1 Estimation of Stochastic Volatility Models for the Purpose of Option Pricing

Mikhail Chernov and Eric Ghysels

In this paper we review the recent advances on the estimation of stochastic volatility (henceforth SV) models for the purpose of option pricing. For SV models, it is impossible to estimate the risk-neutral parameters using only the data from the underlying asset pricing process. We focus on the recent attempts which try to exploit optimally the information in the panel data of options and the information in the underlying fundamental. The fulfillment of this goal, which we have not yet fully accomplished, yields risk neutral parameters as well as the temporal dynamics under both measures.

Introduction

The literature on the subject of stochastic volatility models and on option pricing is huge. There are several recent surveys, notably by Bates (1996b) and Ghysels *et al.* (1996), which cover many of the early and more recent developments. This paper complements the recent literature and focuses on one very specific, and very important topic, namely the estimation of SV models for the purpose of option pricing. It raises several issues, some old and some new ones, which have been addressed with new methods of estimation. In order to implement option pricing formula one has to know the estimates of the parameters under the risk-neutral measure. In the world of Black and Scholes (1973) the no arbitrage portfolio argument determines an exact mapping between the so called physical measure and the risk neutral one. Both the risk neutral and physical parameters can be recovered from observing the dynamics of the underlying fundamental process. For SV models, however, the preference structure of agents is required to make a transition to the risk-neutral world because volatility is not a traded asset. It is therefore impossible to estimate the risk-neutral parameters using only the data from the asset pricing process underlying the option contracts.

The estimation of SV models typically involves the underlying fundamental asset prices. Estimation methods such as Generalized, Simulated and Efficient Methods of Moments (respectively GMM, SMM and EMM), Bayesian MCMC, as well as other simulated likelihood procedures were applied to a multitude of models belonging to the SV class. The differences in model specification, like the time structure (discrete versus. continuous) and the specification of the volatility process, make it very hard to compare the empirical results.

From the time series properties of the fundamental we can price an option only if we are willing to assume that the volatility risk is idiosyncratic as in Hull and White (1987). However, multiple studies find evidence of a nonzero volatility risk premium (see Bates (1996b) for references). This implies in turn that one needs some extra input to make the transition from the physical to the risk neutral measure. Observing only the underlying fundamental and estimating SV models with this information will not deliver derivative security pricing. One solution is to use the options data instead of the underlying fundamental for SV model estimation. It yields directly an estimate of the required risk-neutral parameters. Options data are essentially panel data, i.e. frequent (e.g. daily) observations through time of quotes or transaction prices are recorded. The moneyness and time to maturity of contracts create the cross-sectional heterogeneity. The use of option cross-sections yields estimates of the parameters, but it does not accommodate well the temporal dynamics of the model. Relying exclusively on the cross-section essentially makes the SV model a sophisticated version of the ad-hoc implementation of the Black-Scholes model. This paper focuses on the recent attempts to reconcile these shortcomings by incorporating the time series of the options data into the estimation strategy. Ultimately the goal is to exploit optimally the information in the panel data of options and the information in the underlying fundamental, realizing that some of the information is redundant. The fulfillment of this goal yields risk neutral parameters as well as the temporal dynamics under both measures.

Most of the recent attempts to address these issues use the variants of the Heston (1993) SV model which yields analytical option pricing formula. His approach was extended by many authors and generalized by Bakshi and Madan (1998) and Duffie *et al.* (1998). It applies to jump-diffusions of the affine class.[1] This imposes a certain uniformity on the empirical studies in option pricing because the models are easier to compare as they are nested in the general model specification described in Duffie *et al.* (1998). The parallel development of estimation techniques, including the Simulated Method of Moments (SMM) of Duffie and Singleton (1993) and its recent extension in Gallant and Tauchen (1996) known as the Efficient Method of Moments (EMM), allows one to deal with the issue of filtering the non-observable spot volatility and to achieve the efficiency of maximum likelihood without the knowledge of the likelihood function. Finally, typical econometric diagnostics are not enough. The size of option pricing errors is the ultimate test and object of interest. The out-of-sample performance of any model therefore becomes very important.

Table 1.1 provides a summary of the papers on the subject of estimating SV models for the purpose of option pricing. In the Table we provide information about the estimation method, the type of data used, the volatility filter and the type of

1. The Bakshi and Madan (1998) approach extends beyond the affine class. They, however, do not specify the mapping between the objective and risk-neutral measures.

model estimated. The content of the Table will be a guidance for the discussion through the remainder of the paper. Section 1.1 provides a summary review of SV models which have been considered. The next section discusses estimation techniques. Section 1.3 surveys approaches to volatility filtering and section 1.4 reviews the empirical findings. The last section concludes.

1.1 Characterizations of the Risk Premium

The continuous time affine class of SV models is the most popular because it features analytical solutions for pricing option contracts. Therefore, the majority of empirical studies adopt a specification which is affine.[2] This motivates us to provide some details about this class of models. Comparisons with non-affine models will be addressed in the next section when we discuss particular implementations.

We adopt a simplified version of the Duffie *et al.* (1998) description of an affine jump-diffusion. We assume that the logarithm of an asset price is the first factor in the following system:

$$dX_t = \mu(X_t, t)dt + \sigma(X_t, t)dW_t + dZ_t \tag{1.1}$$

All the models we consider here can be described by a three-dimensional vector X_t with up to four independent Brownian innovation shocks W_t. In particular,

$$\mu(x,t) = \begin{pmatrix} \mu_0 \\ \theta_2 \\ \theta_3 \end{pmatrix} + \begin{pmatrix} 0 & \mu_1 - \frac{1}{2} & c\left(\mu_2 - \frac{1}{2}\right) \\ 0 & -\kappa_2 & 0 \\ 0 & 0 & -\kappa_3 \end{pmatrix} x \tag{1.2}$$

$$\sigma(x,t)\sigma(x,t)^\top = \underbrace{\begin{pmatrix} 1 & \rho_2\sigma_2 & 0 \\ \rho_2\sigma_2 & \sigma_2^2 & 0 \\ 0 & 0 & 0 \end{pmatrix}}_{\Sigma_2} x_2 + \underbrace{\begin{pmatrix} c & 0 & \rho_3\sigma_3 \\ 0 & 0 & 0 \\ \rho_3\sigma_3 & 0 & d\sigma_3^2 \end{pmatrix}}_{\Sigma_3} x_3 \tag{1.3}$$

where x_i is the i^{th} component of the vector x. The univariate jump part is described in the following way:

$$dZ_t = -\lambda(X_t, t)\mu_J dt + \log(1 + J_t)dq_t \tag{1.4}$$

2. It is worth noting that there are compelling empirical reasons to focus on the affine class of models. For instance, Benzoni (1998) finds only marginal differences between the performance of the Heston (1993) model, which is in this class, and the Scott (1987) model, which is not affine.

Table 1.1
Estimation Methodology. The recent empirical work on the SV model estimation for the purpose of option pricing is summarized. The class of an adopted model, type of data used for estimation, estimation method, spot volatility filter and an approach to the estimated volatility dynamics assessment are reported for each paper. Notations: *iv* - BS implied volatility, *giv* - volatility implied from an SV model, *sv* - spot volatility, *o* - objective measure parameters, *n* - risk-neutral measure parameters.

Paper	Model	Type of data	Estimation	*SV* filtering	*SV* dynamics
Pastorello *et al* (1994)	Hull-White	daily *iv* time series	indirect inference	—	—
Bates (1996, 1998)	affine	weekly options panel (transactions)	NL-GLS, options group specific + individual errors	weekly *giv* together with *n*	estimation constrained by the *sv* pdf, but applied to the *giv*
Nandi (1996)	affine	intradaily options panel	NL-GLS, AR(1) errors	daily *giv* together with *n*	—
Bakshi *et al* (1997)	affine	intradaily calls cross-section	NL-OLS	daily *giv* together with *n*	—
Bakshi *et al* (1998)	affine	daily puts panel	SMM	daily *giv* given *n*	—
Jiang and van der Sluis (1998)	discrete time	daily returns + daily calls cross-section	EMM for *o*, NL-OLS for *n* given *o*	daily *sv* reprojected from returns	fully model consistent
Chernov and Ghysels (1998)	affine	daily returns jointly with *iv* time series	EMM	daily *sv* reprojected from returns and *iv*	fully model consistent
Benzoni (1998)	affine and Scott	daily returns + daily calls panel	EMM for *o*, NL-OLS/ SMM for *n* given *o*	daily *giv* with *n* given *o*/ Kalman filter	—/ linearized model consistent
Pan (1998)	affine	daily returns jointly with calls time series	GMM with *giv* used for *sv* (L-GMM)	daily *giv* together with *o*+*n*	*giv* vs. the realized volatility
Poteshman (1998)	non-para-metric	daily options panel+returns	EM-type algorithm+ kernel reg.	daily *giv* together with *o*+*n*	—

$$\log(1+J_t) \sim N(\log(1+\mu_J) - \frac{1}{2}\sigma_J^2, \sigma_J^2) \quad (1.5)$$
$$dq_t \sim Poi(\lambda(X_t, t)dt) \quad (1.6)$$
$$\lambda(x,t) = \lambda_0 + \lambda_2 x_2 + \lambda_3 x_3 \quad (1.7)$$

Finally, the risk-free interest rate is determined by:

$$r(x,t) = r_0 + r_2 x_2 + r_3 x_3 \quad (1.8)$$

The most general constant interest rate model is the one with $r_2 = r_3 = 0$, $c = d = 1$ and four innovation shocks was considered by Bates (1998). The most frequently used models are nested within the above model specification and can be obtained by setting some of the parameters equal to zero. For example, $c = \mu_1 = \theta_3 = \kappa_3 = \sigma_3 = \rho_3 = \lambda_3 = \lambda_2 = 0$ yields the model introduced in Bates (1996a). In addition, setting $\lambda_0 = 0$ yields the original Heston (1993) model. Scott (1997) considers the stochastic interest rate model with $c = \rho_3 = r_0 = \lambda_2 = \lambda_3 = 0$ and $d = r_2 = r_3 = 1$. Finally, Bakshi *et al.* (1997, 1998) investigate a slightly less general model with $r_2 = 0$.[3] All of these models have as many innovation shocks as factors.

In order to price options we need to characterize the transformation from the objective measure P to the risk-neutral measure P^*. If $\exp(X_{1t})$ is the underlying asset price, then we know from financial theory that the drift of the asset price process under P^* should be equal to $r(X_t, t)$. One can not identify restrictions on the drifts of other factors because they are not traded. Consequently, one must rely on some (parametric) specification of the transformation between P and P^*. One commonly used mapping is:

$$\mu^*(x,t) = \mu(x,t) - \Sigma\eta x \quad (1.9)$$
$$\theta(\psi) = (1+\mu_J)\exp\left\{\frac{\sigma_J^2}{2}\psi(\psi-1)\right\} \quad (1.10)$$
$$\lambda^*(x,t) = \lambda(x,t)\theta(\zeta) \quad (1.11)$$
$$\theta^*(\psi) = \theta(\psi+\zeta)/\theta(\zeta) \quad (1.12)$$

where Σ is a tensor comprised of Σ_i's, the vector η represents the assumed market prices of the factors risks, the scalar ζ represents the market price of the jump risk and $\theta(\psi)$ is the moment generating function of the jump size, J_t, distribution (see for instance Duffie *et al.* (1998), Ho *et al.* (1996) for details). The models

3. Note that most of the three factor models have the stochastic interest rate as a third factor. The exception is Bates (1998) who considers another latent process as a third factor.

specified in (1.1) through (1.12) are all special cases of the general option pricing theory developed in Bakshi and Madan (1998) and Duffie *et al.* (1998). They provide pricing formula for European-type as well as several other types of option contracts.

The above theoretical setup defines the task an econometrician faces. In particular, one has to estimate the objective measure parameters of the diffusion (1.1) and the market prices of risk $(\eta^\top, \zeta)^\top$. To execute this task one faces two important problems. The presence of unobservable (latent) variables and absence of the closed-form solution of the system of stochastic differential equations in (1.1) do not allow one to use the maximum likelihood estimation. This problem is well known and discussed in detail in the econometrics literature on the estimation of SV models (see Ghysels *et al.* (1996) for further details). Since Maximum Likelihood is not available, parameter estimates will not be most efficient. Hence, we need to find an estimation procedure with minimal efficiency losses. The second problem one faces is that in order to estimate the prices of factor risk one has to use the options data to estimate these risks.[4] We noted already that it is unclear at this point how one can utilize the available data optimally. For instance, the options cross-section will provide an excellent fit to the current date, but does not contain any information about the dynamics of the system. For instance, Bakshi *et al.* (1997) specify their model under P^* right away and re-fit it every day, foregoing any possibility to predict more than one step ahead. We know that the dynamics can be uncovered from the underlying returns time series. It is not obvious nor straightforward, however, how to combine options and returns data and which options data from the cross-section to select for this combination. Moreover, the data selection is also inherently related to the estimation methodology which is used.

1.2 Estimation Methodology

The first factor of the X_t vector process appearing in (1.1) is the logarithm of the stock price process. As is the case for any stochastic volatility process, there is no analytical expression for the likelihood function for X_{1t}. Hence, estimation has to proceed along different principles, such as method of moments or simulated likelihood methods. The Efficient Method of Moments of Gallant and Tauchen

4. The alternative approach is to determine the measure transformation via the general equilibrium consumption-based model argument as for instance in Bates(1996a, 1998). However, unless one is willing to estimate the relevant parameters from, say monthly consumption data, one can estimate the risk-neutralized parameters only from the options data. See also Scott (1997) for the case when the S&P 500 index is the fundamental process.

(1996) is at the intersection of full-blown maximum likelihood and moment-based estimation. More specifically, the EMM estimation procedure relies on moments, which are selected from a score of an auxiliary model which converges (in terms of the Sobolev norm) asymptotically to the true probability density of the data. Benzoni (1998) and Jiang and van der Sluis (1998) apply EMM using returns data. Chernov and Ghysels (1998) estimate the Heston (1993) model using the joint process of daily returns and at-the-money options implied volatility time series data. The estimation strategy proposed by Chernov and Ghysels (1998), in principle, applies to more complex models as well, involving panels of options data featuring different moneyness and maturity.[5] In such a setting one combines the information in the options cross-section with that in time series dynamics of options and the underlying fundamental (see Chernov *et al.* (1999) for further details). The major drawback of the EMM estimation procedure is that its implementation is computationally demanding and is currently infeasible for auxiliary densities of more than four series. Hence, for large systems (panels) one has to rely on less efficient moment-based procedures.

Several GMM (Hansen (1982)) strategies involving multiple assets were recently proposed and applied in the context of affine diffusion models. Liu (1998) estimates a stochastic interest rate extension of the Heston (1993) model using monthly yields and returns data. The GMM procedure requires analytical expressions for the unconditional moments of the model, which he derives under the assumption that the process is stationary. The stationarity assumption is warranted provided the initial value is drawn from a stationary marginal distribution. It is not feasible, however, to derive this distribution analytically.[6] Therefore, Liu assumes that the starting value is a constant. However, this may lead to the non-stationarity of the stock price process distribution, which in turn violates the regularity conditions required for implementing the GMM estimator and guarantee its standard root-T asymptotics. In one particular case one can derive the marginal distribution, though it involves the unappealing assumption that there is no leverage effect. Ho *et al.* (1996) derive analytical unconditional moments using this restriction. Another application of GMM involving the Heston model is discussed in Pan (1998). Her approach involves daily returns and options data, as in Chernov and Ghysels (1998). Pan derives analytical conditional moments for the model which depend on the

5. Pastorello *et al.* (1994) is an early paper which used exclusively options data and the indirect inference method of Gouriéroux *et al.* (1993), an estimation method very similar to EMM.

6. See the discussion in Hansen and Scheinkman (1995) along the same line with regard to the stationarity of multivariate diffusions.

spot latent volatility process. The presence of unobservable variables precludes the direct implementation of GMM (both in discrete time or in continuous time as suggested by Hansen and Scheinkman (1995)). Pan proposes an elegant modified GMM procedure, coined as L-GMM, which uses the Heston option pricing formula to imply the spot volatility from observed option prices. She invokes the implicit function theorem to obtain the asymptotic covariance estimator of the moment conditions which involve implied volatilities. The fact that latent spot volatilities are replaced by implied ones, still introduces a bias however, as discussed recently by Ledoit and Santa-Clara (1998). Earlier papers have often relied on GMM-type procedures, such as nonlinear ordinary and generalized least squares. Examples include Bates (1996a, 1998), Nandi (1996) and Bakshi *et al.* (1997). In general the estimation procedures in these papers are consistent but do not fully exploit the optimal implementation of GMM.

The aforementioned difficulties encountered with the applications of GMM can be resolved with the Simulated Method of Moments estimation procedure. Duffie and Singleton (1993) derive the asymptotic properties of SMM estimators in the context of time-homogeneous Markov processes. While there are many similarities with GMM, the fact that simulations are used to compute moment conditions has several advantages. In addition to resolving the problems of starting values and the presence of latent processes it also allows for moment conditions which are analytically intractable.[7] This last feature is particular relevant here, since moment conditions based on options prices cannot be computed analytically. While there are clear advantages to the use of SMM, there is a cost associated with the simulation uncertainty which entails a loss of efficiency. One can easily control such losses, however, by augmenting the number of simulations (see e.g. Duffie and Singleton (1993) for further details). For example, even a small number of simulated paths, say ten, increases confidence intervals by no more than 5% relative to GMM with the same moment conditions. The recent applications of SMM to options valuation include Bakshi *et al.* (1998), who use options panel data to estimate a Heston-type model with stochastic interest rates and jumps, and Benzoni (1998), who uses the SMM procedure as a second stage estimator of the market price of volatility risk in the context of the Heston (1993) and Scott (1987) models.

An interesting deviation from the Method of Moments approach is suggested by Poteshman (1998). His approach does not involve a parametric specification for μ and σ in a bivariate version of (1.1). The procedure involves iterations between

7. Please note that EMM and SMM are closely related procedures. The EMM procedure is a particular SMM estimator where the choice of moment conditions is guided by an auxiliary model.

implying volatilities from the options cross-section and non-parametric estimation of the model based on the time-series of these volatilities via an EM-type algorithm. Poteshman (1998) assumes constant risk-free interest rate, to avoid non-parametric estimation of the interest rate process, and imposes a constant correlation coefficient as well. His method has similarities with Pastorello *et al.* (1994) who also iterate between simulated instantaneous volatilities and the ML estimation of the model parameters.

1.3 Volatility Filtering

The distinguishing feature of SV models is the latent volatility process. While it is possible to estimate SV models, using for instance the underlying fundamental, without filtering the volatility process, it is impossible to price options based only on the estimated parameters. Therefore, filtering volatility is intimately related to the pricing of options. Early studies of SV models considered filters based on returns data (see Ghysels *et al.* (1996) for a discussion). Several filters involving exclusively the return process have been proposed in the literature. Harvey *et al.* (1994) suggested to make use of the Kalman filter based on a discrete time SV model. Nelson and Foster (1994) showed how diffusion limit arguments applied to the class of EGARCH models provided a justification of EGARCH models as filters of the instantaneous volatility. French *et al.* (1987) suggest volatility filters based on squared returns. Along the same lines Nelson and Foster (1996) provide the theoretical foundations for using rolling regressions. Some attempts were made to extend these filters to a multivariate context, see in particular Harvey *et al.* (1994), Jacquier *et al.* (1995) and Nelson (1996). These multivariate extensions all involve exclusively return series and cannot accommodate derivative security market information.

The Kalman filter is known to be suboptimal because the volatility filtering problem is inherently non-Gaussian and nonlinear. The exact filter was derived by Jacquier *et al.* (1994) using Bayesian Markov Chain Monte Carlo methods. One advantage of the Bayesian MCMC methods is that filtering spot volatility is a natural by-product of estimation. The EMM estimation procedure has similar features. Gallant and Tauchen (1998) proposed an extension of EMM which yields asymptotically unbiased filters as well, based on a reprojection procedure (which will be described shortly). Bayesian MCMC methods are not easy to extend to options data, either involving simultaneously returns data and options or a panel of options. The EMM procedure in its generic form can be multivariate, (i) only involving a vector of returns, (ii) only involving a vector of options, and (iii) a mixture of the previous two.

Most studies involving options data proposed filters, which are generalized implied volatilities based on inverting the SV option pricing formula (see Table 1.1). There are at least two disadvantages to the use of implied volatilities as filters. First, implied volatilities obtained from European options are biased estimates of spot volatility as they are related to the expected volatility over the remaining life of the option contract. Second, the time series dynamics of estimated volatility are muted since implied volatilities are extracted from options cross-sections. Bates (1996a, 1998) tries to remedy this by augmenting the NL-GLS estimation procedure with restrictions pertaining to the dynamics of the spot volatility process. Since his method still involves implied volatilities, however, it does not resolve the bias due to replacing spot volatility by implieds. For the same reasons, this bias also affects the methods proposed by Pan (1998) and Poteshman (1998).

To proceed further we need to briefly describe the reprojection method of Gallant and Tauchen (1998). Denote the vector of contemporaneous and lagged observed variables by x_t and the vector of contemporaneous unobserved variables by y_t, Θ is the parameters vector. The filtering problem is equivalent to computing the following conditional expectation:

$$\tilde{y}_t = E(y_t|x_t) = \int y_t p(y_t|x_t, \Theta) dy_t \tag{1.13}$$

This expectation involves the conditional probability density of y_t. If we knew the density implied by the system dynamics, we could estimate it by $\hat{p}(y_t|x_t) = p(y_t|x_t, \hat{\Theta})$. Unfortunately, for SV models there is no analytical expression for the conditional density available. Therefore, we need to estimate this density as $\hat{p}(y_t|x_t) = f_K(\hat{y}_t|\hat{x}_t)$, where $\hat{y}_t, \hat{x}_t$ are simulated from the SV model with parameters set equal to $\hat{\Theta}$ and where f_K is the SNP density of Gallant and Tauchen (1989). Gallant and Long (1997) show that:

$$\lim_{K\to\infty} f_K(\hat{y}_t|\hat{x}_t) = p(y_t|x_t, \hat{\Theta}) \tag{1.14}$$

where the index K is related to the dimension of the SNP density parameter vector and grows at the appropriate rate with the sample size. Hence, as noted earlier, reprojection provides an unbiased estimate of the spot volatility.

The reprojection method for the purpose of option pricing is applied by Chernov and Ghysels (1998) and Jiang and van der Sluis (1998). The latter consider univariate filters involving returns data. Chernov and Ghysels (1998) use the reprojection method in a univariate context, involving at-the-money options, and a bivariate setup involving returns data in addition to options.[8] Their results show that the univariate approach only involving options by and large dominates. A by-product of this finding is that they uncover a remarkably simple volatility extraction filter based on a polynomial lag structure of Black-Scholes implied volatilities. The simplicity of this scheme is rather surprising if one thinks of the complexity of the task. Indeed, alternative filters based on returns involve highly nonlinear functions of returns. Hence, the virtue of using volatility data to predict future spot volatility is that one can limit the filter to a linear structure. It should be noted, however, that the construction of the filter is unfortunately not as simple as running a linear regression model. Since spot volatility is a latent process one can only recover the filter weights via the reprojection procedure.

1.4 Evaluation and Empirical Findings

It is standard econometric practice to evaluate models on the basis of the in-sample fit. Diagnostics like overidentifying restrictions tests are among the most commonly used. Obviously, one also looks at the parameter estimates and their economic interpretation. The empirical studies reported in Table 1.1 use quite a variety of model specifications and data. The results are therefore difficult to compare directly. However, there are some findings which are robust. For instance, there is no doubt that volatility risk is priced and that volatility is negatively correlated with the asset returns (leverage effect). In this regard there are some unsettled issues. For instance, using Heston's model Chernov and Ghysels (1998) find only a small, albeit statistically significant, leverage effect regardless whether it is estimated with S&P 500 returns data or SPX options or both together. In contrast, Bakshi *et al* (1998), Benzoni (1998), Nandi (1996) and Pan (1998), using different sample periods, find much larger (though with wide range) point estimates of the leverage effect parameter for the same model. On the other hand, the parameter estimate of the price of volatility risk in Benzoni (1998) and Pan (1998) implies that the volatility process is non-stationary under risk neutral distribution while it is stationary under the objective measure. Such a discrepancy is not found in Bakshi *et al* (1998), Chernov and Ghysels (1998) and Nandi (1996). At this point, there is no clear

8. Gallant *et al.* (1998) adopt a strategy similar to the bivariate approach in Chernov and Ghysels (1998) though not involving options. Namely, they consider the bivariate process of the fundamental and the daily high/low spread, which provides extra information about the course of volatility.

understanding yet of these widely different results. Both theoretical and empirical developments are necessary here.

At the theoretical level, the appropriate affine model specification analysis in Dai *et al* (1998) becomes important. In particular, they show that the models considered in empirical studies impose unnecessary zero restrictions on the parameters. Removing such restrictions will help uncovering the dynamics present in the data. On the empirical side, the next logical step is to use a genuine panel of options and the fundamental returns jointly to estimate a sufficiently rich model. This line is pursued in Chernov *et al.* (1999).

Statistical criteria for model selection are one thing, financial criteria such as pricing and hedging performance out-of-sample are obviously the ultimate scope of model selection. A model rejected by the data may provide more accurate option pricing compared to one with a good in-sample fit. Hence estimated models should also be evaluated out-of-sample. Typically one uses the relative or absolute mean-squared pricing error, which corresponds to the quadratic utility-based loss function.[9] An alternative way to compare models is to appraise the performance of their volatility filters. Chernov and Ghysels (1998) suggest to substitute different volatility filters into the Black-Scholes formula. This comparison allows one to separate the role of volatility filtering from that of the pricing kernel in the valuation of derivative contracts. This comparison also allows one to see whether pricing formula more advanced than BS add any improvement in pricing performance.

There is a vast literature which studies whether the BS implied volatility is a good forecast of the realized volatility (see Bates (1996b) for a review). The premise of these studies is based on the insight of the Hull-White model that implied volatility is the expected value of the average integrated spot volatility. Most of these studies find that implied volatility is biased upwards as compared to the realized volatility. We can explain these findings by the fact that Hull and White (1987) do not model the leverage effect and volatility risk premium, which are found to be present in the data.

Overall, we still find that standard SV models are rejected, on the basis of their statistical fit, by the data whatever the source of observations is, i.e. returns or options. There is improvement in the out-of-sample option pricing performance according to some criteria, like volatility filtering and compared to naive applica-

9. It would be interesting to use the representative investor utility as a loss-function. For instance, the constant interest rate affine models can be supported by the economy with logarithmic preferences (see Bates(1996a)). Brandt (1998) uses the exponential utility to evaluate hedging errors.

tions of the Black-Scholes model. These improvements are not always spectacular and may not justify the additional complications that SV models bring along. Moreover, all studies involving a bivariate model find that one latent factor is not enough to explain the variability in the fundamental process and options prices. A second latent variable or a jump component are necessary to improve pricing performance. Bates (1998) made a first contribution in this direction by considering a three-factor model with jumps. Chernov *et al* (1999) and Duffie *et al* (1998) pursue the extensions of this line of research.

Conclusion

We reviewed the recent advances on the estimation of multi-factor latent volatility models for the purpose of option pricing. Important theoretical developments and improvements in econometric methods were the main reason of the recent surge of research interest in this area.

The multitude of theoretical models, of data sets used and of performance evaluation still does not allow us to make general conclusions about model specification and fitting stylized facts like the volatility smile. We tried to assess where progress was made in recent years. While advances were made in the formulation and estimation of affine SV models we still need to go a long way on their practical implementation. We expect that much of the literature will ultimately converge to a consensus SV model which passes the rigorous tests of statistical fit and outperforms conventional ad hoc pricing schemes.

References

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

Bakshi, G., C. Cao, and Z. Chen. 1998. Pricing and hedging long-term options. *Journal of Econometrics*. Forthcoming.

Bakshi, G. and D. Madan. 1998. Spanning and derivative-security valuation. *Journal of Financial Economics*. Forthcoming.

Bates, D. 1996a. Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies*, 9:69–107.

Bates, D. 1996b. Testing option pricing models. In Maddala, G. and C. Rao, editors, *Handbook of Statistics, Vol. 14*. Elsevier Science B.V.

Bates, D. 1998. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*. Forthcoming.

Benzoni, L. 1998. Pricing options under stochastic volatility: An econometric analysis. Working paper, Northwestern University.

Black, F. and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659.

Brandt, M. 1998. Multiperiod hedging of contingent claims. Working paper, University of Pennsylvania.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen. 1999. Estimating Jump Diffusions with a Panel of Options and Returns. Working Paper, Pennsylvania State University.

Chernov, M. and E. Ghysels. 1998. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. Working paper, Pennsylvania State University.

Dai, Q., J. Liu, and K. Singleton. 1998. Admissibility and identification problems of affine asset pricing models. Working paper, Stanford University.

Duffie, D., J. Pan, and K. Singleton. 1998. Transform analysis and option pricing for affine jump-diffusions. Working Paper, Stanford University.

Duffie, D. and K. Singleton. 1993. Simulated moments estimation of Markov models of asset prices. *Econometrica*, 61:929–952.

French, K., G. W. Schwert, and R. Stambaugh. 1987. Expected stock returns and volatility. *Journal of Financial Economics*, 19:3–29.

Gallant, A. R., C.-T. Hsu, and G. Tauchen. 1998. Calibrating volatility diffusions and extracting integrated volatility. Working Paper, Duke University.

Gallant, A. R. and J. Long. 1997. Estimating stochastic differential equations efficiently by minimum chi-square. *Biometrika*, 84:125–141.

Gallant, A. R. and G. Tauchen. 1989. Seminonparametric estimation of conditionally constrained heterogeneous processes: Asset pricing applications. *Econometrica*, 57:1091–1120.

Gallant, A. R. and G. Tauchen. 1996. Which moments to match? *Econometric Theory*, 12:657–681.

Gallant, A. R. and G. Tauchen. 1998. Reprojecting partially observed systems with application to interest rate diffusions. *Journal of American Statistical Association*, 93:10–24.

Ghysels, E., A. Harvey, and E. Renault. 1996. Stochastic volatility. In Maddala, G. and C. Rao, editors, *Handbook of Statistics, Vol. 14*. Elsevier Science B.V.

Gouriéroux, C., A. Monfort, and E. Renault. 1993. Indirect inference. *Journal of Applied Econometrics*, 85:85–118.

Hansen, L. P. 1982. Large sample properties of generalized method of moments estimators. *Econometrica*, 50:1029–1054.

Hansen, L. P. and J. Scheinkman. 1995. Back to the future: Generating moment implications for continuous-time Markov processes. *Econometrica*, 63:767–804.

Harvey, A., E. Ruiz, and N. Shephard. 1994. Multivariate stochastic variance models. *Review of Economic Studies*, 61:247–264.

Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–343.

Ho, M., W. Perraudin, and B. Sørensen. 1996. A continuous-time arbitrage-pricing model with stochastic volatility and jumps. *Journal of Business and Economic Statistics*, 14:31–43.

Hull, J. and A. White. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42:281–300.

Jacquier, E., N. Polson, and P. Rossi. 1994. Bayesian analysis of stochastic volatility models (with discussion). *Journal of Business and Economic Statistics*, 12:371–417.

Jacquier, E., N. Polson, and P. Rossi. 1995. Stochastic volatility: Univariate and multivariate extensions. Working Paper, Boston College.

Jiang, G. and P. van der Sluis. 1998. Pricing stock options under stochastic volatility and interest rates with Efficient Method of Moments estimation. Working Paper, University of Groningen.

Ledoit, O. and P. Santa-Clara. 1998. Relative pricing of options with stochastic volatility. Working Paper, UCLA.

Liu, J. 1998. Portfolio selection in stochastic environments. Working Paper, Stanford University.

Nandi, S. 1996. Pricing and hedging index options under stochastic volatility. Working Paper, Federal Reserve Bank of Atlanta.

Nelson, D. 1996. Asymptotic filtering theory for multivariate arch models. *Journal of Econometrics*, 71:1–47.

Nelson, D. and D. Foster. 1994. Asymptotic filtering theory for univariate arch models. *Econometrica*, 62:1–41.

Nelson, D. and D. Foster. 1996. Continuous record asymptotics for rolling sample variance estimators. *Econometrica*, 64:139–174.

Pan, J. 1998. "Integrated" time-series analysis of spot and option prices. Working Paper, Stanford University.

Pastorello, S., E. Renault, and N. Touzi. 1994. Statistical inference for random variance option pricing. *Journal of Business and Economic Statistics*. Forthcoming.

Poteshman, A. 1998. Estimating a general stochastic variance model from option prices. Working Paper, University of Chicago.

Scott, L. 1987. Option pricing when the variance changes randomly: Theory, estimation and an application. *Journal of Financial and Quantitative Analysis*, 22:419–438.

Scott, L. 1997. Pricing stock options in a jump-diffusion model with stochastic volatility and interets rates: Applications of Fourier inversion methods. *Mathematical Finance*, 7:413–426.