\xi)\ (A_2D_{2\lambda})^{-1}A_2[-D_{2\psi}+\Omega_{12}'\Omega_{11}^{-1}D_{1\psi}](D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1}, \\ H_{22} &= (1+\xi)\ (A_2D_{2\lambda})^{-1}A_2\Omega_\star A_2'(A_2D_{2\lambda})^{-1\prime}, \end{aligned} \tag{26}$$

with

$$\begin{aligned} \Omega_\star \equiv\ & [D_{2\psi}(D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1}D_{2\psi}'-\Omega_{12}'\Omega_{11}^{-1}D_{1\psi}(D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1}D_{2\psi}'- \\ & D_{2\psi}(D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1}D_{1\psi}'\Omega_{11}^{-1}\Omega_{12}+\Omega_{22}]. \end{aligned}$$

The efficient choice of moments in the second step estimation is obtained by setting $A_2=D_{2\lambda}'\Omega_\star^{-1}$; substituting into (26) we obtain

$$\begin{aligned} H_{21} &= (1+\xi)\ (D_{2\lambda}'\Omega_\star^{-1}D_{2\lambda})^{-1}D_{2\lambda}'\Omega_\star^{-1}[-D_{2\psi}+\Omega_{12}'\Omega_{11}^{-1}D_{1\psi}](D_{1\psi}'\Omega_{11}^{-1}D_{1\psi})^{-1}, \\ H_{22} &= (1+\xi)\ (D_{2\lambda}'\Omega_\star^{-1}D_{2\lambda})^{-1}. \end{aligned}$$

Suppose next that λ is estimated as discussed in Section 5.5. In this case we have:

$$\begin{aligned} H_{21} &= \sqrt{1+\xi}\,(A_2 D_{2\lambda})^{-1} A_2[-D_{2\psi} + \Omega_{12}'\Omega_{11}^{-1} D_{1\psi}](D_{1\psi}'\Omega_{11}^{-1} D_{1\psi})^{-1}\,, \\ H_{22} &= (A_2 D_{2\lambda})^{-1} A_2 \Omega_\star A_2' (A_2 D_{2\lambda})^{-1\,\prime}\,. \end{aligned} \tag{27}$$

Substituting $A_2 = 1$ into (27) we get

$$\begin{aligned} H_{21} &= \sqrt{1+\xi}\;D_{2\lambda}^{-1}\,[\,-D_{2\psi} + \Omega_{12}'\Omega_{11}^{-1} D_{1\psi}\,]\,(D_{1\psi}'\Omega_{11}^{-1} D_{1\psi})^{-1}\,, \\ H_{22} &= D_{2\lambda}^{-2}\,\Omega_\star\,. \end{aligned} \tag{28}$$

Furthermore, by the mean value theorem we have:

$$\sqrt{N}[f(S,V,\tau\,;\hat{\psi},\hat{\lambda}) - f(S,V,\tau\,;\psi,\lambda)] = D_{\psi,\lambda} f(S,V,\tau;\tilde{\psi},\tilde{\lambda}) \begin{bmatrix} \sqrt{N}(\hat{\psi}-\psi) \\ \sqrt{N}(\hat{\lambda}-\lambda) \end{bmatrix}$$

where $0 \le \tilde{\psi} \le \hat{\psi}$, $0 \le \tilde{\lambda} \le \hat{\lambda}$, and $D_{\psi,\lambda} f(S,V,\tau;\psi,\lambda)$ denotes the vector of derivatives of the option price with respect to the parameter vector $(\psi\ \lambda)$.

Consistency of the $(\hat{\psi}\ \hat{\lambda})$ estimator yields that, with probability 1,

$$D_{\psi,\lambda} f(S,V,\tau;\tilde{\psi},\tilde{\lambda}) \to D_{\psi,\lambda} f(S,V,\tau;\psi,\lambda) \equiv \Lambda\,,$$

as $N \to \infty$. Therefore, by Slutsky theorem we conclude that:

$$\sqrt{N}\,[\,f(S,V,\tau;\hat{\psi},\hat{\lambda}) - f(S,V,\tau;\psi,\lambda)] \xrightarrow{D} N(0,\Lambda' H \Lambda).$$

□

D The Heston formula

Given the risk-adjusted model (19)-(20), Heston (1993) obtains the following closed-form formula for the European call option price:

$$f(S_t,V_t,\tau;K) = \mathrm{e}^{-d\tau} S_t\,P_1 + \mathrm{e}^{-r\tau} K P_2\,,$$

where

$$\begin{aligned} P_j(X,V,\tau;\ln(K)) &= \frac{1}{2} + \frac{1}{\pi}\int_0^\infty \mathrm{Real}\left(\frac{\mathrm{e}^{-i\Phi\ln(K)} F_j(X,V,\tau;\Phi)}{i\,\Phi}\right) d\,\Phi\,, \\ F_j(X,V,\tau;\Phi) &= \exp\{C(\tau;\Phi) + D(\tau;\Phi)\,V + i\,\Phi X\}\,, \\ C(\tau;\Phi) &= r\,\Phi i\tau + \frac{\alpha}{\eta^2}\left\{(\beta_j - \rho\eta\Phi i + \gamma)\tau - 2\,\ln\left(\frac{1-g\mathrm{e}^{\gamma\tau}}{1-g}\right)\right\}\,, \\ D(\tau;\Phi) &= \frac{\beta_j - \rho\eta\Phi i + \gamma}{\eta^2}\left(\frac{1-\mathrm{e}^{\gamma\tau}}{1-g\mathrm{e}^{\gamma\tau}}\right)\,, \\ g &= \frac{\beta_j - \rho\eta\Phi i + \gamma}{\beta_j - \rho\eta\Phi i - \gamma}\,, \\ \gamma &= \sqrt{(\rho\eta\Phi i - \beta_j)^2 - \eta^2(2\,z_j\,\Phi i - \Phi^2)}\,, \\ z_1 &= 1/2,\ \ z_2 = -1/2,\ \ \beta_1 = \beta + \lambda - \rho\,\eta,\ \ \beta_2 = \beta + \lambda,\ \ X = \ln(\mathrm{e}^{-d\tau} S_t)\,, \end{aligned}$$

and r, d are respectively the instantaneous risk-free interest rate and dividend yield on the underlying stock.

Tables and Figures

Table 1: Summary statistics. Data on daily rates of return of the S&P 500 index, 01/02/1953 to 12/31/1996 (N=11,076 observations). All figures expressed on a daily basis in percentage form.

Mean	0.0301
Std. Dev.	0.8324
Skewness	-2.0353
Kurtosis	60.6019

Autocorrelation of Returns:

1st	*2nd*	*3rd*	*4th*	*5th*	*6th*
0.1240	-0.0320	-0.0084	-0.0056	0.0222	-0.0131

Augmented Dickey Fuller test for the presence of unit roots. The test is based on the regression:

$$\Delta X_t = \mu + \delta t + \gamma X_{t-1} + \sum_{j=1}^{12} \psi_j \Delta X_{t-j} + \varepsilon_t.$$

	S&P 500 daily prices	*S&P 500 daily returns*
Augmented D. F.	2.85	-29.65
5% critical value	-3.41	-3.41
1% critical value	-3.96	-3.96

Table 2: SNP model estimates. Data on daily rates of return of the S&P 500 index, 01/02/1953 to 12/31/1996, filtered using an MA(1) model (N=11,076 observations). Parameter estimates are expressed in percentage form on a daily basis and refer to the following model:

$$f_K(r_t|x_t;\xi) = \left(\nu + (1-\nu\,) \times \frac{[P_K(z_t,x_t)]^2}{\int_{\mathrm{R}}[P_K(z_t,x_t)]^2\phi(u)du}\right)\frac{\phi(z_t)}{\sqrt{h_t}}\,, \qquad \nu = 0.01,$$

where $\phi(.)$ is the standard normal density,

$$\begin{aligned}
z_t &= \frac{r_t - \mu_t}{\sqrt{h_t}}\,, \\
\mu_t &= \phi_0 + c\,h_t\,, \\
\ln h_t &= \omega + \sum_{i=1}^{p}\beta_i \ln h_{t-i} + (1+\alpha_1 L + \ldots + \alpha_q L^q)\,[\,\theta_1 z_{t-1} + \theta_2\,(b(z_{t-1}) - \sqrt{2/\pi})\,]\,, \\
b(z) &= |z| \;\text{ for } |z| \geq \pi/2K, \quad b(z) = (\pi/2 - \cos(Kz))/K \;\text{ for } |z| < \pi/2K\,, \quad K = 100\,, \\
P_K(z,x) &= \sum_{i=0}^{K_z} a_i(x) z^i = \sum_{i=0}^{K_z}\left(\sum_{|j|=0}^{K_x} a_{ij}x^j\right) z^i\,, \qquad a_{00} = 1\,.
\end{aligned}$$

Parameter	*Estimate*	*Standard Error*
ϕ_0	0.0546	0.0394
c	0.0315	0.0331
ω	3.5526	1.4211
α	-0.4367	0.0624
β_1	0.9880	0.0028
θ_1	-0.1407	0.0304
θ_2	0.3003	0.0269
a_{10}	-0.0489	0.0495
a_{20}	-0.2480	0.0314
a_{30}	-0.0021	0.0242
a_{40}	0.1213	0.0195
a_{50}	-0.0177	0.0161
a_{60}	-0.0504	0.0100
a_{70}	0.0087	0.0100
a_{80}	0.0509	0.0097

Table 3: EMM estimates of the continuous-time, stochastic volatility models. Estimates are for the sample period 01/02/1953 to 12/31/1996. Standard errors are reported in brackets. Parameter estimates are expressed in percentage form on a daily basis, and refer to the following models:

$$\begin{aligned} \frac{dS_t}{S_t} &= (\mu + c\,V_t)\,dt + \sqrt{V_t}\,dW_{1,t}\,, \\ & \text{log}-\text{variance}: \\ d\ln V_t &= (\alpha - \beta \ln V_t)dt + \eta\, dW_{2,t}\,, \\ & \text{square}-\text{root}: \\ dV_t &= (\alpha - \beta\, V_t)dt + \eta\sqrt{V_t}\,dW_{2,t}\,, \\ & \text{corr}(dW_{1,t}, dW_{2,t}) = \rho. \end{aligned}$$

Parameter	log-variance:		square-root:	
	$\rho = 0$	$\rho \neq 0$	$\rho = 0$	$\rho \neq 0$
μ	0.0900 (0.0182)	0.0252 (0.0117)	0.0913 (0.0172)	0.0202 (0.0126)
c	-0.1446 (0.0551)	0.0149 (0.0255)	-0.1484 (0.0523)	0.0256 (0.0282)
α	-0.0076 (0.0037)	-0.0136 (0.0030)	0.0027 (0.0012)	0.0081 (0.0016)
β	0.0074 (0.0034)	0.0160 (0.0030)	0.0068 (0.0031)	0.0156 (0.0032)
η	0.0524 (0.0115)	0.1206 (0.0093)	0.0308 (0.0065)	0.0782 (0.0074)
ρ		-0.5778 (0.0433)		-0.5973 (0.0448)
χ^2 [*d.f.*] (*p-value*)	116.98 [10] ($< 10^{-5}$)	31.53 [9] (0.00024)	116.21 [10] ($< 10^{-5}$)	30.74 [9] (0.00033)

Table 4: SMM estimate of λ. The estimation is performed using 3 moment conditions, corresponding to the maturities of 3, 4 and 5 weeks on European call options on the S&P 500 index, sampled from 1986 to 1996.

λ	*(Standard Error)*
-0.3493	(0.0926)
Test for over-identifying restrictions χ^2_2, (p-value):	
4.58	(0.1013)

Table 5: In-sample pricing errors based on stochastic volatility and Black-Scholes prices. Data on European call options on the S&P 500 index expiring in March 1997. Black-Scholes prices are evaluated at the EMM estimate of the returns standard deviation, σ_{EMM}, and at the implied standard deviation which minimizes the sum of squared residuals between Black-Scholes and market prices, σ_{impl}. Stochastic volatility prices are evaluated for both the log-variance and the Heston model. Variance estimates are obtained from S&P 500 returns (Kalman filter and reprojection method) and option prices (SV implied volatilities.) Pricing errors are computed in the form of sum of squared residuals, normalized by the index level (*SSR*), average dollar absolute deviation from market prices (*DEV*), average dollar absolute deviation from market prices, normalized by the option market prices (*DEV%*).

	SSR	*DEV*	*DEV%*
Black-Scholes, σ_{EMM}	10.34×10^{-4}	\$1.12	11.08%
Black-Scholes, σ_{impl}	9.41×10^{-4}	\$1.05	11.59%
SV log-v Model, Kalman Filter, $\lambda = 0$	31.10×10^{-4}	\$2.16	22.73%
SV log-v Model, Kalman Filter	9.35×10^{-4}	\$1.03	13.66%
SV log-v Model, reprojection method, $\lambda = 0$	37.54×10^{-4}	\$2.36	23.45%
SV log-v Model, reprojection method	10.98×10^{-4}	\$1.06	13.10%
SV log-v Model, Implied Volatilities, $\lambda = 0$	3.65×10^{-4}	\$0.59	6.55%
SV log-v Model, Implied Volatilities	3.63×10^{-4}	\$0.59	6.51%
SV log-v Model, $\rho = 0$, *Kalman Filter*	14.84×10^{-4}	\$1.40	14.77%
SV log-v Model, $\rho = 0$, *reprojection method*	15.73×10^{-4}	\$1.32	16.18%
SV log-v Model, $\rho = 0$, *Implied Volatilities*	4.50×10^{-4}	\$0.70	8.00%
SV Heston Model, Implied Volatilities, $\lambda = 0$	3.74×10^{-4}	\$0.60	6.54%
SV Heston Model, Implied Volatilities	3.67×10^{-4}	\$0.59	6.53%

Table 6: Estimate of λ obtained minimizing the sum of squared residuals between market and stochastic volatility prices. Data on European call options on the S&P 500 index expiring in March 1997, closing option prices sampled daily from December 30, 1996, to March 14, 1997. Variance estimates are obtained from S&P 500 returns (Kalman filter method).

λ	*(Standard Error)*
-0.1372	(0.1418)

Table 7: Out-of-sample pricing errors based on stochastic volatility and Black-Scholes prices. Data on European call options on the S&P 500 index expiring in June 1997. Black-Scholes prices are evaluated at the EMM estimate of the returns standard deviation, σ_{EMM}, and at the implied standard deviation which minimizes the sum of squared residuals between Black-Scholes and market prices, σ_{impl}. Stochastic volatility prices are evaluated for both the log-variance and the Heston model. Variance estimates are obtained from S&P 500 returns (Kalman filter and reprojection method) and option prices (day-before and same-day SV implied volatilities.) Pricing errors are computed in the form of sum of squared residuals, normalized by the index level (*SSR*), average dollar absolute deviation from market prices (*DEV*), average dollar absolute deviation from market prices, normalized by the option market prices (*DEV%*).

	SSR	*DEV*	*DEV%*
Black-Scholes, σ_{EMM}	15.56×10^{-4}	\$1.51	15.55%
Black-Scholes, σ_{impl}	12.05×10^{-4}	\$1.33	17.02%
SV log-v Model, Kalman Filter	24.74×10^{-4}	\$1.94	31.63%
SV log-v Model, reprojection method	24.55×10^{-4}	\$1.93	47.94%
SV log-v Model, Day-Before Imp. Vol.	5.99×10^{-4}	\$0.86	11.01%
SV log-v Model, Same-Day Imp. Vol.	3.06×10^{-4}	\$0.54	8.13%
SV Heston Model, Day-Before Imp. Vol.	6.57×10^{-4}	\$0.92	11.09%
SV Heston Model, Same-Day Imp. Vol.	3.74×10^{-4}	\$0.64	7.01%

Table 8: Linear projection of the log-variance SV implied standard deviations over the return-based estimates of volatilities: Kalman filter estimates (KF) and reprojected volatility (RPJ). Standard errors, robustified for the presence of heteroskedasticity and autocorrelation, in brackets.

$$\text{Imp. Vol.} = b_1 + b_2 \times \text{Vol. Est.} + \varepsilon.$$

Parameter	*Volatility Estimate*	
	KF	*RPJ*
b_1	0.0614	0.0829
	(0.0151)	(0.0183)
b_2	0.5062	0.3507
	(0.1147)	(0.1555)
R^2	0.20	0.09

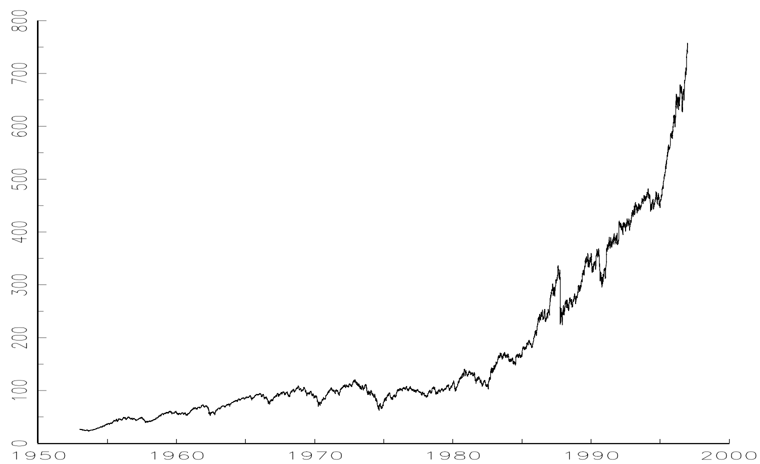

Panel 1: S&P 500 price index.

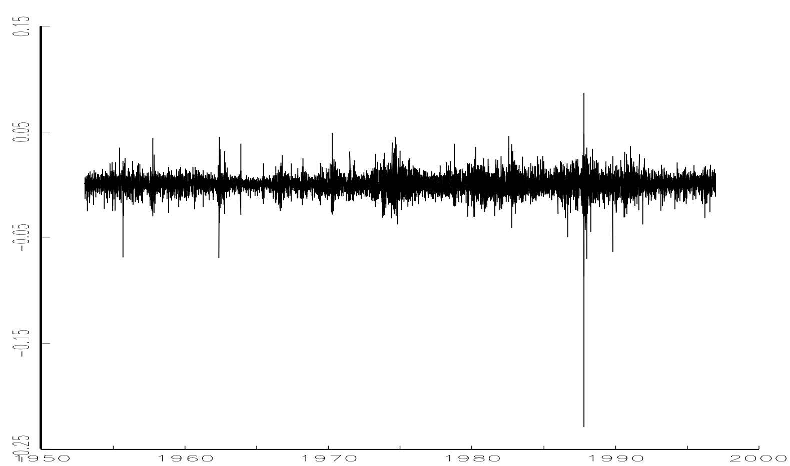

Panel 2: S&P 500 daily rate of return.

Figure 1: S&P 500 prices and returns: 01/02/1953 to 12/31/1996.

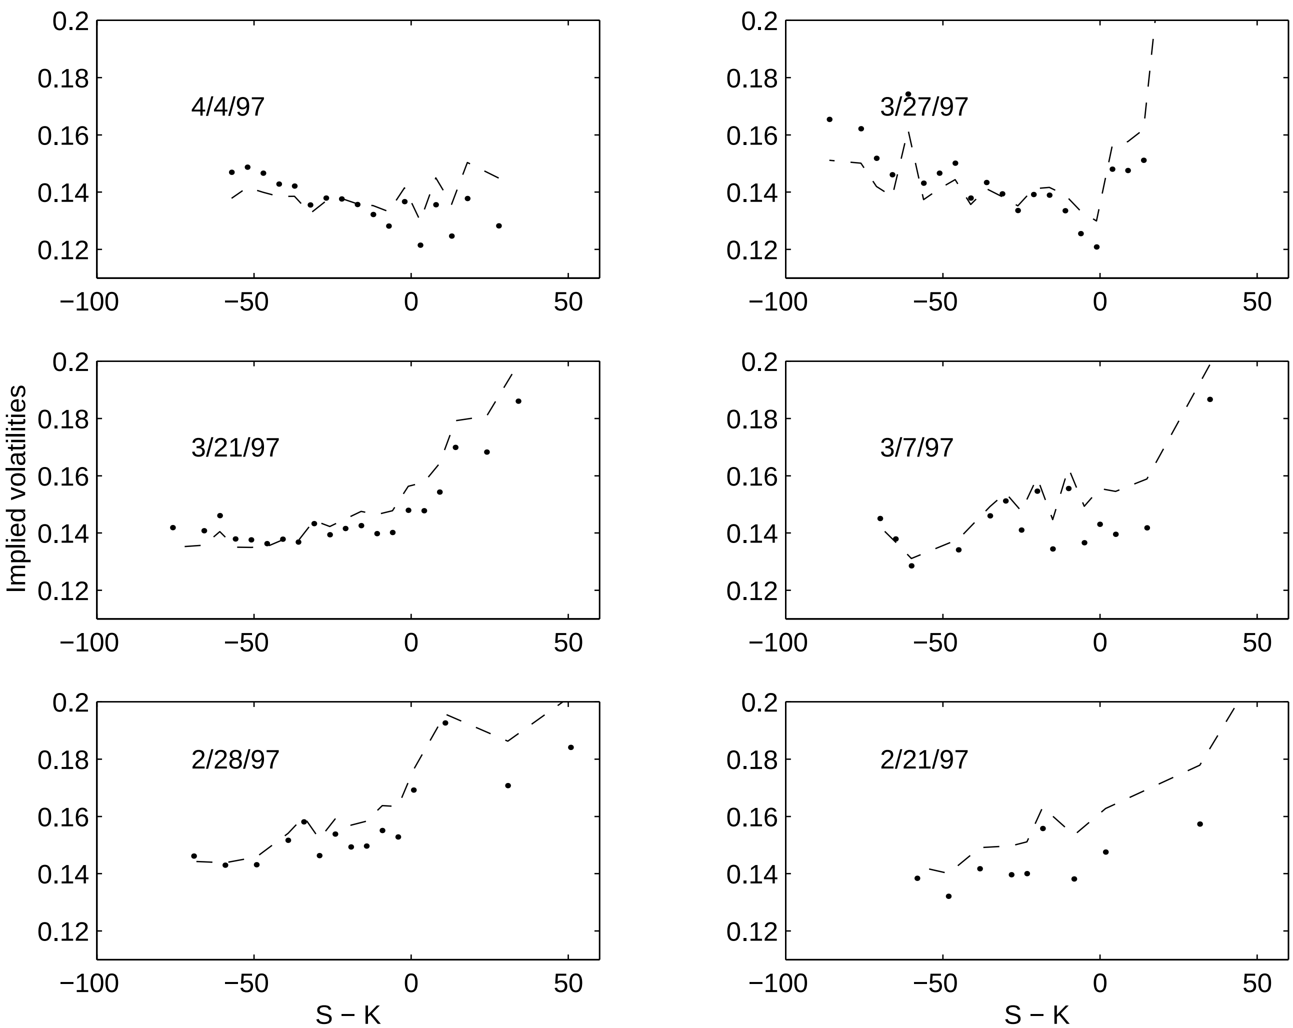

Figure 2: SV and Black-Scholes implied volatilities. Daily standard deviations are converted to an annual frequency multiplying by $\sqrt{252}$. Log-variance SV implied volatilities are computed numerically assuming a market price for variance risk $\lambda = -0.1335$. Implied volatilities are computed using call options expiring in April 1997.

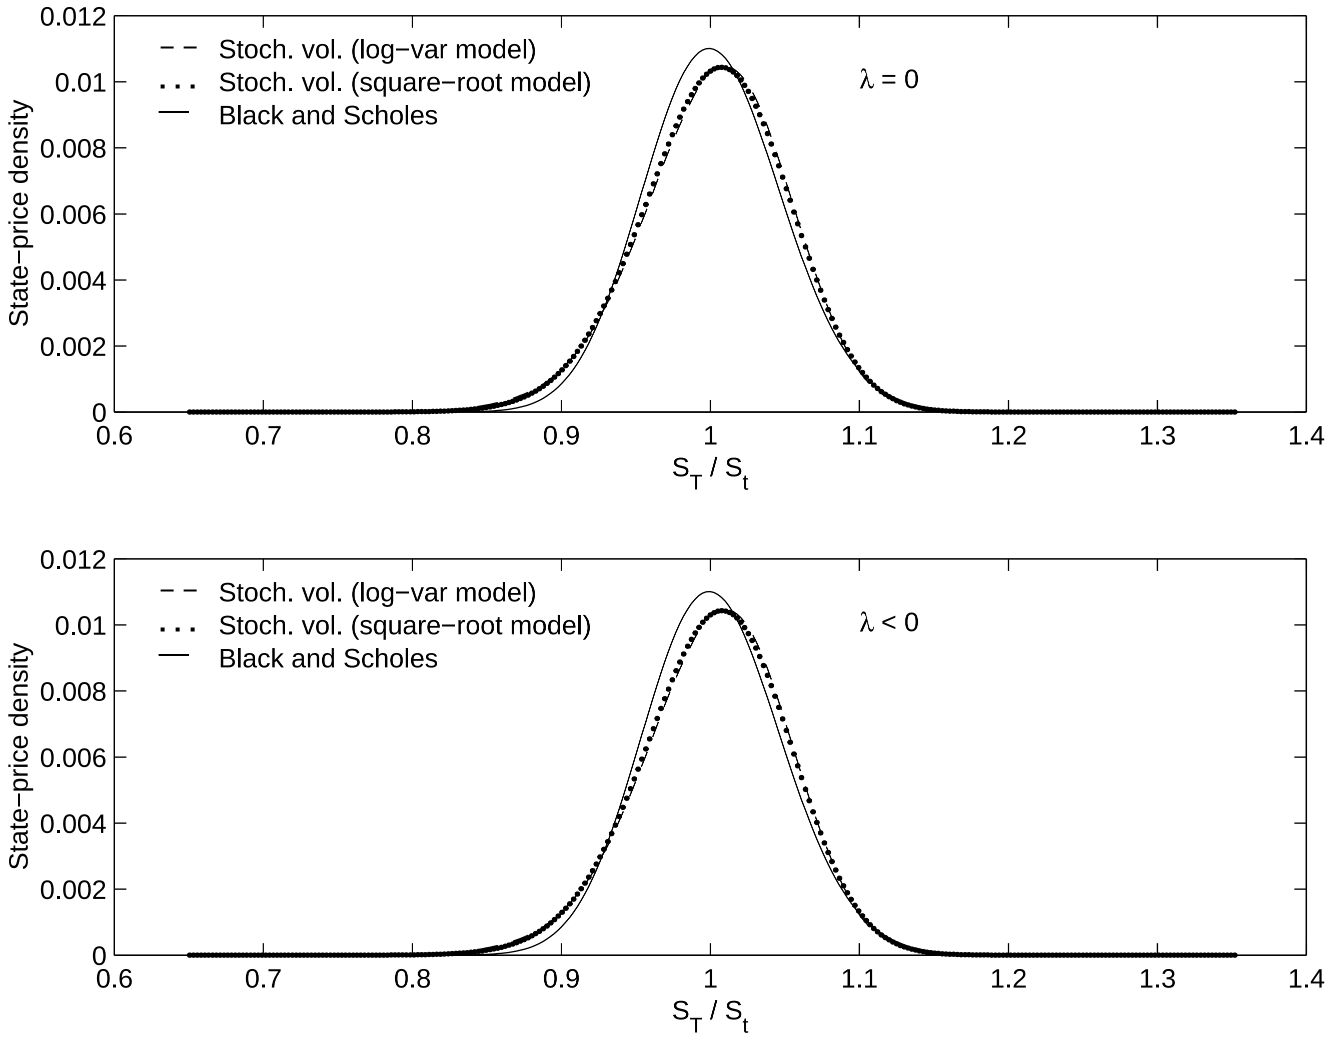

Figure 3: State-price densities on March 21, 1997, relative to options expiring in April 1997. Black-Scholes, log-variance and square-root models. In the log-variance and square-root models, SPDs are evaluated at the same-day SV implied volatilities. The risk premium for variance risk is first set equal to zero (top panel), and then to its estimated value (bottom): $\lambda = -0.1335$ (log-variance model), $\lambda = -0.0278$ (square-root model).

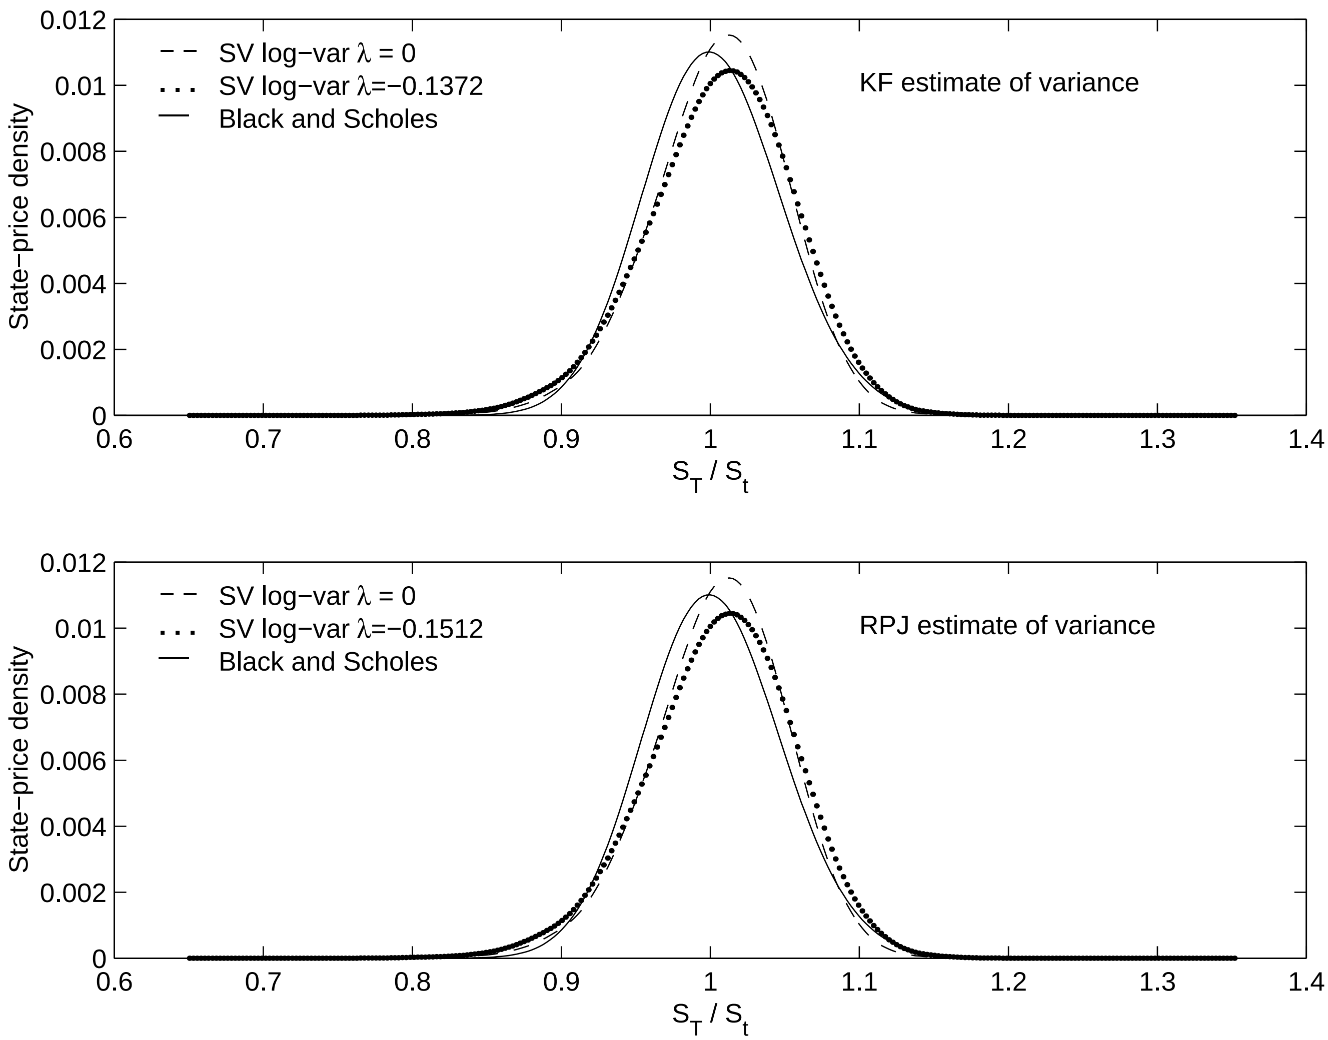

Figure 4: State-price densities on March 21, 1997 relative to options expiring in April 1997. Black-Scholes and log-variance models. The market price of variance risk is set equal to 0 first, and then to its estimated value. In the log-variance model, SPDs are evaluated using variance estimates based on the Kalman filter (top panel) and the reprojection method (bottom).

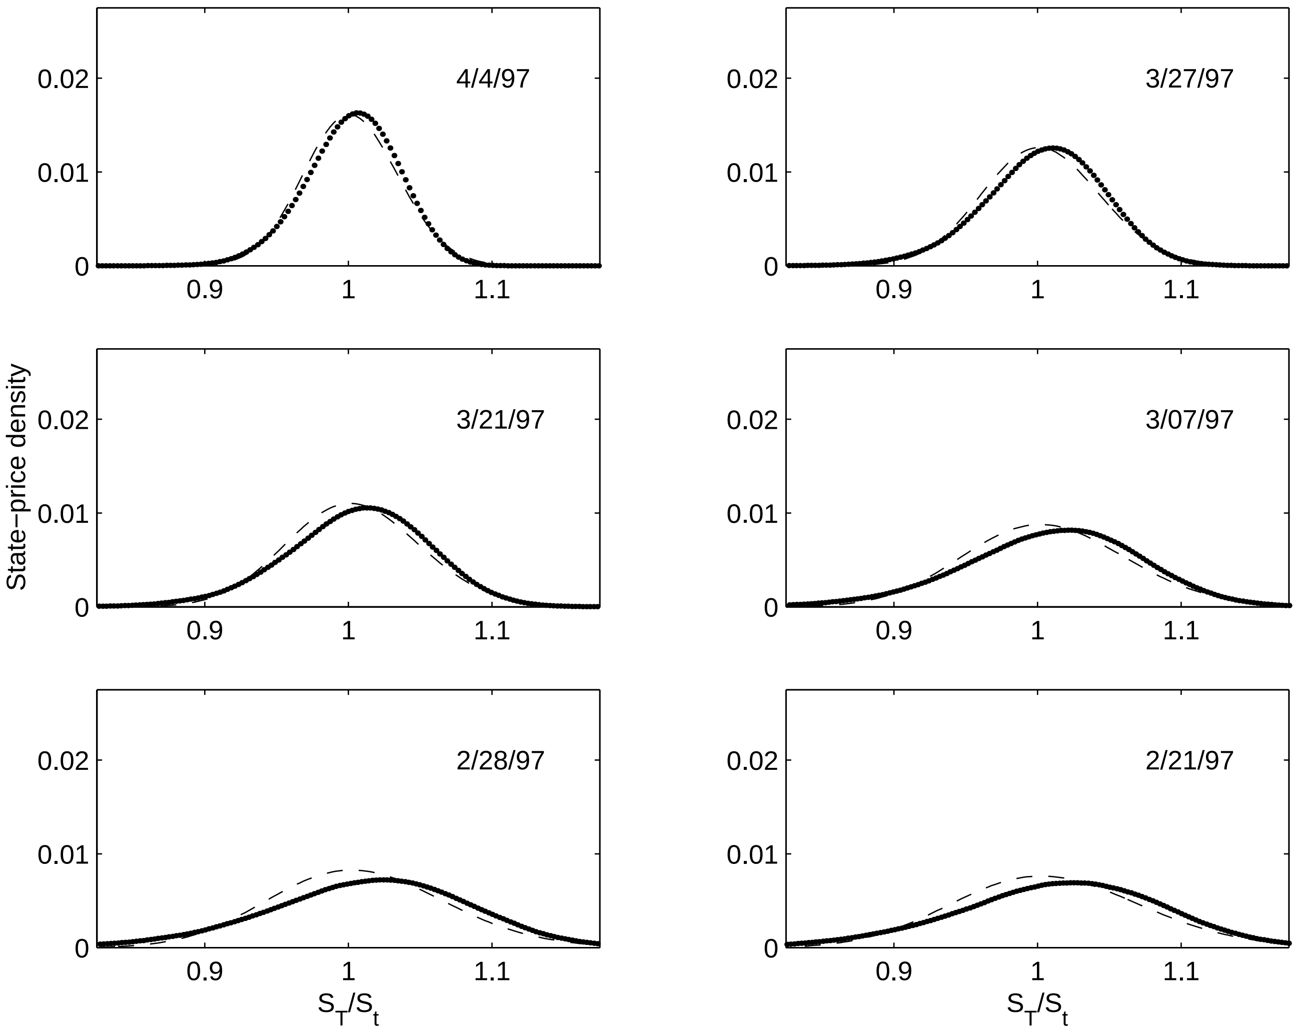

Figure 5: Black-Scholes (- -) and stochastic volatility, log-variance model, (**...**) state-price densities for options expiring April 1997.

References

Aït-Sahalia, Y. (1996): "Nonparametric Pricing of Interest Rate Derivative Securities," *Econometrica*, vol. 64, p. 527-560.

Aït-Sahalia, Y. and A. W. Lo (1998): "Nonparametric Estimation of State-Price Densities Implicit In Financial Asset Prices," *Journal of Finance*, vol. 53, n. 2, p. 499-547.

Aït-Sahalia, Y., Y. Wang and F. Yared (2001): "Do Option Markets Correctly Price the Probabilities of Movements of the Underlying Asset?," *Journal of Econometrics*, vol. 102, 67-110.

Ames, W. F. (1977): "Numerical Methods for Partial Differential Equations," Academic Press.

Andersen, T. G., L. Benzoni and J. Lund (2001): "An Empirical Investigation of Continuous-Time Equity Returns Models," *Journal of Finance*, forthcoming.

Andersen, T. G. and J. Lund (1997): "Estimating Continuous-Time Stochastic Volatility Models of the Short Term Interest Rate," *Journal of Econometrics*, vol. 77, n. 2, p. 343-377.

Bakshi, G., C. Cao and Z. Chen (1997): "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, vol. 52, n. 5, p. 2003-2049.

Bakshi, G., C. Cao and Z. Chen (2000): "Pricing and Hedging Long-Term Options," *Journal of Econometrics*, vol. 94, n. 1-2, p. 277-318.

Bakshi, G., and N. Kapadia (2001), "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies*, forthcoming.

Bandi, F. M. and T. Nguyen (2000): "On the Functional Estimation of Jump-diffusion Processes," Working Paper, University of Chicago and Yale University.

Bandi, F. and P. C. B. Phillips (1998): "Econometric Estimation of Diffusion Models," Working paper, Cowles Foundation, Yale University.

Bates, D. S. (1991): "The Crash of '87: Was It Expected? The Evidence from Options Markets," *Journal of Finance*, vol. 46, n. 3, p. 1009-1044.

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

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

Bates, D. S. (2000): "Post-'87 Crash Fears in S&P 500 Futures Options," *Journal of Econometrics*, vol. 94, n. 1-2, p. 181-238.

Benzoni, L. (2001): "Essays in Empirical Finance," Ph.D. dissertation, Northwestern University.

Bollerslev, T., R. F. Engle and D. B. Nelson (1994): "ARCH Models," in *Handbook of Econometrics*, vol. 4, p. 2959-3038, eds. R. F. Engle and D. McFadden, Amsterdam: North-Holland.

Bollerslev, T. and H. Zhou (2001): "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics*, forthcoming.

Breeden, D. and R. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, vol. 51, p. 621-651.

Buraschi, A. and J. Jackwert (2001): "The Price of a Smile: Hedging and Spanning in Option Markets," *The Review of Financial Studies*, vol. 14, n. 2, p. 495-527.

Canina, L. and S. Figlewski (1993): "The Informational Content of Implied Volatility," *The Review of Financial Studies,* vol. 6, n. 3, p. 659-681.

Carrasco, M., E. Ghysels, M. Chernov, J.-P. Florens (2001): "Estimating Diffusions With a Continuum of Moment Conditions," Working Paper, University of North Carolina.

Chacko, G. and L. M. Viceira (1999): "Spectral GMM Estimation of Continuous-Time Processes," Working Paper, Harvard University.

Chernov, M. (2000): "Implied Volatilities as Forecasts of Future Volatility, Time-Varying Risk Premia and Return Variability," Working Paper, Columbia University.

Chernov, M., R. Gallant, E. Ghysels and G. Tauchen (1999): "A New Class of Stochastic Volatility Models with Jumps: Theory and Estimation," Working Paper, Penn State University, University of North Carolina, Duke University.

Chernov, M., R. Gallant, E. Ghysels and G. Tauchen (2000): "Alternative Models for Stock Price Dynamics," Working Paper, Penn State University, University of North Carolina, Duke University.

Chernov, M. and E. Ghysels (2000): "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation," *The Journal of Financial Economics*, vol. 56, n. 3, p. 407-458.

Christensen, B. and N. Prabhala (1998), "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics,* vol. 50, p. 125-150.

Conley, T. G., L. P. Hansen, E. G. J. Luttmer and J. Scheinkman (1997): "Short-Term Interest Rates as Subordinated Diffusions," *The Review of Financial Studies*, vol. 10, n. 3, p. 525-577.

Coval, J. D., and T. Shumway (2001), Epected Option Returns, *Journal of Finance* vol. 56, p. 983-1009.

Cox, J. C., J. E. Ingersoll, Jr. and S. A. Ross (1985): "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, vol. 53, p. 363-384.

Das, S. R. and R. K. Sundaram (1999): "Of Smiles and Smirks: A Term-Structure Perspective," *Journal of Financial and Quantitative Analysis*, vol. 34, p. 211-239.

Day, T. E. and C. M. Lewis (1992): "Stock Market Volatility and the Information Content of Stock Index Options," *Journal of Econometrics*, vol. 52, p. 267-287.

Duffie, D., J. Pan and K. J. Singleton (2000): "Transform Analysis and Option Pricing for Affine Jump Diffusions," *Econometrica*, vol. 68, n. 6, p. 1343-1376.

Duffie, D. and K. J. Singleton (1993): "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, vol. 61, p. 929-952.

Dumas, B., Fleming, J. and R. E. Whaley (1998): "Implied Volatility Functions: Empirical Tests," *Journal of Finance*, vol. 53, n. 6, p. 2059-2106.

Elerian, O., S. Chib and N. Shephard (1998): "Likelihood Inference for Discretely Observed Non-linear Diffusions," forthcoming *Econometrica.*

Eraker, B. (2000), "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," Working paper, University of Chicago.

Eraker, B. (2001): "MCMC Analysis of Diffusion Models with Application to Finance," *Journal of Business and Economic Statistics*, 19, 177-191.

Eraker, B., M. S. Johannes and N. G. Polson (2001): "The Impact of Jumps on Volatility and Returns," *Journal of Finance*, forthcoming.

Gallant, A. R., D. Hsieh and G. Tauchen (1997): "Estimation of Stochastic Volatility Models with Diagnostics," *Journal of Econometrics*, vol. 81, p. 159-192.

Gallant, A. R., C. Hsu and G. Tauchen (1999): "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance," *The Review of Economics and Statistics*, vol. 81, n. 4, p. 617-631.

Gallant, A. R. and J. R. Long (1997): "Estimating Stochastic Differential Equations Efficiently by Minimum Chi-Square," *Biometrika*, vol. 84, p. 125-142.

Gallant, A. R. and D. W. Nychka (1987): "Semi-Nonparametric Maximum Likelihood Estimation," *Econometrica*, vol. 55, p. 363-390.

Gallant, A. R., P. E. Rossi and G. Tauchen (1992): "Stock Prices and Volume," *Review of Financial Studies*, vol. 5, p. 199-242.

Gallant, A. R. and G. Tauchen (1996): "Which Moments to Match?," *Econometric Theory*, vol. 12, p. 657-681.

Gallant, A. R. and G. Tauchen (1997): "Estimation of Continuous-Time Models for Stock Returns and Interest Rates," *Macroeconomic Dynamics*, vol. 1, p. 137-168.

Gallant, A. R. and G. Tauchen (1998): "Reprojecting Partially Observed Systems With Application to Interest Rate Diffusions," *Journal of the American Statistical Association*, vol. 93, p. 10-24.

Ghysels, E., A. C. Harvey and E. Renault (1998): "Stochastic Volatility," in *Handbook of Statistics*, vol. 14: "Statistical Methods in Finance," p. 119-191, eds. G. S. Maddala and C. R. Rao, Amsterdam: North-Holland.

Glosten, L. R., R. Jagannathan and D. Runkle (1993): "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, vol. 48, p. 1779-1801.

Gourieroux, C., A. Monfort and E. Renault (1993): "Indirect Inference," *Journal of Applied Econometrics*, vol. 8, p. s85-s118.

Guo, D. (1998) "The Risk Premium of Volatility Implicit in Currency Options," *Journal of Business and Economic Statistics*, vol. 16, n. 4, p. 498-507.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moment Estimators," *Econometrica*, vol. 50, p. 1029-1056.

Hansen, L. P. and J. A. Scheinkman (1995): "Back to the Future: Generating Moment Implication for Continuous-Time Markov Processes," *Econometrica*, vol. 63, p. 767-804.

Harvey, A. C. and N. Shephard (1996): "Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns," *Journal of Business and Economic Statistics*, vol. 14, n. 4, p. 429-434.

Heaton, J. C. (1995): "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," *Econometrica*, vol. 63, p. 681-717.

Heston, S. (1993): "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *The Review of Financial Studies*, vol. 6, n. 2, p. 327-343.

Ho, M. S., W. R.M. Perraudin and B. E. Sørensen (1996): "A Continuous-Time Arbitrage Pricing Model With Stochastic Volatility and Jumps," *Journal of Business and Economic Statistics*, vol. 14, n. 1, p. 31-43.

Hull, J. and A. White (1987): "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, vol. XLII, n. 2, p. 281-300.

Jacquier, E. and R. Jarrow (2000): "Bayesian Analysis of Contingent Claim Model Error," *Journal of Econometrics*, vol. 94, n. 1-2, p. 145-180.

Jacquier, E., N. G. Polson and P. E. Rossi (1994): "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics*, vol. 12, n. 4, p. 371-389.

Jiang, G. J. and J. L. Knight (1997): "A Nonparametric Approach to the Estimation of Diffusion Processes, with an Application to a Short-Term Interest Rate Model," *Econometric Theory*, vol. 13, n. 5, p. 615-645.

Jiang, G. J. and J. L. Knight (1999): "Efficient Estimation of the Continuous Time Stochastic Volatility Model Via the Empirical Characteristic Function," *Journal of Business and Economic Statistics*, forthcoming.

Jiang, G. and P. van der Sluis (1999): "Index Option Pricing Models with Stochastic Volatility and Stochastic Interest Rates," *European Finance Review*, vol. 3, n. 3, p. 273-310.

Johannes, M. (1999): "A Non-Parametric Approach to Jumps in Interest Rates," Working Paper, Columbia University.

Johannes, M., R. Kumar and N. G. Polson (1998): "State Dependent Jump Models: How do US Equity Markets Jump?," Working Paper, University of Chicago.

Jones, C. S. (1998): "A Simple Bayesian Approach to the Analysis of Markov Diffusion Processes," Working Paper, The Wharton School, University of Pennsylvania.

Jones, C. S. (1999): "The Dynamics of Stochastic Volatility: Evidence from Underlying and Option Markets," Working Paper, Rochester University.

Jones, C. S. (2001): "A Nonlinear Factor Analysis of S&P 500 Index Option Returns," Working Paper, Rochester University.

Kapadia, N. (1998): "Do Equity Options Price Volatility Risk? Empirical Evidence," Working Paper, University of Massachusetts.

Kim, S., N. Shephard and S. Chib (1998): "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, vol. 65, p. 361-393.

Kloeden, P. E. and E. Platen (1992): "Numerical Solutions of Stochastic Differential Equations," Springer-Verlag.

Lamoureux, C. G. and W. D. Lastrapes (1993): "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities," *Review of Financial Studies*, vol. 6, n. 2, p. 293-326.

Liu, M. and H. H. Zhang (1997): "Specification Tests in the Efficient Method of Moments Framework with Application to the Stochastic Volatility Models," Working Paper, The Chinese University of Hong Kong and Carnegie Mellon University.

Lo, A. (1988): "Maximum Likelihood Estimation of Generalized Itô Processes With Discretely Sampled Data," *Econometric Theory*, vol. 4, p. 231-247.

Lo, A. and A. C. MacKinlay (1990): "An Econometric Analysis of Nonsynchronous Trading," *Journal of Econometrics*, vol. 45, p. 181-211.

Meddahi N. (2001): "An Eigenfunction Approach for Volatility Modeling," Working Paper, Univeristé de Montréal.

Melino A. and S. M. Turnbull (1990): "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, vol. 45, p. 239-265.

Merton, R. C. (1973): "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, n. 4 , p. 141-183.

Naik, V. (1993): "Option Valuation and Hedging Strategies with Jumps in the Volatility of Asset Returns," *Journal of Finance*, vol. 48, n. 5, p. 1969-1984.

Nelson, D. B. (1991): "Conditional Heteroskedasticity in Asset Returns: a New Approach," *Econometrica*, vol. 59, n. 2 , p. 347-370.

Newey, W. (1984): "A Method of Moments Interpretation of Sequential Estimators," *Economics Letters*, vol. 14, p. 201-206.

Newey, W. and K. West (1987): "A Simple Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, vol. 55, p. 703-708.

Pan, J. (2001), "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics*, forthcoming.

Pastorello, S., E. Renault and N. Touzi (2000): "Statistical Inference for Random Variance Option Pricing," *Journal of Business and Economic Statistics*, vol. 18, p. 358-367.

Pham, H. and N. Touzi (1997): "Equilibrium State Prices in a Stochstic Volatility Model," *Mathematical Finance*, vol. 6, n. 2, p. 215-236.

Poteshman, A. M. (1998): "Estimating a General Stochastic Variance Model from Option Prices," Working Paper, University of Illinois at Urbana-Champaign.

Poteshman, A. M. (2000): "Forecasting Future Varience from Option Prices," Working Paper, University of Illinois at Urbana-Champaign.

Renault, E. (1996): "Econometric Models of Option Pricing Errors," Working Paper, Université des Sciences Sociales, Toulouse.

Rubinstein, M. (1994): "Implied Binomial Trees," *Journal of Finance*, vol. 49, p. 771-818.

Santa-Clara, P. and S. Yan (2001): "Risk and Return in the Stock Market: Lessons from Option Data," Working Paper, UCLA and University of Arizona.

Singleton, K. J. (2001): "Estimation of Affine Asset Pricing Models Using the Empirical Characteristic Function," *Journal of Econometrics,* vol. 102, n.1, p. 111-141.

Stanton R. H. (1997): "A Nonparametric Model of Term Structure Dynamics and the Market Price of Interest Rate Risk," *Journal of Finance*, vol. 52, p. 1973-2002.

van der Sluis, P. (1997): "Computationally Attractive Stability Tests for the Efficient Method of Moments," Working Paper, University of Amsterdam.

Willmott, P., J. Dewinne and S. Howison (1993): "Option Pricing: Mathematical Models and Computation," Oxford Financial Press.