Journal of Econometrics 160 (2011) 22–32

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Estimation of objective and risk-neutral distributions based on moments of integrated volatility☆

René Garcia[a,*], Marc-André Lewis[b], Sergio Pastorello[c], Éric Renault[d]

[a] Edhec Business School, France
[b] Banque Nationale du Canada, Canada
[c] Dipartimento di Scienze Economiche, Università di Bologna, Italy
[d] University of North Carolina at Chapel Hill, CIRANO and CIREQ, United States

ARTICLE INFO

Article history:
Available online 6 March 2010

JEL classification:
C1
C5
G1

Keywords:
Realized volatility
Implied volatility
Volatility risk premium
Moments of integrated volatility
Objective distribution
Risk-neutral distribution

ABSTRACT

In this paper, we present an estimation procedure which uses both option prices and high-frequency spot price feeds to estimate jointly the objective and risk-neutral parameters of stochastic volatility models. The procedure is based on a method of moments that uses analytical expressions for the moments of the integrated volatility and series expansions of option prices and implied volatilities. This results in an easily implementable and rapid estimation technique. An extensive Monte Carlo study compares various procedures and shows the efficiency of our approach. Empirical applications to the Deutsche mark–US dollar exchange rate futures and the S&P 500 index provide evidence that the method delivers results that are in line with the ones obtained in previous studies where much more involved estimation procedures were used.

1. Introduction

In continuous-time modeling in finance, stochastic processes for asset prices are combined with an absence of arbitrage argument to obtain the prices of derivative assets. Therefore, statistical inference on continuous-time models of asset prices can and should combine two sources of information, namely the price history of the underlying assets on which derivative contracts are written and the price history of the derivative securities themselves. However, statistical modeling poses a challenge. A joint model needs to be specified, not only for the objective probability distribution which governs the random shocks observed in the economy, but also for the risk-neutral probability distribution, which allows us to compute derivative asset prices as expectations of discounted payoffs. Since the two distributions have to be equivalent, there exists a link between the two through an integral martingale representation which includes the innovations associated with the primitive asset price processes and the risk premia associated with these sources of uncertainty. Moreover, state variables, observable or latent, may affect the drift and diffusion coefficients of the primitive assets and the corresponding risk premia.

The main contribution of this paper is to propose a new methodology for an integrated analysis of spot and option prices. It is based on simple generalized method-of-moment (GMM) estimators of both the parameters of the asset price and state variable processes and the corresponding risk premia. To focus on the issue of the joint specification of an objective probability distribution and a risk-neutral one, we will restrict ourselves to the case of one state variable which will capture the stochastic feature of the volatility process of the underlying asset. We will adopt a popular affine diffusion model where volatility is parameterized as follows:

$$dV_t = k(\theta - V_t)dt + \gamma\sqrt{V_t}dW_t^{\sigma},$$

where V_t is a latent state variable with an innovation governed by a Brownian motion W_t^{σ}. This innovation can be correlated (with a

☆ This paper was started while the second author was a post-doctoral fellow at CIRANO. We would like to thank two anonymous referees and the editors for their very helpful comments and suggestions. We also benefited from comments by Nour Meddahi and conference participants at the Canadian Econometric Study Group (CESG) in Waterloo (2001), the Econometric Society Winter Meeting in Atlanta (2002), and the International Statistical Symposium and Bernoulli Society EAPR Conference in Taipei (2002). The first author gratefully acknowledges financial support from the Fonds Québécois de la recherche sur la société et la culture (FQRSC), the Social Sciences and Humanities Research Council of Canada (SSHRC), the Network of Centres of Excellence MITACS, Hydro-Québec and the Bank of Canada. The third author gratefully acknowledges financial support from the MIUR Prin 2005 program.

* Corresponding address: Département de Sciences Économiques, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, Québec, H3C 3J7, Canada.
E-mail address: rene.garcia@umontreal.ca (R. Garcia).

0304-4076/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2010.03.011

coefficient ρ) with the innovation of the primitive asset price process governed by W_t^S:

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t^S.$$

In a seminal paper, Hull and White (1987) have shown that, in the particular case where $\rho = 0$, the arbitrage-free option price is a conditional expectation of the Black and Scholes (1973) (BS) price, where the constant volatility parameter σ^2 is replaced by the so-called mean integrated volatility: $\frac{1}{T-t}\mathcal{V}_{t,T} = \frac{1}{T-t}\int_t^T \sigma_s^2 ds$ and where the conditional expectation is computed with respect to the risk-neutral conditional probability distribution of $\mathcal{V}_{t,T}$ given σ_t. Heston (1993) has extended the analytical treatment of this option pricing formula to the case where ρ is different from zero, allowing for leverage effects and the presence of risk premia. However, with or without correlation, the option pricing formula involves the computation of a conditional expectation of a highly nonlinear integral function of the volatility process.

To simplify this computation, we propose to use an expansion of the option pricing formula in the neighborhood of $\gamma = 0$, as in Lewis (2000), which corresponds to the Black–Scholes deterministic volatility case. The coefficients of this expansion are well-defined functions of the conditional moments of the joint distribution of the underlying asset returns and integrated volatilities, which we also derive analytically. These analytical expansions will allow us to compute very quickly implied volatilities which are functions of the parameters of the processes and of the risk premia. An integrated GMM approach using intraday returns for computing approximate integrated volatilities (see the pioneering papers of Andersen and Bollerslev, 1998; Barndorff-Nielsen and Shephard, 2001, 2004) and option prices for computing implied volatilities allow us to estimate jointly the parameters of the processes and the volatility risk premium λ.

The main attractive feature of our method is its simplicity once analytical expressions for the various conditional moments of interest are available. These expressions were derived by Lewis (2001) using a recursive method, and also by Bollerslev and Zhou (2002) for the first two moments. The great advantage of the affine diffusion model is precisely to allow an analytical treatment of the conditional moments of interest. Bollerslev and Zhou (2002) have developed such a GMM approach based on the first two moments of integrated volatility to estimate the objective parameters of stochastic volatility and jump-diffusion models. In our estimation, we add moment conditions based on the third moment of integrated volatility. We hope that using these moments will help to better identify certain coefficients, and in particular to effectively estimate the asymmetry coefficient ρ.[1]

Recently, Bollerslev et al. (2011) adopted a very similar approach to ours, but considered a so-called model-free approach to recover implied volatilities. There is clearly a trade-off between model-free and model-based approaches to recover implied volatilities. While a model-free approach is robust to misspecification, it requires theoretically continuous strikes for option prices or practically a very liquid market like the S & P 500 option market.[2] Model-based approaches like ours may be sensitive to misspecification but they require only a few option prices. It may be the only way to proceed for options on individual stocks or less liquid option markets in general. It should also be emphasized that model-free implied volatilities are used ultimately to estimate parameters of a volatility model and the corresponding risk premium. For efficiency reasons, it may make sense to use the less noisy model-based implied volatilities, given of course that the model is well-specified.

Only few studies have estimated jointly the risk-neutral and objective parameters, and the estimation methods used are generally much more involved. Pastorello et al. (2000) proposed an iterative estimation procedure that used option and returns information to provide an estimate of the objective parameters in the absence of risk premia. Poteshman (1998) extends their methodology to include correlation between returns and volatility, a non-zero price of volatility risk, and flexible nonparametric specifications for this price of risk as well as the drift and diffusion functions of the volatility process. Chernov and Ghysels (2000) use the Efficient Method of Moments (EMM), a procedure that estimates the parameters of the structural model through a seminonparametric auxiliary density. Pan (2002) uses the Fourier transform to derive a set of moment conditions pertaining to implied states and jointly estimate jump-diffusion models using option and spot prices.[3] Pastorello et al. (2003) propose a general methodology of iterative and recursive estimation in structural non-adaptive models which nests all the previous implied state approaches. Finally, Aït-Sahalia and Kimmel (2007) propose a maximum likelihood approach, using closed-from approximations to the likelihood function of the joint observations on the underlying asset and option prices.

Compared to all these methods, the main advantage of our method is its simplicity and computational efficiency. We show through an extensive Monte Carlo that the estimation procedure works well for both the no-leverage model ($\rho = 0$) and the leverage model ($\rho \neq 0$). Of course the selected moment conditions differ between the two models. Due to the presence of a correlation parameter in the leverage model, we include moment conditions involving the cross-product of returns with either integrated volatility or implied volatility. In the no-leverage case, the moment conditions are based only on moments of integrated volatility and of implied volatility.

Finally, we provide an empirical illustration of our method for each model. The no-leverage model is applied to the Deutsche mark–US dollar exchange rate futures market. We use 5-min returns on the exchange rate futures and daily option prices on the same futures to compute the moments and implement our methodology. For the leverage model, we use 5-min returns on the S & P 500 index and daily option prices on the same index. Results show the presence of a significant volatility risk premium.

The rest of the paper is organized as follows. In Section 2, we present the general methodology, and show how to construct two blocks of moment conditions for the estimation of the models, one based on the high-frequency return measures, another on the implied volatility obtained as power series in the volatility of volatility parameter γ. Section 3 describes the moment conditions for the first block of moment conditions, while Section 4 explains how to use option price expansions to define model-specific

[1] Meddahi (2002) also derives explicit formulas for conditional and unconditional moments of the continuous-time Eigenfunction Stochastic Volatility (ESV) models of Meddahi (2001), which include as special cases the log-normal, affine and GARCH diffusion models.

[2] For model-free approaches see in particular Britten-Jones and Neuberger (2000), Lynch and Panigirtzoglou (2003) and Jiang and Tian (2005). The latter study shows how to implement the model-free implied volatility using observed option prices. They characterize the truncation errors when a finite range of strike prices are available in practice. To calculate the model-free implied volatility, they use a curve-fitting method and extrapolation from endpoint implied volatilities.

[3] Duffie et al. (2000) have extended the moment computations to the case of affine jump-diffusion models (where jumps are captured by Poisson components), while Barndorff-Nielsen and Shephard (2001) have put forward the so-called Ornstein–Uhlenbeck-like processes with a general Levy innovation. The general statistical methodology that we develop in this paper could be extended to these more general settings if a specification is chosen for the risk premia of the various jump components.

implied volatilities, and how these implied volatilities can lead to the estimation of the parameters. Section 5 presents a Monte Carlo study for two stochastic volatility models, with and without leverage, for several sets of parameter values. In Section 6, we provide two empirical illustrations of the methodology. Appendices A and B contain the expressions for the moment of integrated volatility and for the cross-moments in the leverage model.

2. A general outline of the method

As stated in the introduction, two different but equivalent sets of bivariate stochastic processes are to be considered here. The objective process is taken to be the affine stochastic volatility process

$$d\begin{bmatrix} S_t \\ V_t \end{bmatrix} = \begin{bmatrix} \mu_t S_t \\ \kappa(\theta - V_t) \end{bmatrix} dt + \sqrt{V_t}\begin{bmatrix} S_t & 0 \\ \gamma\rho & \gamma\sqrt{1-\rho^2} \end{bmatrix}\begin{bmatrix} dW_t^1 \\ dW_t^2 \end{bmatrix}, \quad (1)$$

where S_t and V_t are the price and volatility processes. The affine model for the volatility appearing in the returns process was studied by Heston (1993), Duffie et al. (2000) and Meddahi and Renault (2004). The risk-neutral process is taken to be

$$d\begin{bmatrix} S_t \\ V_t \end{bmatrix} = \begin{bmatrix} r_t S_t \\ \kappa^*(\theta^* - V_t) \end{bmatrix} dt + \sqrt{V_t}\begin{bmatrix} S_t & 0 \\ \gamma\rho & \gamma\sqrt{1-\rho^2} \end{bmatrix}\begin{bmatrix} d\tilde{W}_t^1 \\ d\tilde{W}_t^2 \end{bmatrix}, \quad (2)$$

where, by virtue of Girsanov theorem, only the parameters κ and θ are modified in the passage from one measure to the other. We follow Heston (1993) and specify the risk premium structure as: $\kappa^* = \kappa - \lambda$; $\kappa^*\theta^* = \kappa\theta$, the volatility risk premium being parameterized by λ. For such models, the objective parameters to be estimated are[4]: $\beta = (\kappa, \theta, \gamma, \rho)$. In order to define the risk-neutral set of parameters $\beta^* = (\kappa^*, \theta^*, \gamma, \rho)$, one must have the additional parameter λ, since we shall assume the short rate r_t to be observed. We will denote by ψ the vector of parameters comprised of β and λ.

In the next sections, we will show that high-frequency measures of returns can be used to measure the integrated volatility $\mathcal{V}_{t,T}$. Lewis (2001) proposes a method to compute conditional moments of the integrated volatility in affine stochastic volatility models. Using these, it is possible to construct a set of moment conditions $f_1(\mathcal{V}_{t,T}, \beta)$, which is such that $\mathrm{E}\left[f_1(\mathcal{V}_{t,T}, \beta)\right] = 0$. We will denote by $\hat{\beta}$ the estimator based on the set of moment conditions f_1.

Moreover, it is possible to define model-specific implied volatilities $V_t^{imp}(\beta, \lambda, \{c_{obs}\})$, with $\{c_{obs}\}$ being the set of observed option prices. These implied volatilities, that are not to be confounded with Black–Scholes implied volatilities, are defined to be the point-in-time volatility which gives, for given values of the risk-neutral parameters β^*, the observed option price. We use these implied volatilities to construct a second set of moment conditions $f_2(V_t^{imp}(\beta, \lambda, \{c_{obs}\}), \beta)$, which depends in a very nonlinear way on the parameters β and λ. It is thus possible to construct a joint set of moment conditions

$$\mathrm{E}\begin{bmatrix} f_1(\mathcal{V}_{t,T}, \beta) \\ f_2(V_t^{imp}(\beta, \lambda, \{c_{obs}\}), \beta) \end{bmatrix} = 0, \quad (3)$$

which we use to estimate by GMM the objective parameters β and the risk premium λ. We will call $\hat{\psi}$ the estimator based on the joint set of moment conditions f_1 and f_2.

4 Since the drift term μ_t does not matter for option pricing purposes, we do not specify it explicitly. Moreover, the inference method we will use for the objective parameters is robust to its specification.

3. Estimating objective parameters from high-frequency returns

From the seminal works of Andersen and Bollerslev (1998) and Barndorff-Nielsen and Shephard (2001), we know that high-frequency intraday data on returns can be used to obtain indirect information on the otherwise unobservable volatility process. The logarithmic price of an asset is assumed to obey the stochastic differential equation

$$dp_t = \mu(p_t, V_t, t)dt + \sqrt{V_t}dW_t,$$

where V_t is the squared-volatility process (which could be stochastic, particularly of the affine type we discussed above) and W_t is a standard brownian motion. If the drift and diffusion coefficients are sufficiently regular to guarantee the existence of a unique strong solution to the SDE, then, by the theory of quadratic variation, we have

$$\mathrm{plim}_{N\to\infty}\sum_{i=1}^{N}\left[p_{t+\frac{i}{N}(T-t)} - p_{t+\frac{i-1}{N}(T-t)}\right]^2 \longrightarrow \int_t^T V_s ds \equiv \mathcal{V}_{t,T},$$

and $\mathcal{V}_{t,T}$ is referred to as the integrated volatility of the process V_t from time t to T. Andersen et al. (2001a,b, 2003) offer a characterization of the distributional features of daily realized returns volatilities constructed from high-frequency five-min returns for foreign exchange and individual stocks. The finiteness of the number of measures induces a systematic error in the integrated volatility measure, and, in fact, the quadratic variation estimator will be a biased estimator of the integrated volatility if the drift term is not zero, this bias falling as the number of measures increases.

Bollerslev and Zhou (2002) use such an aggregation of returns to obtain integrated volatility time series from which they estimate by GMM the parameters of Heston's (1993) stochastic volatility model. They base their estimation on a set of conditional moments of the integrated volatility, where they add to the basic conditional mean and second moment various lag-one and lag-one squared counterparts. In constructing estimates of the objective parameters of the stochastic volatility process, we follow their basic approach but introduce a new set of moment conditions involving higher moments of the integrated volatility, in particular its skewness. Lewis (2001) derives analytically all conditional moments of the integrated volatility for the class of affine stochastic volatility models (which includes the Heston (1993) and the Hull and White (1987) models).[5]

Some attention has to be devoted to information sets. Following the notation of Bollerslev and Zhou (2002), we shall define the filtration $\mathcal{F}_t = \sigma\{V_s, s \le t\}$, that is, the sigma algebra generated by the instantaneous volatility process. Our moment conditions for the integrated volatility are originally conditional on this filtration. Since only the integrated volatility is observable, we need to introduce the discrete filtration $\mathcal{G}_t = \sigma\left\{\mathcal{V}_{s-1,s}, s = 0, 1, 2, \ldots, t\right\}$, which is the sigma algebra of integrated volatilities. Integrated volatilities are not observable, but realized volatilities are. As Bollerslev and Zhou (2002), we ignore the discretization noise. Corradi and Distaso (2006) provide a theoretical foundation for the approach that ignores the noise in a double asymptotic setting. Evidently, the filtration $\mathcal{G}_t$ is nested in the finer $\mathcal{F}_t$. This enables one to rewrite moment conditions in terms of the coarser filtration using the law of iterated expectations: $\mathrm{E}\left[\mathrm{E}\left(\cdot|\mathcal{F}_t\right)|\mathcal{G}_t\right] = \mathrm{E}\left(\cdot|\mathcal{G}_t\right)$.

5 Duffie et al. (2000) provide analytical expressions for the instantaneous volatility process for such models. Bollerslev and Zhou (2002) derived analytical expressions for the mean and variance of the integrated volatility in Feller-type volatility models. To our knowledge, higher moments of the integrated volatility were not previously computed. Zhou (2003) characterized the Itô conditional moment generator for affine jump-diffusion models, and other nonlinear quadratic variance and semiparametric flexible jump models.

3.1. The no-leverage model

In the case where there is no correlation between returns and volatility ($\rho = 0$), we use the following set of moment conditions:

$$m_{1t}(\beta) = \left[\mathcal{V}^k_{t+1,t+2} - \mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right]\right], \quad k = 1, 2, 3. \tag{4}$$

The three conditional moment restrictions (4) can be expressed in terms of observed integrated volatilities because we have (see Appendix A) closed-form formulas for $\mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right]$ in terms of $\mathrm{E}\left[\mathcal{V}^k_{t,t+1}|\mathcal{G}_t\right]$, for $k = 1, 2, 3$. We use each of the resulting three orthogonality conditions with two instruments, a constant and $\mathcal{V}^k_{t-1,t}$, which results in six unconditional moment restrictions.[6]

$$f_{1t}(\beta) = \begin{bmatrix} \mathcal{V}^k_{t+1,t+2} - \mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right] \\ \left(\mathcal{V}^k_{t+1,t+2} - \mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right]\right)\mathcal{V}^k_{t-1,t} \end{bmatrix}, \quad k = 1, 2, 3. \tag{5}$$

3.2. The leverage model

In the leverage model, the correlation coefficient between the two Brownian motions, ρ, appears as an additional parameter. In principle, it could be identified using only the marginal moments of the integrated and the spot volatility (because the implied spot volatilities depend on both λ and ρ). In practice, these moment expressions are not able to accurately identify the parameter ρ. We add some cross-moments between the log returns[7] and the integrated volatility to the moment conditions in (4). The set of twelve moment conditions used in the estimation of the objective parameters is given below. To derive the closed-form expression of the cross-moments, we used the recurrence formula provided in Appendix A of Lewis (2001). These expressions are given in Appendix B.

$$f_{1t}(\beta) = \begin{bmatrix} \mathcal{V}^k_{t+1,t+2} - \mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right] \\ (\mathcal{V}^k_{t+1,t+2} - \mathrm{E}\left[\mathcal{V}^k_{t+1,t+2}|\mathcal{G}_t\right])\mathcal{V}^k_{t-1,t}, \; k = 1, 2, 3 \\ (p_{t+1} - p_t)\mathcal{V}_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)\mathcal{V}_{t+1,t+2}|\mathcal{G}_t\right] \\ ((p_{t+1} - p_t)\mathcal{V}_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)\mathcal{V}_{t+1,t+2}|\mathcal{G}_t\right])(p_{t-1} - p_{t-2})\mathcal{V}_{t-1,t} \\ (p_{t+1} - p_t)^2\mathcal{V}_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)^2\mathcal{V}_{t+1,t+2}|\mathcal{G}_t\right] \\ ((p_{t+1} - p_t)^2\mathcal{V}_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)^2\mathcal{V}_{t+1,t+2}|\mathcal{G}_t\right])(p_{t-1} - p_{t-2})^2\mathcal{V}_{t-1,t} \\ (p_{t+1} - p_t)\mathcal{V}^2_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)\mathcal{V}^2_{t+1,t+2}|\mathcal{G}_t\right] \\ ((p_{t+1} - p_t)\mathcal{V}^2_{t+1,t+2} - \mathrm{E}\left[(p_{t+1} - p_t)\mathcal{V}^2_{t+1,t+2}|\mathcal{G}_t\right])(p_{t-1} - p_{t-2})\mathcal{V}^2_{t-1,t} \end{bmatrix}. \tag{6}$$

4. Using implied volatilities to link objective and risk-neutral parameters

Implied volatilities are usually computed by inverting the Black–Scholes formula, but they can also be defined for more elaborate pricing models involving additional parameters. Inversion of the pricing formula in the volatility parameter can then only be done for given values of these parameters. The value of implied volatility will therefore depend both on the option price and the parameter values.

However, semi-closed form option pricing formulas are generally difficult to invert and one has to use numerical procedures which are computationally intensive and whose precision has to be controlled. Moreover, implementing integral solutions such as the Heston's formula can be very delicate due to divergences of the integrand in regions of the parameter space. One way to avoid both problems is to rewrite option pricing formulas as power series around values of the parameters for which the model can be analytically solved (*i.e* it has an explicit form in terms of elementary and special functions; not an integral one). This avenue is followed by Lewis (2000).[8]

4.1. Series expansions and inversion of option pricing formulas

Since option prices are continuously differentiable at any order in the volatility of volatility parameter γ, one can expand the pricing formula around a fixed γ, which we will set to zero, as it corresponds to a deterministic volatility model that we can solve analytically. Generally, at date t options will have prices $c(S_t, V_t, K, t, T, r, \beta^*)$, where S_t and V_t are the underlying asset's price and volatility, K the strike price, T the expiration date, and β^* are the parameters of the risk-neutral distribution. The Taylor expansion of c around $\gamma = 0$ has the general form

$$c(S_t, V_t, K, t, T, r, \beta^*) = \sum_{j=0}^{\infty} \mu_j(S_t, V_t, K, t, T, r, \beta^*_{-\gamma})\gamma^j, \tag{7}$$

where $\beta^*_{-\gamma}$ denotes the vector of risk-neutral parameters without the volatility of volatility coefficient γ.

A series expansion like (7) represents a rapid and simple option price computation tool, and it also is usually straightforward to invert, so that one can define implied volatilities and compute them very efficiently. Option prices being strictly increasing functions of the volatility, the inversion is always possible. Assume that given an observed option price c_t^{obs}, the volatility V^{imp} admits the expansion

$$V^{imp}(S_t, c_t^{obs}, K, t, T, r, \beta^*) = \sum_{j=0}^{\infty} \nu_j(S_t, c_t^{obs}, K, t, T, r, \beta^*_{-\gamma})\gamma^j.$$

Such a development makes sense, because if γ goes to zero, the leverage model becomes a deterministic volatility model (in fact, if κ^* goes to zero, we recover the Black–Scholes model). If the observed option price is correctly priced by (7) with V_t set to the implied volatility $V^{imp}(S_t, c_t^{obs}, K, t, T, r, \beta^*)$, then

$$\begin{aligned} c_t^{obs} &= c[S_t, V^{imp}(S_t, c_t^{obs}, K, t, T, r, \beta^*), K, t, T, r, \beta^*] \\ &= \sum_{j=0}^{\infty} \mu_j \left[S_t, \sum_{k=0}^{\infty} \nu_k(S_t, c_t^{obs}, K, t, T, r, \beta^*_{-\gamma})\gamma^k, K, t, T, r, \beta^*_{-\gamma}\right]\gamma^j. \end{aligned}$$

Since we have explicit expressions for the μ_js, we can Taylor expand the right-hand side and collect terms by powers of γ. By doing so, we can define the coefficients $\tilde{\nu}_j(S_t, c_t^{obs}, K, t, T, r, \beta^*_{-\gamma}, \{\nu_k\})$ such that

$$c_t^{obs} = \sum_{j=0}^{\infty} \tilde{\nu}_j(S_t, c_t^{obs}, K, t, T, r, \beta^*_{-\gamma}, \{\nu_k\})\gamma^j.$$

This equation is solved by imposing the conditions

$$\tilde{\nu}_0 = c_t^{obs}$$
$$\tilde{\nu}_j = 0 \quad \forall j \geq 1,$$

which form a triangular system of polynomial equations for the ν_k. This system is easily solved order by order in k. Notice that a similar result could be obtained by starting from a Taylor expansion of

[6] Due to the MA(1) structure of the error terms in (5), the optimal weighting matrix for GMM estimation entails the estimation of the variance and only the first-order autocorrelations of the moment conditions.

[7] Bollerslev and Zhou (2002) express these moment conditions in terms of the log price of the asset, which is nonstationary (see (Bollerslev and Zhou, 2002), footnote 17 on p. 61). Attempts to run the GMM estimation using the conditions based on the price instead of the returns tended to exhibit an erratic behavior: estimates on the boundary of the admissible region, many error flags raised by the optimization software, and so on.

[8] Medvedev and Scaillet (2007) develop similar Taylor expansions for short-term implied volatilities for jump-diffusion stochastic volatility models.

the usual Black–Scholes implied volatility, denoted by V_{BS}, instead of (7).

We now provide an illustration of the inversion methodology with the leverage model, which encompasses the no-leverage one as a special case. The starting point is to derive the Taylor expansion of the call price, or equivalently of the corresponding Black–Scholes implied volatility. To this end, we follow Lewis (2000, chap. 3). To simplify the notation, we will keep the underlying asset price and volatility as the only explicit arguments in the formulas of option prices and Black–Scholes volatilities; also, let us denote with $\mathrm{E}_t(\cdot)$ the conditional expectation operator $\mathrm{E}(\cdot|S_t, V_t)$, taken with respect to the risk-neutral distribution.

Romano and Touzi (1997) have shown that, under the stochastic volatility process (1), the call option price $c(S_t, V_t)$ can be written as

$$c(S_t, V_t) = \mathrm{E}_t\left[c_{BS}(S_t\Omega_{t,T}, (1-\rho^2)\overline{\mathcal{V}}_{t,T})\right], \quad (8)$$

where

$$\Omega_{t,T} = \exp\left(\rho\int_t^T \sqrt{V_s}\mathrm{d}W_s^2 - \frac{1}{2}\rho^2\int_t^T V_s\mathrm{d}s\right) \quad \text{and}$$

$$\overline{\mathcal{V}}_{t,T} = \frac{1}{T-t}\int_t^T V_s\mathrm{d}s.$$

Using the Girsanov theorem and some simple stochastic calculus it easy to show that

$$\mathrm{E}_t\Omega_{t,T} = 1 \quad \text{and} \quad \mathrm{E}_t\overline{\mathcal{V}}_{t,T} = \theta^* + (V_t - \theta^*)\frac{1-\mathrm{e}^{-\kappa^*(T-t)}}{\kappa^*(T-t)}.$$

We can now expand the option pricing formula around these mean values[9]:

$$c(S_t, V_t) = \sum_{p,q=0}^{\infty}\frac{1}{p!q!}\frac{\partial^{p+q}c_{BS}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T})}{\partial S^q\partial V^p}S_t^q(1-\rho^2)^p\mathrm{E}_t \times\left[(\Omega_{t,T}-1)^q(\overline{\mathcal{V}}_{t,T} - \mathrm{E}_t\overline{\mathcal{V}}_{t,T})^p\right].$$

Defining

$$R_{pq}(S, V, T) \equiv S^q\frac{\partial^{p+q}c_{BS}(S, V)}{\partial V^p\partial S^q}\left[\frac{\partial c_{BS}(S, V)}{\partial V}\right]^{-1},$$

this can be rewritten as

$$c(S_t, V_t) = \frac{\partial c_{BS}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T})}{\partial V}\sum_{p,q=0}^{\infty}\frac{(1-\rho^2)^p R_{pq}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T})}{p!q!}\mathrm{E}_t \times\left[(\Omega_{t,T}-1)^q(\overline{\mathcal{V}}_{t,T} - \mathrm{E}_t\overline{\mathcal{V}}_{t,T})^p\right].$$

Let $\tau = T - t$ denote the time to maturity of the option. Lewis (2000) shows that the first values of R_{pq} can be expressed in terms of the moneyness $X = \log[S/K\mathrm{e}^{-r\tau}]$ and the expected integrated volatility $\mathrm{E}_t\mathcal{V}_{t,T}$ as follows:

$$R_{20} = \tau\left[\frac{X^2}{2(\mathrm{E}_t\mathcal{V}_{t,T})^2} - \frac{1}{2\mathrm{E}_t\mathcal{V}_{t,T}} - \frac{1}{8}\right]$$

$$R_{11} = -\frac{X}{\mathrm{E}_t\mathcal{V}_{t,T}} + \frac{1}{2}$$

$$R_{12} = \left[\frac{X^2}{(\mathrm{E}_t\mathcal{V}_{t,T})^2} - \frac{X+1}{\mathrm{E}_t\mathcal{V}_{t,T}} + \frac{1}{4}\right]$$

$$R_{22} = \tau\left[\frac{X^4}{2(\mathrm{E}_t\mathcal{V}_{t,T})^4} - \frac{X^3+6X^2}{2(\mathrm{E}_t\mathcal{V}_{t,T})^3} + \frac{3(X+1)}{2(\mathrm{E}_t\mathcal{V}_{t,T})^2} + \frac{X}{8\mathrm{E}_t\mathcal{V}_{t,T}} - \frac{1}{32}\right].$$

It is also convenient to define the following quantities:

$$J_1 = \rho\frac{[\theta^* + (1+\kappa^*\tau)(\theta^* - V_t)]\mathrm{e}^{-\kappa^*\tau} + [\kappa^*\tau - 2]\theta^* + V_t}{(\kappa^*)^2}$$

$$J_3 = \frac{\theta^* - 2V_t + \mathrm{e}^{2\kappa^*\tau}[(-5+2\kappa^*\tau)\theta^* + 2V_t] + 4\mathrm{e}^{\kappa^*\tau}[\theta^* + \kappa^*\tau\theta^* - \kappa^*\tau V_t]}{4\mathrm{e}^{2\kappa^*\tau}(\kappa^*)^3}$$

$$J_4 = \rho^2\frac{[6 + 2\mathrm{e}^{\kappa^*\tau}(-3+\kappa^*\tau) + \kappa^*\tau(4+\kappa^*\tau)]\theta^* + [-2 + 2\mathrm{e}^{\kappa^*\tau} - \kappa^*\tau(2+\kappa^*\tau)]V_t}{2\mathrm{e}^{\kappa^*\tau}(\kappa^*)^3}$$

if $\kappa^* \neq 0$, and

$$J_1 = \frac{\rho V_t\tau^2}{2}, \qquad J_3 = \frac{V_t\tau^3}{6}, \qquad J_4 = \frac{\rho^2 V_t\tau^3}{6}$$

if $\kappa^* = 0$. Using these expressions, one can approximate the call option price by truncating the expansion at the order two:

$$c(S_t, V_t) = c_{BS}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T}) + \frac{J_1R_{11}}{\tau}\frac{\partial c_{BS}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T})}{\partial V} + \left[\frac{J_4R_{12}}{\tau} + \frac{2J_3R_{20} + J_1^2R_{22}}{2\tau^2}\right]\frac{\partial c_{BS}(S_t, \mathrm{E}_t\overline{\mathcal{V}}_{t,T})}{\partial V}\gamma^2.$$

Lewis (2000) shows that the neglected terms are $O(\gamma^3)$. Alternatively, by starting from a series expansion of the Black–Scholes implied volatility, evaluating the relevant coefficients and truncating at the order two, one arrives at the following approximation:

$$V_{BS}(S_t, V_t) = \mathrm{E}_t\overline{\mathcal{V}}_{t,T} + \frac{J_1R_{11}}{\tau}\gamma + \left[\frac{J_2 + J_4R_{12}}{\tau} + \frac{2J_3R_{20} + J_1^2(R_{22} - R_{11}^2R_{20})}{\tau^2}\right]\gamma^2$$

whose approximation error is still $O(\gamma^3)$. Lewis (2000), however, shows that for plausible values of the arguments the latter approximation is more accurate than the former, in particular for options far out-of-the-money and far in-the money. Following his suggestion, in our empirical analysis we use the Black–Scholes implied volatility series expansion.

We now turn to the extraction of an implied volatility coherent with the leverage model from the observed Black–Scholes implied volatility. Specifically, we want to identify the coefficients ν_i of the approximation

$$V^{imp}(S_t, V_{BS,t}, \beta^*) = \nu_0(S_t, V_{BS,t}, \beta^*_{-\gamma}) + \nu_1(S_t, V_{BS,t}, \beta^*_{-\gamma})\gamma + \nu_2(S_t, V_{BS,t}, \beta^*_{-\gamma})\gamma^2$$

whose errors are $O(\gamma^3)$, and where once again we suppress for simplicity the dependence of V^{imp} and ν_i from the option's characteristics, while keeping explicit only the observed underlying price S_t, the option's Black–Scholes volatility $V_{BS,t}$, and the risk-neutral parameters β^*. Computing the coefficients ν_0, ν_1 and ν_2 is a straightforward but tedious exercise. As pointed out earlier, all we need to do is to insert this series expansion for V_t in the series expansion for V_{BS}, further expand with respect to γ around 0, and recursively compute the values of the coefficients that ensure the equality between the truncated series above and the observed Black–Scholes volatility.[10]

4.2. Moment conditions for implied volatilities

Given daily option prices and some initial values for the objective parameters β_0 and the risk premium λ_0, we are able to construct an implied volatility time series. If the risk premium was suitably chosen and the objective parameters were the true ones, every option on a given day would have the same implied volatility, so that the implied volatility surface would be flat, with a value

9 Hull and White (1987) propose a similar series expansion for the case in which the stochastic volatility is independent of the stock price and follows a geometric Brownian motion.

10 The coefficients ν_0, ν_1, ν_2 of the inverse series for the spot volatility $V_t = \nu_0 + \nu_1\gamma + \nu_2\gamma^2$ are provided in a document available on the following Web site: http://www.sceco.umontreal.ca/renegarcia/articles.htm.

equal to the point-in-time volatility V_t. Since this will not usually be the case, one could use some daily mean value of the observed implied volatilities $V_t^{imp}(\beta_0, \lambda_0)$ to generate the implied volatility time series.[11]

For the no-leverage model, we then consider the daily volatility series and use the moments of the instantaneous volatility in the moment conditions $f_{2t}(\beta)$. Each of these conditional moment conditions are used with two instrumental variables (which are a constant and V_t^k for the moment condition involving $\mathrm{E}\left[V_{t+1}^k|\mathcal{F}_t\right]$), resulting in six unconditional moment conditions:

$$f_{2t}(V, \beta) = \begin{bmatrix} (V_{t+1}^k - \mathrm{E}(V_{t+1}^k|\mathcal{F}_t)) \\ (V_{t+1}^k - \mathrm{E}(V_{t+1}^k|\mathcal{F}_t))V_t^k \end{bmatrix}, \quad k = 1, 2, 3. \tag{9}$$

These moment conditions are used in (3) to obtain estimates of the β and λ parameters with a joint GMM procedure.

We proceed in the same way for the leverage model and use a similar set of moment conditions as (6) with V replacing $\mathcal{V}$ and $\mathcal{F}$ replacing $\mathcal{G}$. Expressions for the cross-moments are provided in Appendix B.

5. A Monte Carlo study

In order to assess how the GMM estimation method described in the previous sections performs, we conducted a Monte Carlo study for both the leverage and no-leverage models. It should be mentioned that any affine model that admits a formulation of the option pricing formula in terms of a power series could be considered with the same methodology.

5.1. The no-leverage model

We study the same three sets of parameters chosen by Bollerslev and Zhou (2002), in order to be able to compare our respective results for the estimation of $\hat{\beta}$. We combine them with different values of the volatility risk premium λ[12] to obtain four parameter sets A, B, C and D.

Note that these are daily parameters. As in Bollerslev and Zhou, we normalized θ so that the yearly volatility is $\sqrt{240 \times \theta}$, with 240 being the number of days per year we chose. This means that a yearly volatility of 7.74% is associated with $\theta = 0.25$. Each day is further subdivided in 80 five-min periods. The quadratic variations $\mathcal{V}_{t,t+1}$ are aggregated over these 80 periods, giving daily integrated volatilities, whereas the option prices are computed from the mid-day price and spot volatility of the underlying asset (i.e. the 40th observation in each day). In turn, each 5-min interval is actually subdivided in ten 30 s subintervals, and the SDE is simulated using the finest grid.

Under the no-leverage model, option prices can be expressed as mean values of the Black–Scholes price over the integrated volatility distribution.[13] Option prices are obtained by simulating volatility trajectories to approximate the integrated volatility distribution. We wanted to avoid using our expansion for the option pricing formula in order to validate its use. However, if enough trajectories are simulated, the simulated and series expansion price are almost indistinguishable (at least in the region where the expansion is valid and has enough precision).[14]

The risk premium structure is chosen as in the Heston's (1993) paper, that is, risk-neutral (denoted by stars) and objective parameters are related by $\kappa^* = \kappa - \lambda$, $\theta^*\kappa^* = \theta\kappa$ and $\gamma^* = \gamma$. Usually, one would expect λ to be positive, so that the asymptotic volatility is higher, meaning that option prices will also be.

The GMM estimation procedure was conducted with a Newey-West kernel with a lag length of two.[15] Results of the estimations are provided in Table 1. Statistics were obtained by estimating parameters over 5000 independent sets of 4-year data (960 observations).

A comparison of the $\hat{\beta}$ estimates with the GMM estimates of Bollerslev and Zhou (2002) reveals that the RMSE are not universally smaller with the moment conditions that we specified. The difference between the two sets is the introduction in our estimators of the third-moment conditions, while in their selection of instruments they included the lagged squared integrated volatility. It appears that the root mean-square error (RMSE) with our moment conditions is lower for the mean-reversion parameter κ, while it is higher for θ. The evidence is mixed for γ.

The third-moment condition seems to perform better for higher volatility of volatility but worse for highly persistent processes. A closer look at the third-moment formula of integrated volatility in Appendices A and B shows that when κ is close to zero a number of terms involving θ disappear in N_T, hence a potential loss of identification for this parameter. Moreover, a γ close to zero will potentially reduce the information about the other parameters that appear inside the brackets since higher powers of γ pre-multiply the whole expressions in the higher-order moments.

The RMSE of γ generally deteriorates with the estimator $\hat{\psi}$. The added moment conditions where implied volatility is recovered through a Taylor expansion around $\gamma = 0$ may make this parameter harder to recover.

Finally, the volatility risk premium λ is also nicely recovered. Its RMSE for $\hat{\psi}$ remains quite small, except maybe for the last configuration of parameters (Panel D). This may be due to the fact that in this case the process is close to violating the condition making the zero boundary inaccessible ($2\kappa\theta \geq \gamma^2$), with a high volatility of volatility parameter. The error is also relatively large for the other parameters in this case.

5.2. The leverage model

The leverage model requires prices and spot volatilities to be observed at the beginning of the period over which the quadratic variation is computed. For this reason, we modified the sampling scheme to use opening prices instead of mid-day prices. Moreover, the previous Monte Carlo strategy does not work in the leverage model.[16] Therefore, we evaluate option prices by adding a random

11 Whether strongly out or in-the-money options should be included in this mean value computation can be debated. Nevertheless, we chose to do so both for the Monte Carlo generated option prices and for the empirical applications. In the latter case, we also compared the results to an approach where we chose only one option per day, either the nearest to the money or the highest volume one.

12 Under CRRA preferences, where a representative investor has a power utility over wealth, a zero correlation between returns and volatility innovations will imply a zero risk premium. However, volatility could be correlated with aggregate consumption and carry a non-zero price of risk. Also, a non-zero premium could be rationalized in an international model in which the volatility price of risk is linked to the exchange risk.

13 Explicit expressions for the option prices and implied volatility series up to the sixth order in γ can be found in the document posted on the above-mentioned web site.

14 Lewis (2000), on page 80, provides a graph in which it can be seen that the approximation error associated with the expansion increases with the moneyness.

15 We checked that results were quite insensitive to this choice.

16 Basically, the problem is that the in-the-money option prices simulated in the way we proceeded violate often the lower bound $C \geq S - K\mathrm{e}^{-r\tau}$, which makes it impossible to compute the associated Black and Scholes volatility. In the no-leverage model this problem was absent because C could be computed as the expectation of $C_{BS}(S, \mathcal{V}_{t,t+\tau})$ with respect to the distribution of the integrated volatility. A similar expression also exists for the leverage model (see formula (8)), but in this case C_{BS} is not evaluated in S, but in S rescaled by a function of the volatility trajectory. Of course increasing the number of trajectories may attenuate the issue, but it does not provide a complete solution, and it significantly increases the computational burden.

Table 1
Parameter estimates of the no-leverage model.

	$\hat{\beta}$			$\hat{\psi}$			
	κ	θ	γ	κ	θ	γ	λ
Parameter set A							
True value	0.1000	0.2500	0.1000	0.1000	0.2500	0.1000	0.0500
Mean	0.0990	0.2502	0.1007	0.1027	0.2463	0.1011	0.0537
Median	0.0977	0.2493	0.1008	0.1026	0.2458	0.0999	0.0537
Std. error	0.0186	0.0166	0.0064	0.0096	0.0167	0.0078	0.0070
RMSE	0.0186	0.0166	0.0064	0.0099	0.0171	0.0078	0.0080
Parameter set B							
True value	0.1000	0.2500	0.1000	0.1000	0.2500	0.1000	0.0200
Mean	0.0996	0.2493	0.1008	0.1038	0.2444	0.0994	0.0244
Median	0.0983	0.2487	0.1009	0.1040	0.2439	0.0994	0.0242
Std. error	0.0187	0.0164	0.0062	0.0097	0.0163	0.0060	0.0070
RMSE	0.0187	0.0164	0.0063	0.0104	0.0172	0.0060	0.0083
Parameter set C							
True value	0.0300	0.2500	0.1000	0.0300	0.2500	0.1000	0.0100
Mean	0.0344	0.2537	0.1003	0.0350	0.2355	0.0953	0.0163
Median	0.0343	0.2388	0.1005	0.0365	0.2362	0.0973	0.0152
Std. error	0.0125	0.0848	0.0076	0.0095	0.0354	0.0157	0.0071
RMSE	0.0133	0.0849	0.0076	0.0107	0.0382	0.0164	0.0095
Parameter set D							
True value	0.1000	0.2500	0.2000	0.1000	0.2500	0.2000	0.0500
Mean	0.1069	0.2463	0.1966	0.0999	0.2383	0.1876	0.0575
Median	0.1062	0.2430	0.1966	0.1029	0.2368	0.1906	0.0571
Std. error	0.0225	0.0406	0.0104	0.0259	0.0365	0.0304	0.0160
RMSE	0.0235	0.0408	0.0109	0.0259	0.0383	0.0329	0.0177

Note: The estimator $\hat{\beta}$ is based on the set of moment conditions f_{1t} defined in (5). For the estimator $\hat{\psi}$, we use in addition the set of moment conditions f_{2t} defined in (9) for a joint estimation as defined in (3).

error term to the series expansion formula in implied volatility. We studied the characteristics of the error terms implicitly introduced in the option prices by the MC strategy used to evaluate them in the no-leverage model, and found out that they looked fairly similar to $\mathcal{N}(0, \omega^2)$ added to the implied B&S volatilities, with $\omega = 4 \times 10^{-7}$. This is the distribution from which the random noises in option prices (via the implied B&S volatilities) are drawn in the experiments concerning the leverage model.

Optimal weighting matrices have been computed in a second step using Newey-West with two lags for both estimators $\hat{\beta}$ and $\hat{\psi}$. We considered the same four sets of values for β and λ as in the no-leverage experiments, but only one value for the parameter ρ equal to −0.5. The results are provided in Table 2.

For each configuration of the parameters, the MC experiment consists of 5000 replications of 960 observations each. In order to identify the correlation parameter ρ, these experiments considered six additional cross-moment conditions in addition to those used in the experiments conducted in no-leverage model for the $\hat{\beta}$ estimator. An additional set of six moment conditions defined in (9) is used for the estimator $\hat{\psi}$. For both estimators, ρ is clearly the most difficult parameter to estimate with the uniformly highest – by a very significant margin – RMSE among all the parameters. In terms of bias (both mean and median), the estimator $\hat{\beta}$ tends to underestimate the leverage, roughly by 20% of the true value. This bias disappears for the estimator $\hat{\psi}$.[17] The performance seems to worsen a bit when γ increases (parameter set 4a). For the parameters κ,θ and γ the RMSEs are generally slightly higher than in the no-leverage case.

An important issue for the inversion procedure recovering the implied volatility is the positivity of the extracted volatility. Since the inversion rests on a truncated expansion, it is legitimate to ask whether the recovered volatility is positive. Although it is impossible to ensure this positivity when such a truncation takes place, we verified in our Monte Carlo that even for a second-order expansion none of the extracted volatilities were negative. However, in our experiments there are no misspecification issues since the data are simulated under the true model. With actual data, this will most likely not be the case. To investigate this aspect of the problem, we set up some Monte Carlo experiments identical to those illustrated above, but in which some errors were added to the simulated option prices through the corresponding Black–Scholes volatilities.[18] These experiments showed that for reasonable values of the variance of the measurement errors, only a small percentage of recovered volatilities were negative. Moreover, these negative volatilities did not have any significant impact on the bias of the estimators. This suggests that in most applications negative recovered volatilities are not a crucial issue, and can be neglected without harm.

Alternatively, we could impose the positivity of the extracted volatilities by suitably constraining the parameter space during the optimization of the GMM objective function. In practice, this amounts to impose one nonlinear constraint on the parameters for each observed option. It should be noted that this problem is fre-

[17] One might think that this behavior is due to the fact that the quadratic variations are only estimates of the true integrated volatilities, and that the measurement errors they embed may lead to the underestimation highlighted above. To check this hypothesis, we estimated β using also the true integrated volatilities. It turns out that the underestimation of ρ is roughly the same as before. Hence, it is not due to the measurement error in quadratic variations. The estimates of β based on the true spot volatilities, $\hat{\beta}_V$, also exhibit a downward bias with respect to ρ, albeit less pronounced. The persistence in the volatility process may cause a finite sample bias. The absence of bias when adding information through option prices seems to confirm this interpretation.

[18] For each option, the random errors added to the simulated Black–Scholes volatility were drawn from a zero-mean Gaussian distribution with a standard error set at some percentage α of the Black–Scholes volatility of the same option evaluated at the unconditional instantaneous volatility. We considered values of α ranging from 5% to 25%.

Table 2
Parameter estimates of the leverage model.

Estimator	Parameter	Parameter set 1a				Parameter set 2a			
		True Value	Mean	Median	RMSE	True Value	Mean	Median	RMSE
$\hat{\beta}$	κ	0.100	0.103	0.100	0.023	0.100	0.104	0.101	0.024
	θ	0.250	0.247	0.247	0.016	0.250	0.247	0.246	0.017
	γ	0.100	0.101	0.101	0.007	0.100	0.101	0.102	0.007
	ρ	−0.500	−0.455	−0.453	0.135	−0.500	−0.451	−0.449	0.135
$\hat{\psi}$	κ	0.100	0.099	0.100	0.011	0.100	0.099	0.100	0.010
	θ	0.250	0.247	0.246	0.020	0.250	0.246	0.246	0.019
	γ	0.100	0.097	0.098	0.008	0.100	0.097	0.098	0.008
	ρ	−0.500	−0.498	−0.500	0.086	−0.500	−0.498	−0.500	0.082
	λ	0.050	0.051	0.051	0.007	0.020	0.021	0.021	0.007
Estimator	Parameter	Parameter set 3a				Parameter set 4a			
		True Value	Mean	Median	RMSE	True Value	Mean	Median	RMSE
$\hat{\beta}$	κ	0.030	0.036	0.034	0.015	0.100	0.112	0.109	0.027
	θ	0.250	0.241	0.238	0.051	0.250	0.241	0.239	0.032
	γ	0.100	0.101	0.101	0.007	0.200	0.198	0.197	0.011
	ρ	−0.500	−0.424	−0.425	0.223	−0.500	−0.445	−0.443	0.155
$\hat{\psi}$	κ	0.030	0.029	0.030	0.008	0.100	0.105	0.109	0.026
	θ	0.250	0.239	0.249	0.059	0.250	0.228	0.224	0.057
	γ	0.100	0.091	0.095	0.017	0.200	0.184	0.190	0.030
	ρ	−0.500	−0.498	−0.500	0.129	−0.500	−0.488	−0.493	0.126
	λ	0.010	0.019	0.013	0.028	0.050	0.065	0.063	0.027

Note: The estimator $\hat{\beta}$ is based on the set of moment conditions f_{1t} defined in (6). For the estimator $\hat{\psi}$, we use a similar set of moment conditions as (6) with V replacing $\mathcal{VF}$ replacing $\mathcal{G}$. We use in addition the set of moment conditions f_{2t} defined in (9) for a joint estimation as defined in (3).

quently observed when the model to be estimated features an observable variable which is a function of some latent variable and the parameters, which is typically the case in derivative pricing and term structure modeling if one ignores the presence of measurement errors. When the parameters are estimated by maximum likelihood, this hugely complicates the estimation because the objective function may become undefined when the extracted latent variable violates some boundary condition. In our framework, the problem is easier to solve because the moment conditions are well-defined even for negative volatilities. Nonlinear constraints can be imposed directly on the objective function through a smooth penalty. We investigated this avenue at length, but the results were disappointing because the parameter estimates were much less reasonable, with mostly negative estimates of the volatility risk premium parameter λ and too low values of the volatility of volatility coefficient γ. Therefore we decided to ignore the negativity of the recovered volatilities since our Monte Carlo showed that it is essentially without consequences. We will see that our empirical results are fairly reasonable and compare well to estimates obtained in other studies with different estimation methodologies.

6. Empirical illustrations

This section provides two applications of the proposed estimators, For the no-leverage model, we use a sample of high-frequency data on the Deutsche mark–US dollar (DM/$) futures contract and of daily call options on the same contracts. The spot price of the S&P500 index and daily call prices on the index are used to estimate the parameters of the leverage model.

6.1. *Estimation of the no-leverage model*

Data on exchange rate are often used to estimate a no-leverage model. In Bollerslev and Zhou (2002) and Bates (1996) estimates of the parameter ρ are close to zero for the DM/$ exchange rate data either on the exchange rate themselves or on the corresponding currency options, confirming the appropriateness of a no-leverage specification.

Our data are for the DM/$ futures contract and were obtained from Tickdata.com. We form five-min log returns for the DM/$ futures contract on the spot exchange rate. The sample begins in January 1984 and ends in December 1998, for a total of 3753 daily observations. For each date 80 five-min intervals were observed, which were used to compute the quadratic variation measures.

The option data set consists of daily data on options on DM/$ futures from the Chicago Mercantile Exchange (CME). At each date, a number of options were observed on the DM/$ futures, corresponding to different strike prices and maturity dates. Options with less than 15 days to maturity and with low transaction volume (less than 10 contracts) were dropped to avoid the inclusion of outliers. The number of options observed at each date is not constant – rather, it initially increases and subsequently declines towards the end of the sample – but it is never smaller than five. Notice that since our observations consist of options on a futures contract, the basic Hull and White formula and the volatility of volatility series expansions for the spot volatility have to be adjusted accordingly. In practice, as shown by Black (1976), if we denote by t the current date and by T the maturity date of the option, it is sufficient to use the standard pricing formula but with the discounted price $\exp[-r_t(T-t)]S_t$ of the underlying asset instead of the price S_t. As a risk-free interest rate we used the three-month Treasury bill rate on the secondary market.

The top panel of Table 3 reports the GMM estimates of the parameters of the no-leverage model on the sample of returns and option prices on the DM/$ futures. First, we report the estimates obtained with $\hat{\beta}$, which exploits 6 moment conditions for the integrated volatility, and hence does not allow us to estimate the risk premium parameter λ, and then the estimates with $\hat{\psi}$, which is the joint estimator of all the parameters of the model, and is based on 12 moment conditions, 6 for the integrated volatility and 6 for the filtered spot volatilities computed using the inverse series outlined in the previous sections.

For $\hat{\psi}$ we chose three ways to exploit the information available in option prices. First, we included all acceptable options according to the criteria set above, for a total of 96,599. Second, as in Pan (2002), we selected each day the option closest to the money. Finally, we included the option with the highest volume each day.

Table 3
Estimation results with financial data.

DM/US$ futures data and the no-leverage model						
Est.	#opt.	κ	θ	γ	λ	pv $\mathcal{J}$
$\widehat{\beta}$	–	0.094 (0.043)	0.292 (0.047)	0.222 (0.054)	–	0.158
$\widehat{\psi}$	96 599	0.098 (0.003)	0.293 (0.018)	0.233 (0.012)	0.088 (0.003)	0.005
$\widehat{\psi}$	3753 (Pan)	0.112 (0.031)	0.295 (0.036)	0.246 (0.012)	0.073 (0.021)	0.256
$\widehat{\psi}$	3753 (Vol.)	0.066 (0.002)	0.313 (0.031)	0.219 (0.020)	0.050 (0.006)	0.080

S&P500 index data and the leverage model							
Est.	#opt.	κ	θ	γ	ρ	λ	pv $\mathcal{J}$
$\widehat{\beta}$	–	0.173 (0.053)	0.809 (0.081)	0.713 (0.133)	−0.165 (0.100)	–	0.001
$\widehat{\psi}$	178 916	0.027 (0.050)	0.844 (0.073)	0.186 (0.049)	−0.215 (0.103)	0.017 (0.012)	0.000
$\widehat{\psi}$	2517 (Pan)	0.070 (0.011)	0.800 (0.075)	0.414 (0.049)	−0.319 (0.103)	0.049 (0.012)	0.000
$\widehat{\psi}$	2517 (Vol.)	0.032 (0.006)	0.718 (0.024)	0.336 (0.093)	−0.090 (0.010)	0.022 (0.005)	0.000

In the last two cases, we had a total of 3753 options over the sample. When all options are used for estimation, we extract an average implied volatility, without regard to the moneyness or the liquidity of the options.

Along with the number of options used in the estimation, the table reports the point estimates of the parameters, the associated asymptotic standard errors (computed using a Newey-West weighting scheme with lag two) in parentheses, and the p-value of the corresponding Hansen overidentifying restrictions test.

A first observation is that, for all parameters, the estimates based on futures returns only and those based on futures and option prices are very close to each other when all options are included. Estimates are still close with the at-the-money options, but differ more for the highest-volume options, especially for the κ parameter.

The estimates of the volatility risk premium are of the right sign and of reasonable magnitude. Moreover, the estimated risk premium is higher when all options are considered, which is intuitive since illiquid options are included along with the more liquid ones. It is also true of the at-the-money estimate, which is higher than the premium estimated with the highest-volume option.

The model fares well in terms of the overidentifying statistic $\mathcal{J}$. The p-value is always greater than 0.05 except when all options are included. This is a bit surprising since the one-factor stochastic volatility model has not been successful in describing the exchange rate data in previous studies.

Using high-frequency exchange rate data, Bollerslev and Zhou (2002) find that the one-factor stochastic volatility model does not fully capture the dynamics of the daily DM/$ volatility. The stationarity condition ($2\kappa\theta \geq \gamma^2$) is violated by their parameter estimates in $\hat{\beta}$ but it is not the case with our estimates. Moreover their estimates of κ, θ and γ are fairly higher than our estimates. This could be justified on several grounds. First, they use data on the spot exchange rate, whereas our application uses observations on a futures contract; second, the sample interval is different (we consider a longer interval, beginning slightly less than two years before and ending slightly more than two years after); and third, and most importantly, their quadratic variation measures are computed from 288 five-min returns over the 24-hour cycle, while ours are based on just 80 five-min intervals, and hence tend to be significantly smaller than theirs. Overall they find that a two-factor volatility factor is better supported by the data.

Bates (1996) used short maturity at- and out-of-the-money Deutsche mark foreign currency options to estimate by maximum likelihood the parameters of several jump-diffusion models, among which the stochastic volatility model we estimated. The estimates are relatively close: for the parameter $\kappa\theta$ our value is 0.0274 to compare with 0.031 for their corresponding parameter α; for γ, 0.233 to compare with 0.284 for their corresponding parameter σ_v; finally, for κ^* annualized, 2.4 to compare with 1.30 for their corresponding parameter β^*. Bates (1996) concludes that there is substantial qualitative agreement between implicit and time-series-based distributions, most notably with regard to implicit volatilities as forecasts of future volatility. Our estimates are in fact very close whether or not we include the information about option prices in the estimation. Although the one-factor volatility model is not rejected by the $\mathcal{J}$ statistic, the estimates suggest a very persistent volatility process with too-high a volatility of volatility γ parameter, which may be justified by the omission of a second volatility factor or a jump component.

6.2. Estimation of the leverage model

While a one-factor stochastic volatility model may have some empirical support for exchange rate data, it is overwhelmingly rejected in the empirical literature that estimated this model for equity index returns. There is a consensus that single volatility factor models do not fit the data (see Andersen et al. (2002), Chernov et al. (2003), Eraker et al. (2002), Pan (2002), among others). Several authors augmented affine SV diffusions with jumps.[19] Other authors[20] have shown, however, that SV models with jumps in returns are not able to capture all the empirical features of observed option prices and returns.

The p-values of the $\mathcal{J}$ statistic confirm the strong rejection of the stochastic volatility model. However, to assess our methodology, it is important to compare our results (reported in the second panel of Table 3) with the estimates produced in several of the above-mentioned studies. The GMM estimates are based on the moment conditions described in Section 4. We follow exactly the same format to construct the integrated volatility and implied volatility measures as in the exchange rate futures case. The high-frequency data for the spot S&P 500 index (to build 80 five-min returns per day) were obtained from Tickdata.com for the period from January 4, 1996 till December 30, 2005, that is 2517 daily observations. The number of call options per day varies from a minimum of 30 to a maximum of 130, for a total of 179,176 options.

The estimates of $\hat{\psi}$ with all options included are close to the posterior mean values obtained with a Bayesian method by Eraker et al. (2003), with time-series data on the index only, over the 1980–1999 period: κ in our study is estimated at 0.027, compared to 0.0231 in theirs, θ at 0.844 versus 0.9052, γ at 0.186 versus 0.1434. The largest difference is for the estimate of ρ (−0.215 instead of −0.3974). Our point estimate is low in absolute value with respect to most studies, especially those using option price data. Pan (2002) and Bakshi et al. (1997) obtain estimates around −0.5. However, in the latter study, the authors also compute a sample time-series correlation between daily S&P 500 index returns and daily changes in the implied volatility of the stochastic volatility model over the period 1988–1991. They obtain an estimate of −0.28, which is closer to our estimate also based on moments of the implied volatility for the leverage model.

When we select one option per day (whether near the money or highest volume), the estimates are robust for κ and θ, but much

[19] See in particular Andersen et al. (2001), Bates (1996), Chernov et al. (2003), Eraker et al. (2003), Pan (2002), among others.

[20] Bakshi et al. (1997), Bates (2000), Chernov et al. (2003) and Pan (2002).

more variable for γ and ρ. The estimated values for γ are much higher, while the absolute values for ρ are respectively on the high and the low side. To gauge the estimated value for the volatility risk premium, we can compare it to Pan (2002), who also used return and option data to estimate a jump-diffusion model (which includes both a volatility risk premium and a jump risk premium). The volatility risk premium is estimated at 7.6, which translates in a daily parameter of 0.03, compared to 0.049 for our estimate.

7. Concluding remarks

In this paper, we proposed a joint estimation procedure of objective and risk-neutral parameters for stochastic volatility models. This approach uses both the high-frequency return information available on an underlying asset and the information on options written on this underlying. We applied this procedure to actual return and option data on exchange rate futures on the Deutsche mark–US Dollar and on the S&P 500 index.

Analytical expressions for the moments of integrated volatility in affine stochastic volatility models enabled us to obtain explicit expansions of the implied volatility, a crucial feature of our procedure. The method is computationally simple since no simulations or numerical function inversions are involved. Many extensions of this work can be envisioned. A better specification for stock returns should incorporate jumps. Hence, developing an estimation procedure for jump-diffusion models appears as a natural extension. Introducing other measures such as bi-power variation (see (Barndorff-Nielsen and Shephard, 2004)) and accounting for microstructure noise are also avenues to be explored.

Appendix A. Computation method and expressions for moments of integrated volatility

The first moment or the conditional expectation of the integrated volatility is:

$$\begin{aligned} \mathrm{E}\left[\mathcal{V}_{t,T}|V_t\right] &= \mathrm{E}\left[\int_t^T V_u \mathrm{d}u \,\middle|\, V_t\right] \\ &= \int_t^T \mathrm{E}\left[V_u|V_t\right]\mathrm{d}u \\ &= \int_t^T \left(\theta + \mathrm{e}^{-\kappa(u-t)}\left(V_t-\theta\right)\right)\mathrm{d}u \\ &\equiv a_{T-t}V_t + b_{T-t}, \end{aligned}$$

with $a_{T-t} \equiv \int_t^T \alpha_{u-t}\mathrm{d}u$ and $b_{T-t} \equiv \int_t^T \beta_{u-t}\mathrm{d}u$.

In order to compute higher moments, let us consider the V_t-dependent random variable $\mathrm{E}\left[\mathcal{V}_{t,T}|V_t\right]$:

$$\mathrm{E}\left[\mathcal{V}_{t,T}|V_t\right] = \int_t^T \mathrm{E}_t\left[V_u\right]\mathrm{d}u.$$

If one defines $G(u,t) = \mathrm{E}_t[V_u]$, it is clear from the law of iterated expectations that $G(u,t)$ is a martingale in t. Thus, Itô's lemma implies that

$$\mathrm{dE}\left[\mathcal{V}_{t,T}|V_t\right] = -V_t\mathrm{d}t + a_{T-t}\gamma\sqrt{V_t}\mathrm{d}W_t.$$

Taking integer powers and expectations on each side we obtain

$$\mathrm{E}_t\left[\left(\mathcal{V}_{t,T} - \mathrm{E}[\mathcal{V}_{t,T}]\right)^n\right] = \mathrm{E}_t\left[\left(\int_t^T a_{T-s}\sqrt{V_s}\mathrm{d}W_s\right)^n\right]$$

where by $\mathrm{E}_t[\cdot]$, we mean $\mathrm{E}[\cdot|V_t]$. This formula gives us a way to construct all central moments. The computation of this integral is however far from trivial. The interested reader will find in Lewis (2001) details of the computation for the third and fourth central moments. We will content ourselves of giving the explicit form of the three first central moments for Feller-like stochastic volatility processes (both models considered here are in that class). The variance has the form[21]:

$$\mathrm{E}[(\mathcal{V}_T - \mathrm{E}(\mathcal{V}_T))^2] = A_T V + B_T$$

where

$$A_T = \frac{2\gamma^2\left(\sinh(T\kappa) - T\kappa\right)}{\mathrm{e}^{T\kappa}\kappa^3}$$

$$B_T = \frac{\gamma^2\left(\theta + 4\mathrm{e}^{T\kappa}\theta\left(1+T\kappa\right) + \mathrm{e}^{2T\kappa}\left(-5\theta + 2T\theta\kappa\right)\right)}{2\mathrm{e}^{2T\kappa}\kappa^3}.$$

The third central moment has the form

$$\mathrm{E}[(\mathcal{V}_T - \mathrm{E}(\mathcal{V}_T))^3] = M_T V + N_T$$

with:

$$M_T = \frac{3\gamma^4\left(-1 + 2\mathrm{e}^{3T\kappa} - 2\mathrm{e}^{T\kappa}\left(1+2T\kappa\right) + \mathrm{e}^{2T\kappa}\left(1 - 2T\kappa\left(1+T\kappa\right)\right)\right)}{2\mathrm{e}^{3T\kappa}\kappa^5}$$

$$N_T = \frac{\gamma^4\left(\theta + 6\mathrm{e}^{T\kappa}(\theta + T\theta\kappa) + 2\mathrm{e}^{3T\kappa}(-11\theta + 3T\theta\kappa) + 3\mathrm{e}^{2T\kappa}\theta(5 + 2T\kappa(3+T\kappa))\right)}{2\mathrm{e}^{3T\kappa}\kappa^5}.$$

Finally, the fourth moment can be shown to be:

$$\mathrm{E}[(\mathcal{V}_T - \mathrm{E}(\mathcal{V}_T))^4] = Q_T V^2 + R_t V + S_T$$

with:

$$Q_T = \frac{12\gamma^4\left(-(T\kappa) + \sinh(T\kappa)\right)^2}{\mathrm{e}^{2T\kappa}\kappa^6}$$

$$\begin{aligned} R_T = \frac{\gamma^4}{\mathrm{e}^{4T\kappa}\kappa^7} &\left(3\theta\kappa\left(-1 + \mathrm{e}^{2T\kappa} - 2\mathrm{e}^{T\kappa}T\kappa\right)\left(1 + 4\mathrm{e}^{T\kappa}\left(1+T\kappa\right)\right.\right. \\ &+ \left.\mathrm{e}^{2T\kappa}\left(-5+2T\kappa\right)\right) + \gamma^2\left(-3 + 15\mathrm{e}^{4T\kappa} - 12\mathrm{e}^{2T\kappa}\right. \\ &\times \left(1+T\kappa\right)\left(1+2T\kappa\right) - 6\mathrm{e}^{T\kappa}\left(2+3T\kappa\right) \\ &\left.\left.- 2\mathrm{e}^{3T\kappa}\left(-6 + T\kappa\left(3 + 2T\kappa\left(3+T\kappa\right)\right)\right)\right)\right) \end{aligned}$$

$$\begin{aligned} S_T = \frac{\gamma^4\theta}{4\mathrm{e}^{4T\kappa}\kappa^7} &\left(3\theta\kappa\left(1 + 4\mathrm{e}^{T\kappa}\left(1+T\kappa\right) + \mathrm{e}^{2T\kappa}\left(-5+2T\kappa\right)\right)^2\right. \\ &+ \gamma^2\left(3 + 24\mathrm{e}^{T\kappa}\left(1+T\kappa\right) + 3\mathrm{e}^{4T\kappa}\left(-93+20T\kappa\right)\right. \\ &+ 12\mathrm{e}^{2T\kappa}\left(7 + 2T\kappa\left(5+2T\kappa\right)\right) \\ &\left.\left.+ 8\mathrm{e}^{3T\kappa}\left(21 + T\kappa\left(27 + 2T\kappa\left(6+T\kappa\right)\right)\right)\right)\right). \end{aligned}$$

The explicit expressions for these moments are the ones used in the GMM estimation and option pricing expansions.

Appendix B. Expressions for the cross-moments in the leverage model

For the cross-moments, we use the expressions in Box I. These moments – along with the first three marginal moments of V_t – are used twice: once with the spot volatilities filtered from the option prices, and once with the observed integrated volatilities. In the latter case, we need the expressions of V_t and V_t^2 in terms of $\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{F}_t)$ and $\mathrm{E}(\mathcal{V}_{t,t+1}^2|\mathcal{F}_t)$. The solutions are given by:

$$V_t = \frac{\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{F}_t) - b}{a}$$

$$V_t^2 = \frac{Ab + ab^2 - aB - (A - 2ab)\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{F}_t) + a\mathrm{E}(\mathcal{V}_{t,t+1}^2|\mathcal{F}_t)}{a^3}$$

where a, b, A and B are the coefficients appearing in the conditional expectation and variance of the integrated volatility. Substituting

[21] For simplicity of notation we write the moments at time $t = 0$, and let $V = V_0$.

$$\mathrm{E}[(p_{t+1}-p_t)V_{t+1}|\mathcal{F}_t]=\frac{\gamma(V_t\kappa+\theta(-1+\mathrm{e}^{\kappa}-\kappa))\rho}{\mathrm{e}^{\kappa}\kappa}$$

$$\begin{aligned}\mathrm{E}[(p_{t+1}-p_t)^2V_{t+1}|\mathcal{F}_t] = \frac{1}{2\mathrm{e}^{2\kappa}\kappa^2}\Big[&2(-1+\mathrm{e}^{\kappa})V_t^2\kappa+2V_t(\theta\kappa(2+\mathrm{e}^{2\kappa}+\mathrm{e}^{\kappa}(-3+\kappa))\\&+\gamma^2(1+\mathrm{e}^{\kappa}(-1+\kappa+\kappa^2\rho^2)))+\theta(2(-1+\mathrm{e}^{\kappa})\theta\kappa(1+\mathrm{e}^{\kappa}(-1+\kappa))\\&+\gamma^2(-1+\mathrm{e}^{2\kappa}(1+4\rho^2)-2\mathrm{e}^{\kappa}(\kappa+2\rho^2+2\kappa\rho^2+\kappa^2\rho^2)))\Big]\end{aligned}$$

$$\mathrm{E}[(p_{t+1}-p_t)V_{t+1}^2|\mathcal{F}_t]=\frac{\gamma(2V_t^2\kappa^2+(-1+\mathrm{e}^{\kappa})\theta(-1+\mathrm{e}^{\kappa}-\kappa)(\gamma^2+2\theta\kappa)+V_t(\gamma^2+2\theta\kappa)(-1-2\kappa+\mathrm{e}^{\kappa}(1+\kappa)))\rho}{\mathrm{e}^{2\kappa}\kappa^2}.$$

Box I.

the expression of V_t and V_{t+1} in the first cross-moment above, we get:

$$\mathrm{E}\left[(p_{t+1}-p_t)\frac{\mathrm{E}(\mathcal{V}_{t+1,t+2}|\mathcal{F}_{t+1})-b}{a}\middle|\mathcal{F}_t\right] = \frac{1}{\mathrm{e}^{\kappa}\kappa}\left[\gamma\left(\frac{\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{F}_t)-b}{a}\kappa+\theta(-1+\mathrm{e}^{\kappa}-\kappa)\right)\rho\right].$$

By the Law of Iterated Expectations,

$$\mathrm{E}\left[(p_{t+1}-p_t)\frac{\mathcal{V}_{t+1,t+2}-b}{a}\middle|\mathcal{F}_t\right] = \frac{1}{\mathrm{e}^{\kappa}\kappa}\left[\gamma\left(\frac{\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{F}_t)-b}{a}\kappa+\theta(-1+\mathrm{e}^{\kappa}-\kappa)\right)\rho\right].$$

Finally, the moment conditions have to be computed by conditioning to the "discrete-time" filtration $\mathcal{G}_t=\{(p_s,\mathcal{V}_{s-1,s}),\ s=t,t-1,\ldots\}\subset\mathcal{F}_t$. Iterating the expectations we get:

$$\mathrm{E}\left[(p_{t+1}-p_t)\frac{\mathcal{V}_{t+1,t+2}-b}{a}\middle|\mathcal{G}_t\right] = \frac{1}{\mathrm{e}^{\kappa}\kappa}\left[\gamma\left(\frac{\mathrm{E}(\mathcal{V}_{t,t+1}|\mathcal{G}_t)-b}{a}\kappa+\theta(-1+\mathrm{e}^{\kappa}-\kappa)\right)\rho\right].$$

An analogous procedure leads to implementable moment conditions derived from the expressions of $\mathrm{E}[(p_{t+1}-p_t)^2V_{t+1}|\mathcal{F}_t]$ and $\mathrm{E}[(p_{t+1}-p_t)V_{t+1}^2|\mathcal{F}_t]$ above.

References

Aït-Sahalia, Y., Kimmel, R., 2007. Maximum likelihood estimation of stochastic volatility models. Journal of Financial Economics 83, 413–452.

Andersen, T.G., Benzoni, L., Lund, J., 2001. Towards an empirical foundation of continuous-time equity return models. Journal of Finance 57, 1239–1284.

Andersen, T.G., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time return models. Journal of Finance 57 (3), 1239–1284.

Andersen, T., Bollerslev, T., 1998. Answering the spektics: yes, standard volatility models do provide accurate forecasts. International Economic Review 53, 219–265.

Andersen, T., Bollerslev, T., Diebold, F.X., Ebens, H., 2001a. The distribution of stock return volatility. Journal of Financial Economics 61, 43–76.

Andersen, T., Bollerslev, T., Diebold, F.X., Labys, P., 2001b. The distribution of exchange rate volatility. Journal of American Statistical Association 96, 42–55.

Andersen, T., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 579–625.

Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. Journal of Finance 94, 277–318.

Barndorff-Nielsen, O.E, Shephard, N., 2001. Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. Journal of the Royal Statistical Society, Series B 63, 167–241.

Barndorff-Nielsen, O.E, Shephard, N., 2004. Power and bipower variation with stochastic jumps. Journal of Financial Econometrics 2, 1–48.

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

Bates, D., 2000. Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics 94, 181–238.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

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

Bollerslev, T., Zhou, H., 2002. Estimating stochastic volatility diffusion using conditional moments of the integrated volatility. Journal of Econometrics 109, 33–65.

Bollerslev, T., Gibson, M., Zhou, H., 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. Jounal of Econometrics 160 (1), 235–245. Finance and Economics Discussion Series. Divisions of Research and Statistics and Monetary Affairs Federal Reserve Board, Washington, DC.

Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. Journal of Finance 55, 839–866.

Chernov, M., Ghysels, E., 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. Journal of Financial Economics 56, 407–458.

Chernov, M., Gallant, R., Ghysels, E., Tauchen, G., 2003. Alternative models of stock price dynamics. Journal of Econometrics 116, 225–257.

Corradi, V., Distaso, W., 2006. Semiparametric comparison of stochastic volatility models using realized measures. Review of Economic Studies 73, 635–667.

Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68 1343–1376.

Eraker, B., Johannes, M.S., Polson, N.G., 2003. The impact of jumps in returns and volatility. Journal of Finance 58, 1269–1300.

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

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

Jiang, G., Tian, Y., 2005. Model-free implied volatility and its information content. Review of Financial Studies 18, 1305–1342.

Lynch, D., Panigirtzoglou, N., 2003. Option implied and realized measures of variance, Working Paper, Monetary Instruments and Markets Division, Bank of England.

Lewis, A.L., 2000. Option Valuation Under Stochastic Volatility. Finance Press.

Lewis, M.-A., 2001. Analytical expressions for the moments of the integrated volatility in affine stochastic volatility models, Working Paper, CIRANO.

Meddahi, N., 2001. An eigenfunction approach for volatility modeling, Working Paper 29-2001, CRDE, Université de Montréal.

Meddahi, N., 2002. Moments of continuous time stochastic volatility models, Working Paper, Université de Montréal.

Meddahi, N., Renault, E., 2004. Temporal aggregation of volatility models. Journal of Econometrics 119, 355–379.

Medvedev, A., Scaillet, O., 2007. Approximation and calibration of short-term implied volatilities under jump-diffusion stochastic volatility. Review of Financial Studies 20 (2), 427–459.

Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study. Journal of Financial Economics 63, 3–50.

Pastorello, S., Renault, E., Touzi, N., 2000. Statistical inference for random variance option pricing. Journal of Business and Economic Statistics 18, 358–367.

Pastorello, S., Patilea, V., Renault, E., 2003. Iterative and recursive estimation in structural nonadaptive models. Journal of Business and Economic Statistics 21–4, 449–482.

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

Romano, M., Touzi, N., 1997. Contingent claim and market completeness in a stochastic volatility model. Mathematical Finance 7 (4), 399–410.

Zhou, H., 2003. Itô conditional moment generator and the estimation of short rate processes. Journal of Financial Econometrics 1, 250–271.