Journal of Banking and Finance 148 (2023) 106745

Contents lists available at ScienceDirect

Journal of Banking and Finance

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

Efficient Quasi-Bayesian Estimation of Affine Option Pricing Models Using Risk-Neutral Cumulants

Riccardo Brignone[a], Luca Gonzato[b,*], Eva Lütkebohmert[a]

[a] Department of Quantitative Finance, Institute for Economic Research, University of Freiburg, Rempartstr. 16, 79098 Freiburg i. Br., Germany
[b] Department of Statistics and Operations Research, University of Vienna, Kolingasse 14-16, Vienna 1090, Austria

ARTICLE INFO

Article history:
Received 8 September 2021
Accepted 5 December 2022
Available online 13 December 2022

JEL classification:
C13
C15
C58
G12

Keywords:
Sequential Monte Carlo
Quasi-Bayesian Estimation
Risk-Neutral Cumulants
Multifactor Affine Models

ABSTRACT

We propose a general, accurate and fast econometric approach for the estimation of affine option pricing models. The algorithm belongs to the class of Laplace-Type Estimation (LTE) techniques and exploits Sequential Monte Carlo (SMC) methods. We employ functions of the risk-neutral cumulants given in closed form to marginalize latent states, and we address parameter estimation by designing a density tempered SMC sampler. We test our algorithm on simulated data by tackling the challenging inference problem of estimating an option pricing model which displays two stochastic volatility factors, allows for co-jumps between price and volatility, and stochastic jump intensity. Furthermore, we consider real data and estimate the model on a large panel of option prices. Numerical studies confirm the accuracy of our estimates and the superiority of the proposed approach compared to its natural benchmark.

1. Introduction

Modern option pricing models can replicate important features of asset price dynamics, such as stochastic volatility, co-jumps, and stochastic jump intensity. These more realistic modeling assumptions, however, come at a cost. The complexity of the models has significantly increased with many unobserved factors, which makes the econometric estimation a complicated task. Therefore, standard methods often fail to produce accurate results, for example, the full frequentist approach suffers from likelihood approximation and difficulties with the optimization procedure (i.e. local maxima). Moreover, from a Bayesian perspective, it is not easy to design a good proposal distribution in a Markov Chain Monte Carlo (MCMC) setting (see e.g. Fulop and Li, 2019). In addition, increased model complexity requires the usage of more informative data rather than only spot returns, which are not sufficient to pin down all the parameters in sophisticated models (see Fulop et al., 2015). For this reason, many authors started using data from derivatives markets to conduct proper inference. This, however, introduces additional estimation complexity due to the highly non-linear relationship between option prices, static parameters, and latent states. Indeed, for optimal statistical inference, the ideal would be to directly consider option panel data. Filtering techniques, which allow to jointly estimate latent states and parameters based on observed option prices, however, are computationally demanding and their extension to multifactor models is not trivial. For example, Hurn et al. (2015) propose a particle filter to estimate the Heston (1993) option pricing model using large cross-sections of option prices, but the implementation is feasible only thanks to the usage of two supercomputers. Du and Luo (2019) show that an unscented Kalman filter approach can be used to simultaneously handle spot and option prices under affine model specifications. Nonetheless, to reduce the computational burden, they are forced to reduce the dataset to only one randomly selected maturity and three option prices per day (weekly data) with a significant loss of information. Furthermore, Dufays et al. (2022) build a particle MCMC algorithm (see Andrieu et al., 2010) for the estimation of several one-factor option pricing models. The likelihood evaluations in this particle filter approach are much faster thanks to the usage of the quantiles of the filtering density, but still computationally very expensive and difficult to extend to multifactor models.

In recent years, estimation approaches have been developed which exploit alternative observation variables that efficiently summarize information from option markets. In this regard, a

* Corresponding author.
E-mail addresses: riccardo.brignone@finance.uni-freiburg.de (R. Brignone), luca.gonzato@univie.ac.at (L. Gonzato), eva.luetkebohmert@finance.uni-freiburg.de (E. Lütkebohmert).

https://doi.org/10.1016/j.jbankfin.2022.106745

very appealing estimation approach proposed by Feunou and Okou (2018) uses risk-neutral cumulants, which can be computed from option prices (Bakshi et al., 2003) and capture the bulk of information embedded in option prices. Under affine model specifications (see Duffie et al., 2003) risk-neutral cumulants have a linear relationship with the state variables. This allows us to compute the quasi-likelihood function through a modified Kalman filter and to obtain parameter estimates through maximum likelihood optimization. Recently, Orlowski (2021) introduced a further refinement based on alternative portfolios of options, which can be seen as functions of risk-neutral cumulants. Both approaches are highly related and have several merits: *i*) they are very general in the sense that they are applicable to all the affine multifactor models; *ii*) they are fast because model implied cumulants can be computed analytically; *iii*) they are extremely accurate because they incorporate information from the option market, which is necessary to pin down parameters governing the jump structure of the model. However, their methods also show some limitations. Firstly, statistical inference is based on likelihood optimization, which is very tricky for models involving many parameters due to the presence of multiple troublesome local maxima. This means that the final output strongly depends on the initial input. Moreover, the Hessian matrix is usually unstable, exacerbating the calculation of the information matrix and, consequently, the standard errors.

In this paper, we present a Quasi-Bayesian (QB) extension of the methods proposed in Feunou and Okou (2018) and Orlowski (2021) which inherits the aforementioned merits and allows surmounting limitations. More specifically, we first integrate out latent states using the modified Kalman filter. We then target the posterior of static parameters by running a density tempered SMC sampler in the spirit of Del Moral et al. (2006) and Duan and Fulop (2015). Since we perform the marginalization step with a quasi-likelihood procedure, our econometric approach leads to a Laplace-type Estimator (LTE), which has been studied in its general formulation by Chernozhukov and Hong (2003) and proved to be useful in many applications (see e.g. Todorov, 2011; Ruge-Murcia, 2012). Compared to the method of Feunou and Okou (2018) our approach can be considered as a more efficient global optimizer and offers several advantages in estimating complex option pricing models. First, the SMC sampler favors likelihood tempering, which means that the algorithm does not get stuck in local maxima. Second, the calculation of any statistics of the posterior (e.g. standard errors, confidence intervals, etc.) is straightforward since we can use results from Chernozhukov and Hong (2003) to give a classical interpretation to our simulation-based output.[1] In particular, from the quasi-posterior distribution we can provide consistent estimates of model parameters (through the quasi-posterior mean) and a sandwich estimator for the asymptotic variance-covariance matrix. Third, given that we perform the marginalization step using a Kalman filter algorithm, we gain considerable computational speed compared with Bayesian alternatives like MCMC methods (see e.g. Eraker, 2004) and SMC^2-type algorithms (see e.g. Chopin et al., 2013; Fulop and Li, 2013; Fulop and Li, 2019).

To illustrate our methodology, we consider an affine dynamic asset pricing model that allows for two stochastic volatility factors, stochastic jump intensity and co-jumps between price and volatility. This model incorporates many popular models as special cases. We want to point out, however, that our proposed method is very general and can be potentially applied to any other affine multifactor model. Before moving to real data applications we perform extensive Monte Carlo studies to check the accuracy, efficiency, and stability of the proposed approach. Compared to the classical quasi-likelihood optimization approach, our new global optimizer allows us to obtain more reliable and very stable parameter estimates.

Then we apply our methodology to real data. We construct historical time series ranging between 05 December 2007 and 04 March 2021 of functions of risk-neutral cumulants of the log-returns with maturity in $1, 2, 3, 6, 9,$ and 12 months. We show that the proposed estimation methodology works well in practice. It produces realistic parameters and filtered factors. A comprehensive econometric analysis of our results highlights the great stability of our methodology, as shown by the acceptance rates obtained during the resample-move steps. The quasi-posterior distributions of static parameters exhibit good convergence, with tight confidence intervals, due to the efficient tempering procedure. Hence, we find that condensing the information contained in option prices into risk-neutral cumulants is an effective way to conduct statistical inference on jump-diffusion option pricing models. It allows drastically reducing the computational effort (compared to competing methods that directly use option prices), without loss of information. We then compare pricing errors obtained through the proposed methodology with those of the plain quasi log-likelihood optimization (used as a benchmark) and find that our approach is more accurate with a sensible reduction of the root mean squared pricing errors across the various strikes and maturities. In summary, numerical results document that the proposed approach outperforms the benchmark in several aspects: *i*) accuracy and ease of computation of the standard errors; *ii*) accuracy in the parameter estimates; *iii*) computing time; *iv*) significant reduction of the option pricing errors, even for a smaller computing time.

Additional empirical studies show that our main model specification outperforms a nested model in terms of log-likelihood and higher-order risk-neutral moments matching, which confirms the importance of introducing more than one volatility factor. Moreover, we highlight the importance of including measures of the risk-neutral kurtosis as additional observed variables for parameter estimation to properly track observed time series of option prices; especially during periods of financial crises. Finally, we show how to estimate the evolution of implied risk premia based on our new approach.

The paper is organized as follows. In Section 2 we present the option pricing model that we consider for numerical experiments throughout the paper, and we detail our econometric approach. Section 3 contains an illustrative example on simulated data. In Section 4 we apply our methodology to a large panel of option prices, and we discuss the economic implications of the estimated parameters and latent factors. In Section 5 we compare the performance of our approach with its natural benchmark, i.e., the Quasi Maximum Likelihood Estimation (QMLE) approach in Feunou and Okou (2018), and we present some additional empirical studies. Section 6 concludes the paper with some final remarks.

2. Econometric Method

In this section we detail our econometric approach. We start by introducing the model setup which is used as an example throughout this paper. Then we review the concept of replicating risk-neutral moments (equivalently, cumulants) using suitable portfolios of options and the computation of cumulants under affine models. Based on this, we introduce a methodology for quasi-likelihood evaluation. Finally, we propose a new QB particle method which serves as an efficient global optimizer for the estimation of model parameters.

[1] We remark that the resulting quasi-posterior is not a valid posterior distribution. Consequently, we are not allowed to use the finite sample posterior as a measure of uncertainty. However, we can follow Chernozhukov and Hong (2003) and derive the asymptotic distribution to conduct proper inference.

2.1. Model

To illustrate our econometric approach we consider an extension of the Double Heston model proposed in Christoffersen et al. (2009), with simultaneous jumps in the variance and price process, and stochastic jump intensity. We refer to this model as Double Jump Double Volatility Stochastic Intensity (DJDVSI) model. Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{Q})$ be a filtered probability space, which supports all the processes we encounter in the sequel. Under the risk-neutral measure $\mathbb{Q}$ the dynamics of the log-returns $X_t = \ln(S_t/S_0)$, where S_t is the asset price at time t, is defined by the following system of stochastic differential equations:

$$dX_t = (r - 0.5(V_{1t} + V_{2t}) - \lambda_t\mu^\star)dt + \sqrt{V_{1t}}dW^x_{1t} + \sqrt{V_{2t}}dW^x_{2t} + J_x dN_t, \quad (1)$$

$$dV_{1t} = k_1(\theta_1 - V_{1t})dt + \sigma_1\sqrt{V_{1t}}dW^v_{1t}, \quad (2)$$

$$dV_{2t} = k_2(\theta_2 - V_{2t})dt + \sigma_2\sqrt{V_{2t}}dW^v_{2t} + J_v dN_t, \quad (3)$$

$$d\lambda_t = k_\lambda(\theta_\lambda - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dW^\lambda_t. \quad (4)$$

where r is the riskfree rate and $\mu^\star = e^{\mu_J + 0.5\sigma_J^2} - 1$ is the convexity adjustment for the jump component. In this specification the log-return process (1) consists of three independent martingales: two diffusive components and a jump component. In particular, V_{1t} and V_{2t} are the two factors driving the instantaneous variance and they evolve as CIR (Cox et al., 1985) processes (with jumps in V_{2t}) with $\mathbb{E}[dW^x_t dW^v_{1t}] = \rho_1 dt$, $\mathbb{E}[dW^x_t dW^v_{2t}] = \rho_2 dt$, which allow to capture the so-called leverage effect. In order to model abrupt changes we consider a compound Poisson process, where $J_x \sim \mathcal{N}(\mu_J, \sigma_J^2)$ dictates the amplitude of a price jump, while $J_v \sim \text{Exp}(\mu_v)$ is an exponentially distributed random variable (J_x and J_v are independent of all the other sources of randomness). The number of jumps is $N_t \sim \text{Poisson}\left(\int_0^t \lambda_s ds\right)$, where λ_t denotes the stochastic jump intensity which is given by (4). It is well documented in the literature (see Fulop et al., 2015; Fulop and Li, 2019; Bardgett et al., 2019 among others) that sudden changes in equity returns and large movements in the variance are likely to occur at the same time. In line with the literature, we therefore choose to generate asset and variance jumps from the same counting process. Finally, the total variance of the stock returns is given by (cfr. Christoffersen et al., 2009, Formula 6)

$$V_t dt = \text{Var}[dS_t/S_t] = (V_{1t} + V_{2t})dt + (\mu_J^2 + \sigma_J^2)\lambda_t dt.$$

As we outline later in this section, in order to implement our methodology we need to derive the risk-neutral Cumulant Generating Function (CGF) of log-returns. The following Proposition 1 shows the Moment Generating Function (MGF) as the solution of a system of Ordinary Differential Equations (ODE). The CGF is then simply obtained by taking the logarithm of the MGF.

Proposition 1. *Given a final date $T > t$ and the time to maturity $\tau = T - t$, the moment generating function of X_T is*

$$\mathbb{E}^{\mathbb{Q}}[e^{uX_T}|\mathcal{F}_t] = \exp\{uX_t + A(u,\tau) + B(u,\tau)V_{1t} + C(u,\tau)V_{2t} + D(u,\tau)\lambda_t\} \quad (5)$$

where the functions A, B, C and D solve the system of ODEs

$$\begin{cases} \frac{\partial A(u,\tau)}{\partial\tau} = ru + k_1\theta_1 B(u,\tau) + k_2\theta_2 C(u,\tau) + k_\lambda\theta_\lambda D(u,\tau), \\ \frac{\partial B(u,\tau)}{\partial\tau} = -\frac{1}{2}(u - u^2) - (k_1 - \rho_1\sigma_1 u)B(u,\tau) + \frac{1}{2}\sigma_1^2 B^2(u,\tau), \\ \frac{\partial C(u,\tau)}{\partial\tau} = -\frac{1}{2}(u - u^2) - (k_2 - \rho_2\sigma_2 u)C(u,\tau) + \frac{1}{2}\sigma_2^2 C^2(u,\tau), \\ \frac{\partial D(u,\tau)}{\partial\tau} = -k_\lambda D(u,\tau) + \frac{1}{2}\sigma_\lambda^2 D^2(u,\tau) + \left(\frac{e^{u\mu_J + u^2\sigma_J^2/2}}{1 - C(u,\tau)\mu_v} - 1\right) \\ \qquad\qquad -(e^{\mu_J + \sigma_J^2/2} - 1)u \end{cases} \quad (6)$$

with initial conditions $A(u,0) = B(u,0) = C(u,0) = D(u,0) = 0$.

Proof. See Appendix A. □

This model specification incorporates as special cases many affine specifications proposed in the option pricing literature (e.g. Heston, 1993, Bates, 1996, Duffie et al., 2000, Christoffersen et al., 2009, Wachter, 2013). However, in our econometric framework it is also straightforward to consider other affine option pricing model specifications such as pure jump variance processes (Barndorff-Nielsen and Shephard, 2001, Bernis et al., 2021), time changed Lévy processes (Carr et al., 2003), stochastic interest rates and other multifactor models (for example Andersen et al., 2015, Gonzato and Sgarra, 2021).

2.2. Risk-Neutral Cumulants

The proposed estimation methodology is based on the idea of condensing the information from option data into measures of risk-neutral variance, skewness, and kurtosis. To this aim, we first show how such measures can be constructed in a model-free manner by using the informative option portfolios proposed by Feunou and Okou (2018) and Orlowski (2021). We then compute the corresponding measures in the present affine model. This provides the basic tools for implementing the modified Kalman filter which we introduce later in this section.

First, we compute the following measures of the variance, skewness, and kurtosis of log-returns from option panels according to (cfr. Orlowski, 2021)

$$\hat{c}^{\text{var}}_{t,\tau} = -2\int_0^{10S_t} p_t(\hat{\sigma}_{IV}(K,\tau))\frac{\partial^2}{\partial K^2}\left(\log\frac{K}{F_t}\right)dK, \quad (7)$$

$$\hat{c}^{\text{skew}}_{t,\tau} = -3\int_0^{10S_t} p_t(\hat{\sigma}_{IV}(K,\tau))\frac{\partial^2}{\partial K^2}\left(\left(\log\frac{K}{F_t}\right)^2 + 2\log\frac{K}{F_t}\right)dK, \quad (8)$$

$$\hat{c}^{\text{kurt}}_{t,\tau} = -4\int_0^{10S_t} p_t(\hat{\sigma}_{IV}(K,\tau))\frac{\partial^2}{\partial K^2}\left(\left(\log\frac{K}{F_t}\right)^3 + 3\left(\left(\log\frac{K}{F_t}\right)^2 + 2\log\frac{K}{F_t}\right)\right)dK, \quad (9)$$

where τ is the option maturity, $F_t = S_t e^{r\tau}$ is the forward index level, and S_t is the time t price of the underlying. $\hat{\sigma}_{IV}(K,\tau)$ and $p_t(\hat{\sigma}_{IV}(K,\tau))$ are the implied volatility and the corresponding price of a put (respectively, call) option with maturity τ and strike $K \le F_t$ ($K > F_t$).

Next, we show how to compute model implied measures of variance, skewness, and kurtosis. We consider model-implied cumulants which can be computed by taking the n-th derivative of the CGF of log-returns. Indeed, if we denote by V_{1t}, V_{2t}, and λ_t the factors on which the distribution of X_t depends in the model (1)–(4), the conditional CGF of log-returns in the affine setting is

$$\psi(u;\tau) = uX_t + A(u,\tau) + B(u,\tau)V_{1t} + C(u,\tau)V_{2t} + D(u,\tau)\lambda_t, \quad (10)$$

where $A(\cdot)$, $B(\cdot)$, $C(\cdot)$, and $D(\cdot)$ solve the system of ODEs given in Proposition 1. It is then possible to retrieve model-implied cumulants by taking the corresponding derivatives evaluated at $u = 0$. We get

$$\begin{aligned} \text{CUM}^{(n)}_{t,\tau} &= \frac{\partial^n\psi(u;\tau)}{\partial u^n}\Big|_{u=0} \\ &= \mathcal{I}X_t + \frac{\partial^n A(u,\tau)}{\partial u^n}\Big|_{u=0} + \frac{\partial^n B(u,\tau)}{\partial u^n}\Big|_{u=0}V_{1t} \\ &\quad + \frac{\partial^n C(u,\tau)}{\partial u^n}\Big|_{u=0}V_{2t} + \frac{\partial^n D(u,\tau)}{\partial u^n}\Big|_{u=0}\lambda_t, \end{aligned} \quad (11)$$

where $\mathcal{I}$ is equal to 1 when $n = 1$ and equal to 0 otherwise. Equation (11) can be solved explicitly in affine models. This fact is very important: model-implied cumulants are obtained through analytical formulas, with the consequence that they are computed in few milliseconds (no numerical methods or special functions are involved) and exactly (no approximations are required at any step). This constitutes a significant advantage over alternative methods which consider option prices as observations and involve time consuming numerical techniques such as numerical solutions to ODE systems or characteristic function inversion (e.g. the unscented Kalman filter proposed in Du and Luo, 2019). In this way, the filtering procedure can be accelerated and is prevented from possible numerical malfunctions.[2]

To avoid illegible long mathematical expressions, we do not report the analytical solution of (11), but instead provide a Matlab® code for the (symbolic) calculation of the quantities $\frac{\partial^n A(u,\tau)}{\partial u^n}\Big|_{u=0}$, $\frac{\partial^n B(u,\tau)}{\partial u^n}\Big|_{u=0}$, $\frac{\partial^n C(u,\tau)}{\partial u^n}\Big|_{u=0}$ and $\frac{\partial^n D(u,\tau)}{\partial u^n}\Big|_{u=0}$ for the model (1)–(4) for $n = \{1, 2, 3\}$ in Appendix B. In Appendix C, we discuss the possibility of extending this approach to the symbolic computation of the required slopes to other affine option pricing models, with particular attention to models with self-exciting jump intensity. Finally, model implied measures of variance, skewness, and kurtosis are given by:

$$\begin{aligned} c^{\text{var}}_{t,\tau} &= -2\,\text{CUM}^{(1)}_{t,\tau} \\ c^{\text{skew}}_{t,\tau} &= -3\big(\text{CUM}^{(2)}_{t,\tau} + 2\,\text{CUM}^{(1)}_{t,\tau}\big), \\ c^{\text{kurt}}_{t,\tau} &= -4\big(\text{CUM}^{(3)}_{t,\tau} + 3\big(\text{CUM}^{(2)}_{t,\tau} + 2\,\text{CUM}^{(1)}_{t,\tau}\big)\big). \end{aligned} \tag{12}$$

2.3. State-Space Formulation

In this section we show how to cast the DJDVSI model into a discrete time state-space model in a similar way as Feunou and Okou (2018). Let us define $c_t = (c^{\text{var}}_{t,\tau}, c^{\text{skew}}_{t,\tau}, c^{\text{kurt}}_{t,\tau})^\top$ and $c^{\text{mkt}}_t = (\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}, \hat{c}^{\text{kurt}}_{t,\tau})^\top$. For a given day t we observe risk-neutral measures of variance, skewness, and kurtosis at J different maturities. Equations (11) and (12) imply that c_t is linearly related to the 3×1 vector of latent factors $F_t = (V_{1t}, V_{2t}, \lambda_t)^\top$. Thus, the measurement equation of our state-space model is

$$c_t = \Gamma_0 + \Gamma_1 F_t + \eta_t, \tag{13}$$

where Γ_0 and Γ_1 are $3J \times 1$ and $3J \times 2$ matrices of coefficients depending on the solution of the derivatives in (11) and the representation in (12). Further, $\eta_t \sim \mathcal{N}(0, \Omega)$ where Ω is a $3J \times 3J$ diagonal covariance matrix.

The transition equation is obtained from a simple Euler approximation with step size Δt of the continuous-time system (2)–(4) and it is defined as

$$F_t = \Phi_0 + \Phi_1 F_{t-1} + \varepsilon_t, \tag{14}$$

where Φ_0 and Φ_1 are 3×1 and 3×3 matrices of coefficients. The transition noise is $\varepsilon_t \sim \mathcal{N}(0, \Sigma(F_{t-1}))$, where ε_t is independent of η_t and $\Sigma(F_{t-1}) = \text{Cov}(F_{t-1})$ is the conditional covariance of F_{t-1}. The complete expressions for Φ_0 and Φ_1 are

$$\Phi_0 = \Delta t \begin{bmatrix} k_1\theta_1 \\ k_2\theta_2 \\ k_\lambda\theta_\lambda \end{bmatrix}, \quad \Phi_1 = I_2 + K_1, \quad K_1 = \Delta t \begin{bmatrix} -k_1 & 0 & 0 \\ 0 & -k_2 & \mu_v \\ 0 & 0 & -k_\lambda \end{bmatrix}, \tag{15}$$

where I_3 is a 3×3 identity matrix. Moreover, the transition noise has the following conditional covariance matrix

$$\begin{aligned} \Sigma(F_{t-1}) &= \text{Cov}(F_{t-1}) \\ &= \Delta t \begin{bmatrix} \sigma_1^2 V_{1,t-1} & 0 & 0 \\ 0 & \sigma_2^2 V_{2,t-1} + 2\mu_v^2 \lambda_{t-1} & 0 \\ 0 & 0 & \sigma_\lambda^2 \lambda_{t-1} \end{bmatrix}. \end{aligned} \tag{16}$$

Note that the standard Kalman filter is not optimal in this setting because the conditional covariance $\text{Cov}(F_{t-1})$ depends on the state itself. Hence, we implement a modified Kalman filter, where the transition noise at time $t-1$ is used as an estimate for the time t iteration. If we define $F_{t|t} = \mathbb{E}[F_t]$ and $P_{t|t} = \text{Cov}[F_t]$, the recursions of the filter are

$$\begin{cases} F_{t+1|t} = \Phi_0 + \Phi_1 F_{t|t} \\ P_{t+1|t} = \Phi_1 P_{t|t} \Phi_1^\top + \Sigma(F_{t|t}) \\ c_{t+1|t} = \Gamma_0 + \Gamma_1 F_{t+1|t} \\ M_{t+1|t} = \Gamma_1 P_{t+1|t} \Gamma_1^\top + \Omega \\ F_{t+1|t+1} = \max\Big(F_{t+1|t} + P_{t+1|t}\Gamma_1^\top M^{-1}_{t+1|t}\big(c^{\text{mkt}}_{t+1} - c_{t+1|t}\big), 0\Big) \\ P_{t+1|t+1} = P_{t+1|t} - P_{t+1|t}\Gamma_1^\top M^{-1}_{t+1|t}\Gamma_1 P_{t+1|t} \end{cases} \tag{17}$$

and the Gaussian quasi log-likelihood is given by

$$-\frac{1}{2}\sum_{t=1}^{T} \ln\big((2\pi)^{3J}\det(M_{t|t-1})\big) + \big(c^{\text{mkt}}_t - c_{t|t-1}\big)^\top M^{-1}_{t|t-1}\big(c^{\text{mkt}}_t - c_{t|t-1}\big). \tag{18}$$

The properties of the aforementioned modified Kalman filter are discussed in Monfort et al. (2017) and the accuracy of the proposed filtering method is investigated in Appendix D.

2.4. A Quasi-Bayesian Particle Method

To conduct inference, we rely on Sequential Monte Carlo (SMC) methods. Denote the vector of model parameters by Θ and all observations and hidden states up to time t by $y_{1:t} = \{c^{\text{mkt}}_t\}_{t=1}^T$ and $x_{1:t} = \{V_{1t}, V_{2t}, \lambda_t\}_{t=1}^T$, respectively. The joint posterior of parameters and hidden states $p(\Theta, x_{1:t} \mid y_{1:t})$ can then be decomposed as follows

$$p(\Theta, x_{1:t} \mid y_{1:t}) = p(x_{1:t} \mid \Theta, y_{1:t}) p(\Theta \mid y_{1:t}), \tag{19}$$

where $p(x_{1:t} \mid \Theta, y_{1:t})$ results from the state filtering problem, and $p(\Theta \mid y_{1:t})$ enables parameter estimation. This suggests a hierarchical structure in which the first task is solved using a modified Kalman filter (outlined in Section 2.3) and the second task is addressed using simulation-based techniques, which we explain below. Given the high information content of option data and the complex structure of state of the art asset pricing models, there is a close link between latent states and fixed parameters. This means that MCMC methods would lead to highly correlated draws and a very slow mixing of the chain (Fulop and Li, 2019). For this reason we target the posterior distribution of static parameters with an SMC sampler as in Del Moral et al. (2006) and Duan and Fulop (2015). Applying Bayes' rule allows us to decompose the posterior distribution of static parameters as follows

$$p(\Theta \mid y_{1:t}) \propto p(y_{1:t} \mid \Theta) p(\Theta), \tag{20}$$

where $p(y_{1:t} \mid \Theta)$ is the likelihood function, and $p(\Theta)$ is the prior density. This decomposition suggests that given a set of parameters (obtained from a prior density) we get an approximation $\hat{p}(y_{1:t} \mid$

[2] We remark that advanced particle filters considering option prices as observed variables are also available. Nevertheless, the likelihood function depends on model implied option prices, meaning that for each likelihood evaluation we have to calculate option prices along three dimensions: for each option (with different strikes and maturities), for each particle, and for each day. This is computationally infeasible due to the necessity of using time consuming numerical techniques to compute option prices.

Θ) of $p(y_{1:t} \mid \Theta)$ and $\hat{p}(x_{1:t} \mid \Theta, y_{1:t})$ of $p(x_{1:t} \mid \Theta, y_{1:t})$ with the modified Kalman filter. We then sample $p(\Theta \mid y_{1:t})$ by using a density tempered SMC sampler. We start from an easy-to-sample distribution, and we then gradually approach the target through a sequence of densities. In particular, we construct a sequence of P densities between the prior and the posterior using a tempering sequence γ_i, $i = 1, \ldots, P$ for $\gamma_1 = 0$ and $\gamma_P = 1$ and

$$\pi_i(\Theta) = \frac{\eta_i(\Theta, y_{1:T})}{\int \eta_i(\Theta, y_{1:T}) d\Theta}, \quad \eta_i(\Theta) = \hat{p}_i(y_{1:T} \mid \Theta)^{\gamma_i} p(\Theta), \tag{21}$$

where $\pi_i(\Theta)$ represents the intermediate posterior at iteration i, and $\hat{p}_i(y_{1:T} \mid \Theta)^{\gamma_i}$ is the (tempered) likelihood estimated by a modified Kalman filter. To move from $\pi_i(\Theta)$ to $\pi_{i+1}(\Theta)$, we reweigh each particle by $\hat{p}_i(y_{1:T} \mid \Theta^{(n)})^{\gamma_{i+1}-\gamma_i}$ for $n = 1, \ldots, N$. Here N is the number of parameter particles and to each particle corresponds a different vector of parameters $\Theta^{(n)}$. The tempering coefficients are chosen to ensure sufficient particle diversity. This can be done by a grid search, where the Effective Sample Size (ESS)[3] is evaluated at the grid points of γ, and the one with the ESS closest to a fixed constant is selected. At this stage, it is essential to avoid particle impoverishment. First, we resample particles proportional to their weights to obtain an equally weighted set. Using a Markov kernel, we then shift the particles to enrich the support without changing the distribution of the particles.

Since using an approximated filter yields a marginal quasi-likelihood approximation, our approach departs from the pseudo-marginal framework. Instead, we assume a QB setting, also known as LTE. We emphasize again that we may not use the estimated quasi-posterior distribution to conduct inference. However, we can easily derive the asymptotic distribution for this purpose. In particular, to give our simulation results a classical interpretation, we compute the quasi-posterior mean as a consistent estimate of true parameters and the sandwich covariance matrix as a measure of uncertainty. The asymptotic distribution of the estimator is then given by

$$\tilde{\Theta} \sim \mathcal{N}\big(\hat{\Theta}, J(\hat{\Theta})^{-1} \Omega(\hat{\Theta}) J(\hat{\Theta})^{-1}\big), \tag{22}$$

where $\hat{\Theta}$ is the quasi-posterior mean. $J(\hat{\Theta})^{-1}$ is the inverse Hessian and can be estimated from the covariance of the quasi-posterior distribution, and $\Omega(\hat{\Theta})$ denotes the covariance of the scores which can be estimated by running the modified Kalman filter at the quasi-posterior mean (see Chernozhukov and Hong, 2003 for more details).

In the next section, we show that our QB-SMC sampler performs very well on a challenging problem from the option pricing literature.

3. Illustrative Example on Simulated Data

To illustrate the performances of our estimation approach, we proceed with a test on simulated data. For this experiment, we fix model parameters close to those estimated in the literature on similar models. We discretize the SDEs in (2)–(4) using the Euler scheme (see Section 2.3) with daily time steps. Based on this, we simulate 30 times the latent factors for a total of 3780 observations, mimicking $15 = 3780/252$ years of daily data. From the paths of V_1, V_2 and λ, we recover the (simulated) time series of the second and third risk-neutral cumulants from (11) for the following six maturities: $\tau \in \{1, 2, 3, 6, 9, 12\}$ months. Next, following the literature (e.g. Li and Zinna, 2018; Fulop and Li, 2019; Du and Luo, 2019), we generate the observation error from a Gaussian random variable with zero mean and diagonal covariance matrix $\Omega = \sigma_e^2 I_{3J}$, where σ_e is the standard deviation of the pricing error. At the end of this procedure, we have 30 simulated time series (for a 15-year period with daily observations) of risk-neutral cumulants for six different maturities.

For each dataset, we run our QB-SMC sampler and the QMLE method based on the numerical maximization in (18) for benchmark comparison. We run our density tempered SMC sampler using $N = 2000$ particles. Given the complex form of the likelihood, we simply give Normal priors to all parameters, but we consider truncated Normal priors to respect the domain of some parameters, see Table 1. The hyperparameters of the prior distributions are chosen to be consistent with those estimated in the recent literature (see e.g. Fulop and Li, 2019; Du and Luo, 2019).

In the resample-move step, we use a Gaussian random walk proposal and move the parameters as one block using a Metropolis-Hastings (MH) kernel. More precisely, after resampling, we rejuvenate the particles' population by proposing new parameters from a Gaussian random walk proposal. To compute the MH acceptance rate, we then run the modified Kalman filter for each new set of parameters using all observations. The mean and covariance of the proposal distribution are fitted to the actual particles' population. We further increase the efficiency of the MCMC step by adjusting the scale of the proposal to obtain an acceptance rate between 0.2 and 0.4. The resample-move step is triggered whenever the ESS is below $N/2$, and we continue to move as long as the number of unique particles is below a threshold $N/2$.

To obtain reliable estimates for the QMLE method, we generate $M = 30\,000$ random model parameters from the priors in Table 1. We compute the quasi log-likelihoods in (18) for each of them, and use the 5 sets of parameter with the highest quasi log-likelihood as initial points for the numerical maximization. Finally, the parameter set with the highest quasi log-likelihood is used as final estimate of the QMLE method (this procedure is similar to that used in Li and Yin, 2014 and Yeap et al., 2018).[4] The number of random model parameter sets and the number of initial points for the optimization are chosen so that the average execution time of QMLE and QB-SMC methods is similar (see below), which allows a fair comparison in terms of accuracy. More specifically, the average computational time for one run of the QB-SMC algorithm in this experiment is 24.78 minutes, while one run of the QMLE takes 27 minutes. All computations were performed on a desktop HP EliteDesk 800 G5 PC.[5]

We collect estimates from both methods and evaluate the accuracy. Let us denote by $\hat{\Theta}_i$ the posterior mean of parameters for the i-th dataset and with $\Theta^\star$ the vector of the true parameters. Mean and Root Mean Squared Error (RMSE) are computed as

$$\text{Mean} = \frac{1}{30}\sum_{i=1}^{30} \hat{\Theta}_i, \quad \text{RMSE} = \sqrt{\frac{1}{30}\sum_{i=1}^{30} (\hat{\Theta}_i - \Theta^\star)^2}.$$

Results are given in Table 2 along with the true parameter values. The results document that for our proposed method, the average of the parameter estimates is close to the true value for all the parameters. Hence, our QB-SMC method outperforms the QMLE approach overall with superior performances in terms of bias (i.e. the difference between mean and true values) and RMSE for all parameters (except for σ_1, where the differences are small anyway). This is particularly evident for those parameters which are notoriously difficult to estimate, such as ρ_1, ρ_2, k_λ, θ_λ, σ_λ, μ_J and σ_J (see e.g. Eraker, 2004; Andersen et al., 2015). The QMLE method produces much higher RMSEs, which is due to the presence of multiple local maxima. In fact, the gradient-based opti-

[3] This quantity is defined by $\text{ESS} = (\sum_{n=1}^N s_i^{(n)})^2 / \sum_{n=1}^N (s_i^{(n)})^2$, where $s_i^{(n)}$ is the weight attached to the n-th particle at iteration i.

[4] Maximization is implemented using the built-in Matlab® function `fmincon`.

[5] In the QB-SMC algorithm, the modified Kalman filter runs in parallel on 8 CPU cores.

Table 1
Priors specification.

Θ	Dist	Support	(μ_0, σ_0)	Θ	Dist	Support	(μ_0, σ_0)
k_1	Tr. Normal	$(0, \infty)$	(6.00, 3.00)	θ_λ	Tr. Normal	$(0, \infty)$	(1.50, 2.00)
θ_1	Tr. Normal	$(0, \infty)$	(0.01, 0.10)	σ_λ	Tr. Normal	$(0, \infty)$	(2.00, 2.00)
σ_1	Tr. Normal	$(0, \infty)$	(0.50, 1.00)	μ_ν	Tr. Normal	$(0, \infty)$	(0.025, 0.05)
ρ_1	Tr. Normal	$(-1, 1)$	(-0.50, 0.50)	μ_J	Normal	$(-\infty, \infty)$	(-0.02, 0.05)
k_2	Tr. Normal	$(0, \infty)$	(6.00, 3.00)	σ_J	Tr. Normal	$(0, \infty)$	(0.03, 0.05)
θ_2	Tr. Normal	$(0, \infty)$	(0.01, 0.10)	σ_{e1}	Tr. Normal	$(0, \infty)$	(0.10, 0.10)
σ_2	Tr. Normal	$(0, \infty)$	(0.50, 1.00)	σ_{e2}	Tr. Normal	$(0, \infty)$	(0.10, 0.10)
ρ_2	Tr. Normal	$(-1, 1)$	(-0.50, 0.50)	σ_{e3}	Tr. Normal	$(0, \infty)$	(0.10, 0.10)
k_λ	Tr. Normal	$(0, \infty)$	(6.00, 2.00)				

Table 2
Mean and RMSE of the parameter estimates from 30 Monte Carlo simulations for the QB-SMC and QMLE. Values in green indicate the best performance, values in red denote the worst performance.

		QB-SMC		QMLE	
Θ	True	Mean	RMSE	Mean	RMSE
k_1	10.50	10.3242	0.2709	10.2673	2.5466
θ_1	0.03	0.0295	0.0006	0.0450	0.0313
σ_1	0.40	0.3841	0.0185	0.3972	0.1373
ρ_1	-0.50	-0.5317	0.0354	-0.5503	0.2232
k_2	1.10	1.1023	0.0046	0.8289	0.4419
θ_2	0.02	0.0219	0.0021	0.0472	0.0425
σ_2	0.15	0.1497	0.0023	0.1594	0.0539
ρ_2	-0.80	-0.8067	0.0146	-0.7012	0.1880
μ_ν	0.05	0.0502	0.0002	0.0649	0.0247
k_λ	0.65	0.6517	0.0071	0.9807	0.9229
θ_λ	5.75	5.7220	0.0439	5.0754	1.8179
σ_λ	2.66	2.6538	0.0085	1.7609	1.3153
μ_J	-0.04	-0.0400	0.0001	-0.0555	0.0268
σ_J	0.05	0.0501	0.0001	0.0407	0.0188

mization strongly depends on the initial point supplied. This translates into a larger variance in the parameter estimates, which is reflected in the higher RMSEs. To reduce RMSEs, the number of eligible starting points would have to be increased, but this can only be achieved at the cost of increased computational effort. In contrast, the QB-SMC method does not suffer from the presence of local maxima, and the final parameter estimates are close to each other across the 30 simulations (as indicated by the smaller RMSEs). As a consequence, this method appears to be more robust and suitable for real data applications.

Finally, to better illustrate the superior performances of the proposed method, we perform the following experiment. We generate a single simulated dataset as described above and compute the true (quasi) log-likelihood as in (18). Then we estimate the model on this single dataset 10 times for $N \in \{500, 1000, 1500, 2000\}$ for QB-SMC and $M \in \{10\,000, 20\,000, 30\,000, 40\,000\}$ for QMLE. For the latter, we consider the 5 parameter sets with the highest log-likelihood as initial points for the numerical optimization. We then take the average (over the 10 repetitions) of the time required to perform the estimation and of the quasi log-likelihood in agreement with the estimated parameters. The difference between the true (quasi) log-likelihood and the average estimated log-likelihood measures the accuracy of the methodology. The results are reported in Table 3 and are striking. The proposed QB-SMC approach outperforms the QMLE method in terms of both running time and accuracy.

4. Empirical Estimation

In this section, we apply the proposed methodology to estimate the model in (1)–(4) on real data. We start by describing how the historical time series of observed risk-neutral measures of variance, skewness, and kurtosis are built from option quotes. We then present estimation results and performances of the method in terms of parameter convergence, statistical efficiency, and latent factor filtering. A discussion on estimated parameters is deferred to the next section, where we compare our results to a nested model and to a benchmark estimation method.

Table 3
Run-time accuracy profiles of QB-SMC and QMLE methods on simulated data. Time is expressed in minutes, the log-likelihood is computed given the estimated parameters. Both results correspond to the average across 10 repetitions on the same simulated dataset. The true quasi log-likelihood is 3.8840E+05.

QB-SMC			QMLE		
N	time	log-like	M	time	log-like
500	7.95	3.8827E+05	10 000	11.23	3.4304E+05
1000	16.09	3.8830E+05	20 000	19.97	3.4803E+05
1500	19.38	3.8836E+05	30 000	27.01	3.5355E+05
2000	24.78	3.8838E+05	40 000	35.37	3.5411E+05

4.1. Data

Our empirical analysis is based on the construction of model-free risk-neutral measures of variance, skewness, and kurtosis for various maturities. OptionMetrics® provides historical time series of implied volatility surfaces for the S&P 500. We consider daily observations from 05 December 2007 to 04 March 2021 for a total of 3457 days and compute $\hat{c}^{\text{var}}_{t,\tau}$, $\hat{c}^{\text{skew}}_{t,\tau}$ and $\hat{c}^{\text{kurt}}_{t,\tau}$ as in (7)–(9).[6] Our dataset contains a similar number of observations as the one in Feunou and Okou (2018) but more recent data. This allows us to consider not only the 2008/2009 global financial crisis, but also other important periods of financial turmoils like the European debt crisis of 2009–2013 and the COVID pandemic of 2020–2021. Considering the following fixed maturities $\tau \in \{1, 2, 3, 6, 9, 12\}$ (expressed in months), we end up with 3457 observations of $\hat{c}^{\text{var}}_{t,\tau}$, $\hat{c}^{\text{skew}}_{t,\tau}$, $\hat{c}^{\text{kurt}}_{t,\tau}$ on six different maturities, i.e., $3457 \times 3 \times 6 = 62226$ total observations. Historical time series of the estimated risk-neutral cumulants are reported in Figure 1. We report summary statistics in Table 4, where we display for each maturity the mean, standard deviation, and the first-lagged autocorrelation of the observed measures of variance, skewness, and kurtosis. From Table 4 we can see that the mean (in absolute value) and the standard deviatios of $\hat{c}^{\text{var}}_{t,\tau}$, $\hat{c}^{\text{skew}}_{t,\tau}$, and $\hat{c}^{\text{kurt}}_{t,\tau}$ increase with maturity. We also notice that $\hat{c}^{\text{var}}_{t,\tau}$, $\hat{c}^{\text{skew}}_{t,\tau}$, and $\hat{c}^{\text{kurt}}_{t,\tau}$ are highly persistent.

A question that might arise is how much information is lost by using the second, third, and fourth risk-neutral cumulants instead of the full option panel. To answer this question, we run the following regression:

$$IV^{\text{mkt}}_{t,\tau,K} = \alpha + \beta_1 \hat{c}^{\text{var}}_{t,\tau} + \beta_2 \hat{c}^{\text{skew}}_{t,\tau} + \beta_3 \hat{c}^{\text{kurt}}_{t,\tau} + \varepsilon_{t,\tau}, \tag{23}$$

[6] For more details on the implementation of these formulas we refer to Orlowski (2021, Section 4.3).

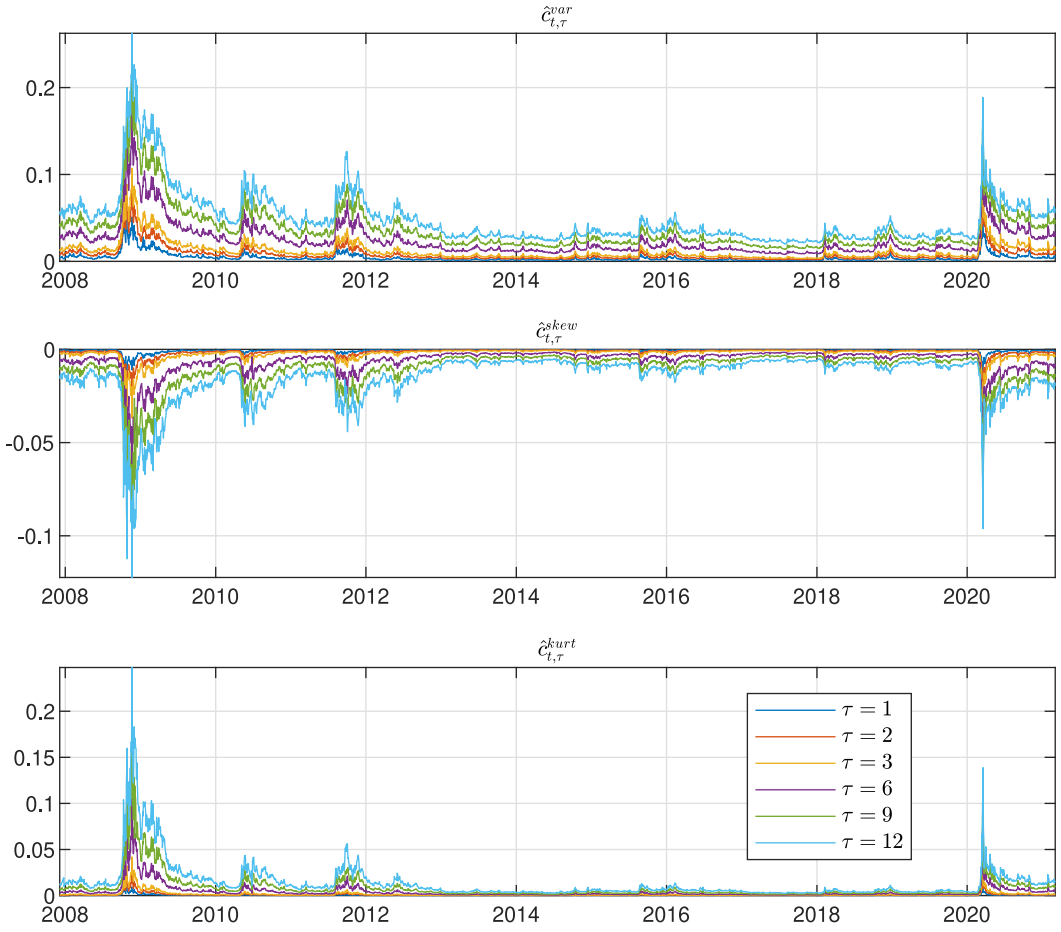

Fig. 1. Historical time series of observed measures of variance, skewness, and kurtosis for 6 different maturities $\tau \in \{1, 2, 3, 6, 9, 12\}$ months for the time period 05 December 2007 to 04 March 2021.

Table 4
Summary statistics for the observed measures of variance, skewness, and kurtosis for 6 different maturities $\tau \in \{1, 2, 3, 6, 9, 12\}$ months for the period between 05 December 2007 and 04 March 2021.

	$\hat{c}^{\text{var}}_{t,1}$	$\hat{c}^{\text{var}}_{t,2}$	$\hat{c}^{\text{var}}_{t,3}$	$\hat{c}^{\text{var}}_{t,6}$	$\hat{c}^{\text{var}}_{t,9}$	$\hat{c}^{\text{var}}_{t,12}$
Mean (in %)	0.4215	0.8012	1.2080	2.4453	3.6696	4.8967
Std. dev. (in %)	0.5342	0.8550	1.1444	1.8767	2.4891	3.0591
AR(1)	0.9606	0.9754	0.9800	0.9867	0.9893	0.9906
	$\hat{c}^{\text{skew}}_{t,1}$	$\hat{c}^{\text{skew}}_{t,2}$	$\hat{c}^{\text{skew}}_{t,3}$	$\hat{c}^{\text{skew}}_{t,6}$	$\hat{c}^{\text{skew}}_{t,9}$	$\hat{c}^{\text{skew}}_{t,12}$
Mean (in %)	-0.0581	-0.1359	-0.2357	-0.6082	-1.0350	-1.5076
Std. dev. (in %)	0.1131	0.2162	0.3277	0.6496	0.9476	1.2598
AR(1)	0.9235	0.9586	0.9575	0.9670	0.9653	0.9782
	$\hat{c}^{\text{kurt}}_{t,1}$	$\hat{c}^{\text{kurt}}_{t,2}$	$\hat{c}^{\text{kurt}}_{t,3}$	$\hat{c}^{\text{kurt}}_{t,6}$	$\hat{c}^{\text{kurt}}_{t,9}$	$\hat{c}^{\text{kurt}}_{t,12}$
Mean (in %)	0.0276	0.0700	0.1345	0.4272	0.8384	1.3782
Std. dev. (in %)	0.0779	0.1780	0.3064	0.7778	1.3348	2.0315
AR(1)	0.9052	0.9321	0.9397	0.9633	0.9676	0.9792

where $IV^{\text{mkt}}_{t,\tau,K}$ denotes the market implied volatility for a maturity τ and strike price K, while ε is the error. The adjusted R^2 from this linear regressions are reported in Table 5 for different levels of moneyness and various maturities. Adjusted R^2 are extremely high, confirming that the second, third, and fourth risk-neutral cumulants explain (on average) about 98% of the variation contained in the implied volatility. Therefore, condensing the information contained in option panels into risk-neutral cumulants represents a viable method for conducting statistical inference. Compared to directly considering option prices, the loss of information is small, but as we illustrate later, such an approach leads to a significant speedup in the calculations.

4.2. Econometric Performance

Next, we move to the empirical results obtained by running the QB-SMC sampler algorithm with $N = 2000$ particles. We take the Gaussian random walk proposal and perform the resample-move step whenever the ESS falls below the preselected threshold of

Table 5
Adjusted R^2 from the regression in (23) for various levels of moneyness and maturities.

$\tau \setminus K$	80	85	90	95	100	105	110	115	120
1	0.8770	0.9293	0.9654	0.9805	0.9849	0.9850	0.9480	0.8560	0.8449
2	0.9789	0.9876	0.9912	0.9900	0.9901	0.9919	0.9868	0.9633	0.9163
3	0.9873	0.9918	0.9931	0.9915	0.9917	0.9913	0.9897	0.9750	0.9395
6	0.9928	0.9941	0.9952	0.9942	0.9942	0.9912	0.9912	0.9883	0.9811
9	0.9939	0.9944	0.9953	0.9944	0.9946	0.9917	0.9910	0.9890	0.9850
12	0.9875	0.9948	0.9955	0.9938	0.9949	0.9920	0.9906	0.9890	0.9864

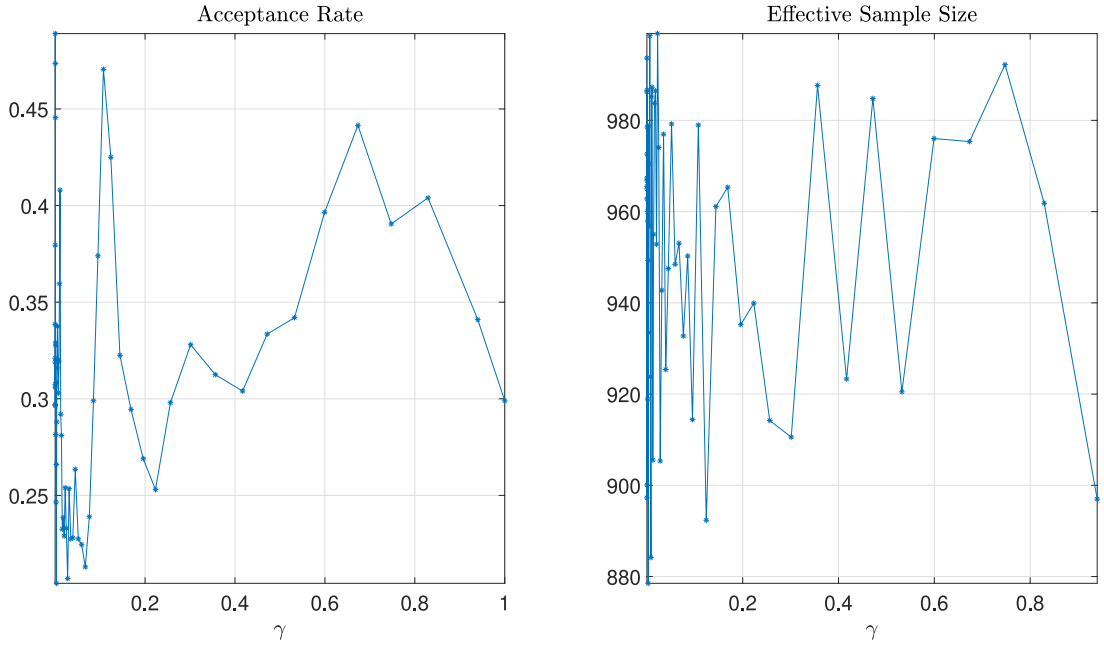

Fig. 2. Acceptance rates (left panel) and ESS (right panel) with respect to γ.

$N/2$, and then move until the number of unique particles is greater than $N/2$. The sampler is initialized with the priors in Table 1.

From an econometric point of view, two important statistics related to the efficiency of our method are the ESS and the acceptance rate. In Figure 2, we present the acceptance rate (left panel) and the ESS (right panel). A few comments are in order. First, the acceptance rate optimally fluctuates between 0.2 and 0.45 ("Goldilocks Principle", Rosenthal, 2011). This means that our algorithm is able to efficiently incorporate the increasing degree of information content revealed by the data. The ESS is required to stay close to a constant, which is usually set equal to $N/2$, where $N = 2000$ is the number of parameter particles. As expected, the ESS fluctuates around half of the number of parameter particles. Second, in Figure 3 we present the tempering procedure, which leads to the progressive shrinkage of confidence bands. More precisely, we see that at the beginning, when the tempering coefficient is equal to 0, we only have the prior information, and the prior distributions have quite large standard errors. However, as the algorithm proceeds, the information contained in the data is reflected in the estimates. This can be seen in the shrinkage of the $(5, 95)$% confidence intervals (red lines). Results are reported for a few selected parameters, and similar plots for the remaining parameters can be found in Figure E.11 in the Appendix. Another important aspect related to the efficiency of the proposed methodology is the impact of the prior distribution on the final estimation output. To address this point, following Chopin et al. (2013, Figure 1), Fulop et al. (2015, Figure A1), Brix et al. (2018, Figure 10), we report the prior and quasi-posterior distributions in Figure 4. Again, results are given only for a subset of the parameters. For the remaining, we refer to Figure E.12 in the Appendix. We observe that in our QB-SMC approach, the priors are flat in the regions of high quasi-posterior density. This means that the prior specification does not impact the final result. Nevertheless, the dispersions of the posterior distributions are very small.

In Figure 5, we plot the filtered variance (central panel) and jump intensity (bottom panel). In periods when prices fall, we observe a rise in variance which is in line with the well known leverage effect. This is also confirmed by our estimates for $\rho_1 = -0.9223$ and $\rho_2 = -0.7834$ (see Section 5).

The grey bars in the bottom panel of Figure 5 represent the jump times in S&P 500 dynamics. Those jumps are detected using the iterated re-weighted least squares technique developed in Callegaro et al. (2017). The algorithm detects 216 jumps in the S&P 500 price in our sample. We note that the filtered jump intensity λ_t, displayed in Figure 5, is high during turbulent periods with many jumps in close succession (e.g. during 2008–2009 and 2020–2021) and reverts to its long run mean in quiet periods. Moreover, we find that the jump intensity is close to 0 in the 2017–2018 period, which is consistent with the fact that the iterated reweighted least squares algorithm does not detect any jumps in this period.

We conclude this section with a discussion of the variance factors and jump parameters. First, we obtain volatility factors with very different statistical properties. V_2 is highly persistent, while V_1 displays low persistent dynamics, as indicated by the estimated speed of mean reversion parameters $k_1 = 11.9764$ and $k_2 = 0.9672$. This result is consistent with findings in Christoffersen et al. (2009), who observe that one factor may exhibit fast mean reversion (describing short-run variance), while the second factor may exhibit relatively low speed of mean reversion to describe long-run variance. Second, the jump parameters are

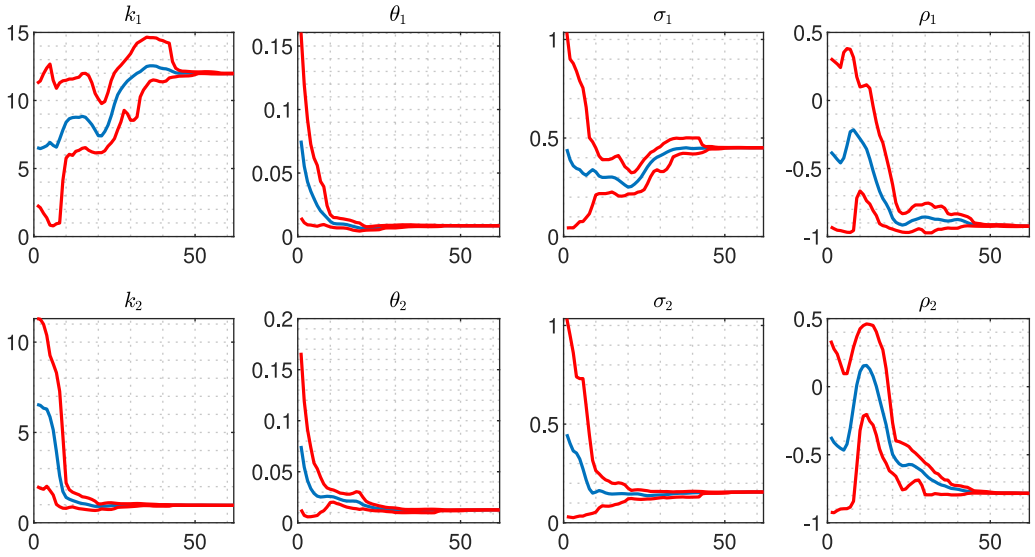

Fig. 3. Bridging the priors and the quasi-posteriors: mean (blue line) and (5, 95)% quantiles (red lines) of some selected parameters at each tempering step.

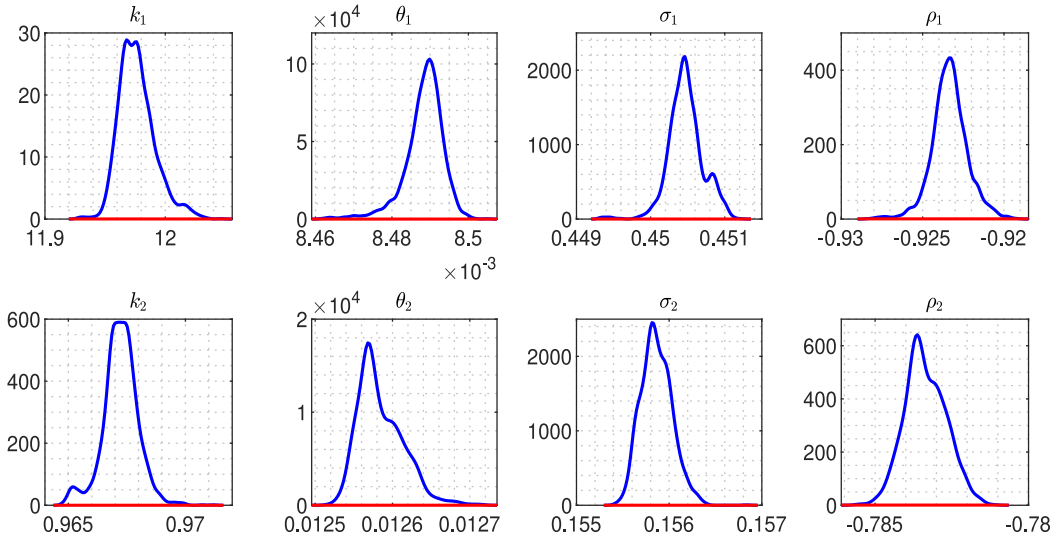

Fig. 4. Full sample quasi-posterior distributions (blue line), obtained with a kernel density estimator, and prior distributions (red line) of some selected parameters.

quite large compared to what is typically found in the literature. As we illustrate later, this is not due to the estimation methodology, as the benchmark QMLE approach yields similar results. An explanation for such large (in absolute value) jump parameters is provided in Section 5.3.

5. Parameter Estimation and Benchmark Comparison

In this section, we conduct several numerical experiments on real data: *i*) a benchmark comparison with the QMLE approach in terms of econometric fit, run-time speed, and option pricing errors; *ii*) a comparison with a nested model in terms of cumulants matching; *iii*) a study on the sensitivity of the parameter estimates (and related option pricing errors) with respect to the choice of the dataset; *iv*) a study of how to incorporate risk premia in our estimation approach. More specifically, for point *iii*) we consider three different datasets: $\mathcal{D}_1$, $\mathcal{D}_2$ and $\mathcal{D}_3$. The first one considers as observation only $\{\hat{c}^{\text{var}}_{t,\tau}\}$, the second one considers $\{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}\}$, while the third one considers $\{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}, \hat{c}^{\text{kurt}}_{t,\tau}\}$. The parameter estimates, maximized log-likelihoods, and related computational times are reported in Table 6. The results are discussed below.

5.1. Benchmark Comparison

We compare the performances of the proposed QB-SMC approach with its natural benchmark, i.e., the QMLE approach based on the numerical maximization of (18). Note that the main drawback of the numerical implementation of the QMLE method is the choice of the starting points fed to the numerical optimizer. In fact, the objective function has multiple local maxima, so the choice of the starting points is crucial. This problem is exacerbated by the large number of parameters to be estimated. The procedure we adopt to implement QMLE is as follows (compare with Li and Yin, 2014; Yeap et al., 2018): *i*) we use the priors in Table 1 and generate 10^5 parameter sets; *ii*) we evaluate the quasi log-likelihood for each parameter set as in (18); *iii*) we take the 10 parameter sets with the highest quasi log-likelihood and run 10 different optimizations using these parameter sets as starting points; *iv*) we take the parameters with the highest quasi log-likelihood as the final estimate. The results are reported in the fourth and fifth columns of Table 6. Several comments are in order. First, computing standard errors for QMLE is very challenging. We have considered several numerical optimizers that provide different (nu-

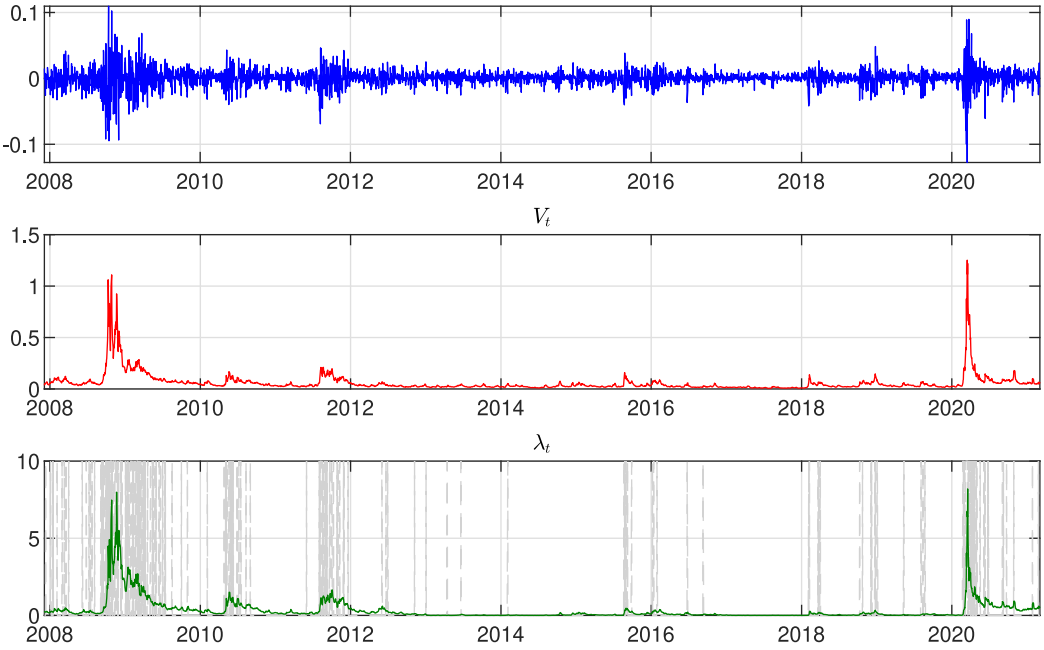

Fig. 5. Historical log-returns (top panel) of S&P 500 index, filtered variance (central panel), filtered jump intensity (bottom panel) along with jump times (indicated with light grey bars) in the S&P 500 price historical time series.

merical) estimates of the Hessian matrix. However, its numerical (pseudo) inversion could not be determined, which precludes the computation of standard errors. This is a first point in favour of our approach, for which we have reliable estimates of the standard error (see 22). These are reported in parenthesis in Table 6. Second, our QB-SMC offers a better trade-off between accuracy and computational effort. For the above settings, QB-SMC is more accurate (as indicated by the highest log-likelihood for both models) and also faster in terms of execution time. This fact gives confidence that the suggested methodology can be interpreted as a more efficient global optimizer that provides a considerably higher log-likelihood at lower computational cost.

Table 6
Parameters estimated through QB-SMC and QMLE approaches for DJDVSI and DJSVSI models and for different datasets $\mathcal{D}_1 = \{\hat{c}^{\text{var}}_{t,\tau}\}$, $\mathcal{D}_2 = \{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}\}$, $\mathcal{D}_3 = \{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}, \hat{c}^{\text{kurt}}_{t,\tau}\}$. Standard errors in parenthesis. Running time is expressed in minutes. Following Feunou and Okou (2018) we fix the pricing errors for the QMLE approach and set them equal to those obtained through the QB-SMC approach.

	DJDVSI				DJSVSI	
	QB-SMC			QMLE	QB-SMC	QMLE
$\Theta \setminus \mathcal{D}$	$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{D}_3$	$\mathcal{D}_3$	$\mathcal{D}_3$	$\mathcal{D}_3$
k_1	13.9812 (0.2277)	12.6197 (0.4218)	11.9764 (0.0139)	13.5657	1.1311 (0.0285)	1.2877
θ_1	0.0158 (0.0008)	0.0133 (0.0008)	0.0085 (0.0000)	0.0065	0.0282 (0.0001)	0.0134
σ_1	0.6622 (0.0117)	0.5787 (0.0131)	0.4505 (0.0001)	0.2509	0.2518 (0.0004)	0.1859
ρ_1	−0.9359 (0.1280)	−0.9858 (0.0059)	−0.9233 (0.0002)	-0.6219	−0.9748 (0.0079)	-0.984
k_2	1.5739 (0.1531)	1.8667 (0.0236)	0.9672 (0.0003)	1.5561	–	–
θ_2	0.0114 (0.0018)	0.0023 (0.0000)	0.0126 (0.0000)	0.0092	–	–
σ_2	0.1871 (0.0019)	0.0920 (0.0007)	0.1559 (0.0001)	0.1691	-	-
ρ_2	−0.9587 (0.0736)	−0.9206 (0.0159)	−0.7834 (0.0006)	-0.5929	–	–
μ_v	0.0566 (0.0122)	0.1017 (0.0041)	0.1404 (0.0000)	0.1686	0.1005 (0.0034)	0.1237
k_λ	2.5593 (0.0599)	3.8644 (0.0511)	3.4321 (0.0045)	3.2843	5.8298 (0.0637)	5.1551
θ_λ	0.3257 (0.0184)	0.4759 (0.0058)	0.2244 (0.0001)	0.2426	0.2649 (0.0030)	0.3218
σ_λ	1.2822 (0.0293)	1.9159 (0.0202)	1.2395 (0.0004)	1.2602	1.7533 (0.0067)	1.7237
μ_J	−0.0119 (0.0455)	−0.1095 (0.0022)	−0.1263 (0.0000)	-0.1463	−0.0806 (0.0001)	-0.0955
σ_J	0.002 (0.0099)	0.0722 (0.0015)	0.2145 (0.0001)	0.2143	0.1829 (0.0025)	0.1986
$\sigma_e^{(1)}$	0.0026 (0.0006)	0.0075 (0.0006)	0.0130 (0.0000)	–	0.0687 (0.0012)	–
$\sigma_e^{(2)}$	–	0.0509 (0.0165)	0.0496 (0.0004)	–	0.0552 (0.0060)	–
$\sigma_e^{(3)}$	–	–	0.1398 (0.0003)	–	0.0851 (0.0112)	–
Log-Like	1.12E+05	2.16E+05	3.21E+05	3.15E+05	3.10E+05	3.07E+05
time	21.51	40.28	28.76	130.79	29.85	79.87

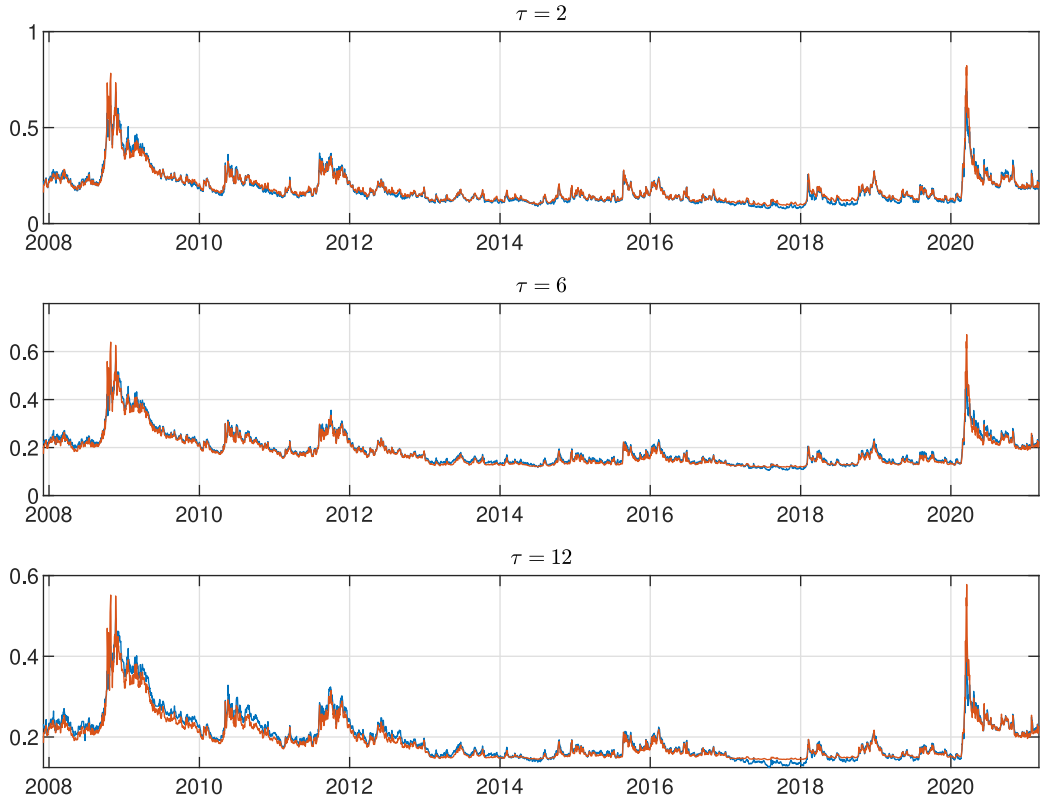

Fig. 6. Implied volatilities comparison. Blue lines represent the time series of at-the-money market implied volatilities, while orange lines represents the model implied volatilities for $\tau = \{2, 6, 12\}$ months. Model implied volatilities are computed from the DJDVSI model using the QB-SMC approach, the parameter estimates are given in the fourth column of Table 6.

Finally, we show how the estimated parameters in Table 6 perform in pricing options. More specifically, we take the model parameters and re-filter the latent states. Using the model parameters and the filtered latent states, we then compute model implied volatility surfaces for each day in our sample. To do this, we first compute the option price and invert the Black-Scholes formula. The option price is computed using the Fourier-Cosine series expansion method (COS) of Fang and Oosterlee (2008). First, we consider the estimates of our QB-SMC approach for the DJDVSI model and show in Figure 6 the time series of market and model implied volatilities (for three different maturities). The two time series are very close, which means that parameters estimated by our proposed approach are highly realistic and can correctly replicate the time series of market implied volatility surfaces. In Table 7, we report the Root Mean Squared Pricing Errors (RMSPEs) for both estimation methods. In the bottom panel of this table, we report the ratio between the RMSPEs of QB-SMC and QMLE approaches. A value less than 1 (highlighted in green) indicates that the QB-SMC performs better than the benchmark, and vice versa when the value is greater than 1. We find that the parameters estimated with QB-SMC perform much better than the benchmark in valuing options. This is quite evident across all maturities and strikes. QB-SMC always outperforms the benchmark except in ten (out of 36) cases, which are highlighted in red in Table 7. The average of the ratios is 0.9190, indicating a generally superior performance of QB-SMC.

To sum up, we have shown that the proposed approach provides parameter estimates that better explain the observed market data, both in terms of log-likelihood and root mean squared pricing errors.

5.2. Comparison with a Nested Model

Next, we compare the performances of the DJDVSI model with a nested specification. More specifically, we consider a two-factor model that results from (1)–(4) with one of the two variance factors turned off. We refer to this model as Double Jump Stochastic Volatility Stochastic Intensity (DJSVSI) model because it features jumps both in the variance and spot returns processes and allows for stochastic jump intensity. First, the QB-SMC also performs very well in estimating this nested model. Our approach yields higher log-likelihood compared to the QMLE approach, as is visible from the sixth and seventh columns of Table 6. Second, we find that parameters describing the variance factor confirm what was recently observed in Dufays et al. (2022, Figure 7): The volatility is very persistent and the leverage effect is really strong. These findings contrast with those of previous studies by Fulop and Li (2019); Du and Luo (2019), where the variance factor is not very persistent and the leverage effect is much less pronounced. The reason for the different results is the use of less informative data (i.e. variance swaps in Fulop and Li, 2019) and a reduced option panel (which ensures that the estimation in Du and Luo, 2019 is computationally tractable). Even in case of the nested model, jump size parameters are quite large in absolute value. We discuss this issue in Section 5.3. Third, the nested model specification fits the data worse compared to the full model, as indicated by the lower log-likelihood. To better illustrate this point, we present a comparison between market- and model-implied risk-neutral cumulants in terms of RMSEs in Table 8. The table shows that the DJDVSI model outperforms the nested one in 15 out of 18 cases. Therefore, the addition of a second variance factor is crucial in order to fit the time series of market risk-neutral cumulants.

Table 7
Root Mean Squared Pricing Errors (RMSPEs) in the DJDVSI model.

	QB-SMC								
$\tau \setminus K$	80	85	90	95	100	105	110	115	120
1	0.0444	0.0389	0.0319	0.0211	0.0240	0.0264	0.0290	0.0360	0.0442
2	0.0221	0.0236	0.0218	0.0164	0.0150	0.0221	0.0270	0.0330	0.0381
3	0.0259	0.0265	0.0237	0.0180	0.0130	0.0185	0.0245	0.0316	0.0378
6	0.0327	0.0304	0.0263	0.0201	0.0132	0.0147	0.0204	0.0275	0.0357
9	0.0323	0.0294	0.0255	0.0202	0.0131	0.0140	0.0198	0.0267	0.0339
12	0.0316	0.0279	0.0243	0.0200	0.0132	0.0144	0.0203	0.0269	0.0334
	QMLE								
$\tau \setminus K$	80	85	90	95	100	105	110	115	120
1	0.0443	0.0425	0.0414	0.0268	0.0248	0.0325	0.0306	0.0354	0.0435
2	0.0238	0.0292	0.0298	0.0204	0.0141	0.0278	0.0319	0.0341	0.0360
3	0.0275	0.0311	0.0298	0.0213	0.0115	0.0219	0.0296	0.0341	0.0375
6	0.0337	0.0330	0.0296	0.0225	0.0127	0.0147	0.0237	0.0311	0.0375
9	0.0337	0.0320	0.0286	0.0229	0.0138	0.0133	0.0211	0.0293	0.0362
12	0.0341	0.0312	0.0279	0.0234	0.0150	0.0137	0.0204	0.0282	0.0351
	QB-SMC/QMLE								
$\tau \setminus K$	80	85	90	95	100	105	110	115	120
1	1.0021	0.9146	0.7699	0.7866	0.9663	0.8122	0.9455	1.0159	1.0159
2	0.9294	0.8079	0.7305	0.8018	1.0585	0.7946	0.8476	0.9685	1.0586
3	0.9412	0.8522	0.7962	0.8432	1.1250	0.8445	0.8255	0.9260	1.0073
6	0.9718	0.9224	0.8873	0.8950	1.0426	0.9993	0.8625	0.8850	0.9522
9	0.9571	0.9208	0.8917	0.8833	0.9479	1.0495	0.9389	0.9099	0.9381
12	0.9278	0.8951	0.8687	0.8556	0.8827	1.0502	0.9955	0.9546	0.9521

Table 8
Root Mean Squared Errors of model implied risk-neutral cumulants for the DJDVSI model (top panel) and for the DJDVSI model (middle panel). In the bottom panel we report the ratio between the RMSEs obtained with the DJDVSI and DJSVSI models, respectively. Values in green (respectively, red) denote best performance for DJDVSI (DJSVSI) model.

	DJDVSI					
	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 6$	$\tau = 9$	$\tau = 12$
$c^{\text{var}}_{t,\tau}$	1.33E-03	2.37E-03	3.28E-03	5.65E-03	7.80E-03	1.01E-02
$c^{\text{skew}}_{t,\tau}$	3.33E-04	3.15E-04	4.84E-04	1.25E-03	2.50E-03	3.87E-03
$c^{\text{kurt}}_{t,\tau}$	3.40E-04	5.52E-04	8.62E-04	1.91E-03	3.87E-03	7.45E-03
	DJSVSI					
	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 6$	$\tau = 9$	$\tau = 12$
$c^{\text{var}}_{t,\tau}$	1.61E-03	2.83E-03	4.17E-03	7.52E-03	1.02E-02	1.29E-02
$c^{\text{skew}}_{t,\tau}$	3.34E-04	3.80E-04	6.11E-04	1.63E-03	3.34E-03	5.44E-03
$c^{\text{kurt}}_{t,\tau}$	2.63E-04	3.13E-04	5.76E-04	1.99E-03	4.51E-03	8.27E-03
	DJDVSI/DJSVSI					
	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 6$	$\tau = 9$	$\tau = 12$
$c^{\text{var}}_{t,\tau}$	0.826	0.838	0.788	0.752	0.764	0.787
$c^{\text{skew}}_{t,\tau}$	0.996	0.831	0.792	0.767	0.750	0.711
$c^{\text{kurt}}_{t,\tau}$	1.294	1.765	1.497	0.958	0.858	0.901

5.3. Sensitivity of parameter estimates to the choice of the dataset

To investigate why the absolute values of the jump parameters are higher than those in the existing literature, we conducted the following experiment. We implemented our QB-SMC based on three different datasets: (i) $\mathcal{D}_1 = \{\hat{c}^{\text{var}}_{t,\tau}\}$ considers only the measure of variance as observed variable; (ii) a dataset that also includes the measure of skewness $\mathcal{D}_2 = \{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}\}$, and (iii) the full dataset $\mathcal{D}_3 = \{\hat{c}^{\text{var}}_{t,\tau}, \hat{c}^{\text{skew}}_{t,\tau}, \hat{c}^{\text{kurt}}_{t,\tau}\}$. Parameter estimates are reported in the second, third, and fourth columns of Table 6. First, with the exception of jump sizes, most parameters remain relatively similar for the different datasets. However, differences in the magnitude of jump sizes are evident. More specifically, when using $\mathcal{D}_1$, the jump sizes are relatively small in absolute terms and are consistent with the literature (e.g. Fulop and Li, 2019). When using $\mathcal{D}_2$, the size of the jumps increases considerably. Finally, when using the full dataset $\mathcal{D}_3$, the jump sizes increase again in absolute value. This is particularly evident for the parameter σ_J. We show that these values of the jump parameters are necessary to match higher order cumulants time series. More specifically, we plotted in Figure 7 market and model implied cumulants for the various datasets considered. We find that using only $\mathcal{D}_1$, the time series of $\{\hat{c}^{\text{var}}_{t,\tau}\}$ matches perfectly, but the fit for the observed measures of higher order cumulants is very poor. Using $\mathcal{D}_2$, the model is able to reproduce both market variance and skewness, but the performance in reproducing the market kurtosis is quite poor. Therefore, in order to match all observed measures of market variance, skewness, and kurtosis, it seems necessary to increase the size of the jumps. In this case, the proposed methodology reproduces the observed time series very well. In summary, the fit of the market implied cumulants improves gradually. Thus, we conclude that large jump sizes are necessary to correctly match the observed time

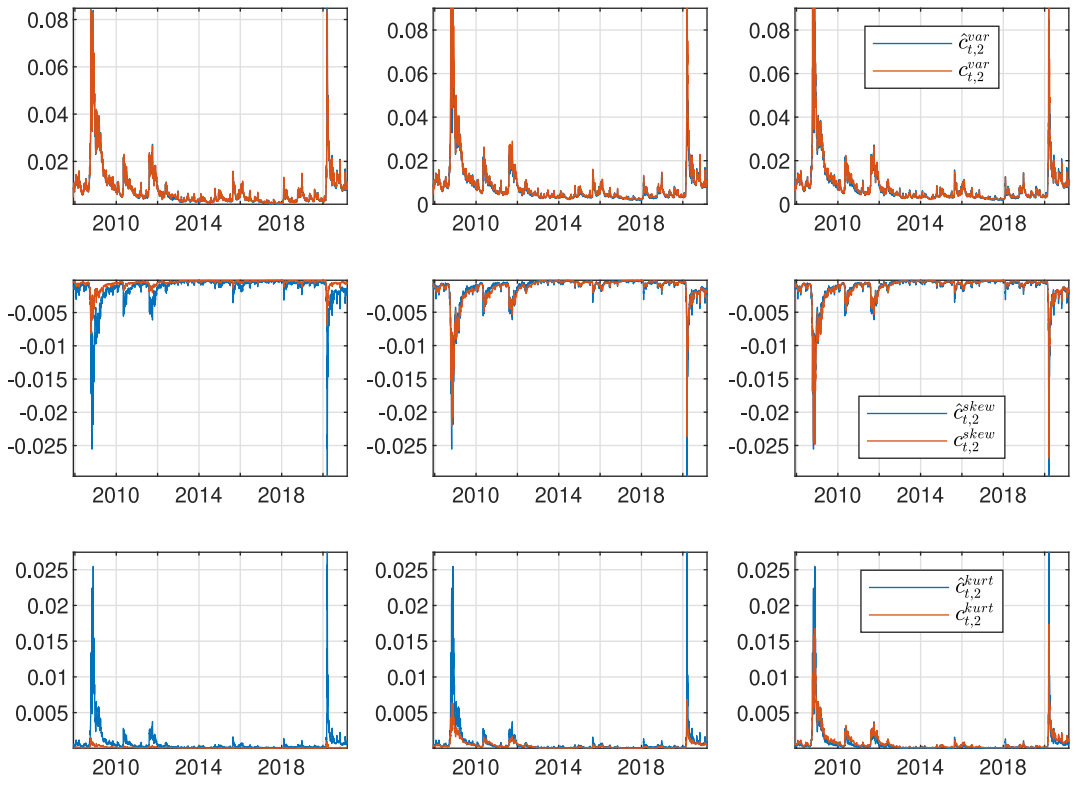

Fig. 7. Cumulants matching for different choices of the dataset: $\mathcal{D}_1$ (first column), $\mathcal{D}_2$ (second column) and $\mathcal{D}_3$ (third column). Blue lines represent the time series of the market implied risk-neutral cumulants, while orange lines represents the model implied risk-neutral cumulants for $\tau = 2$ months. Model implied cumulants are computed from the DJDVSI model using the QB-SMC approach, the parameter estimates are given in Table 6.

series of risk-neutral cumulants. In addition, the third column of Figure 7 provides further evidence of the high precision of our estimates: market- and model-implied cumulants are very close. This provides confidence that the proposed approach performs well and that the DJDVSI is flexible enough to match the time series of risk-neutral cumulants. To keep the paper compact, we report only the results for the case of two months maturity; additional results are available upon request. Next, we examine the impact of the choice of dataset on option pricing. To this end, we compute the RMSPEs for all the estimated parameters for each day of the dataset. Overall, the final RMSPE is 0.0313 for $\mathcal{D}_1$, 0.0274 for $\mathcal{D}_2$, 0.0256 for $\mathcal{D}_3$. From this, it is evident that conducting inference using the risk-neutral variance alone is not sufficient to correctly match the time series of option prices. However, comparing $\mathcal{D}_2$ with $\mathcal{D}_3$, the overall reduction in RMSPE is rather small. Figure 8 shows the running mean of the RMSPE. The main benefit of including the risk-neutral kurtosis in the dataset arises during turbulent periods, such as the 2008 global financial crisis and the recent COVID-19 outbreak. Indeed, during these periods we observe spikes in the time series of $\hat{c}^{\text{kurt}}_{t,\tau}$ (see e.g. Figure 1) that cannot be properly replicated when using $\mathcal{D}_2$ for the estimation. Therefore, $\hat{c}^{\text{kurt}}_{t,\tau}$ is far from being 0 in these periods and plays a crucial role in option pricing. In particular, we observe that during calm periods the difference between using $\mathcal{D}_2$ and $\mathcal{D}_3$ is very small, while during financial crises the increase of RMSPE is much more pronounced. The loss of accuracy is exacerbated for the shortest maturity (1 month).

5.4. Risk Premia

In this subsection, we study how to incorporate risk premia into our estimation framework. To this end, we follow Feunou and Okou (2018) and assume that, in the inferential procedure, the measurement equation (13) is defined under the risk-neutral measure, while the transition equation (14) evolves under the historical measure. Therefore, to move from one measure to the other, we perform a structure-preserving change of measure by shifting the parameters $k_1, \theta_1, k_2, \theta_2, k_\lambda, \theta_\lambda$ and leaving the other parameters unchanged (see the fourth column of Table 6). More precisely, this is an application of Girsanov's theorem, which translates in the following relationships:

$$k_1 = k_1^{\mathbb{P}} + \sigma_1\phi_1^v,\ \theta_1 = k_1^{\mathbb{P}}\theta_1^{\mathbb{P}}/k_1,$$
$$k_2 = k_2^{\mathbb{P}} + \sigma_2\phi_2^v,\ \theta_2 = k_2^{\mathbb{P}}\theta_2^{\mathbb{P}}/k_2,$$
$$k_\lambda = k_\lambda^{\mathbb{P}} + \sigma_\lambda\phi_\lambda,\ \theta_\lambda = k_\lambda^{\mathbb{P}}\theta_\lambda^{\mathbb{P}}/k_\lambda,$$

where the parameters with superscript $\mathbb{P}$ refer to the corresponding coefficients when the dynamics (1)–(4) evolve under the historical measure. ϕ_1^v, ϕ_2^v and ϕ_λ denote the risk premia of the first and second variance factor and the jump intensity, respectively. To estimate ϕ_1^v, ϕ_2^v and ϕ_λ, we run our QB-SMC using $N = 2000$ particles and with the following independent truncated normal priors $\Theta_\phi \sim \mathcal{TN}(\mu_\phi, I_3)$, where $\mu_\phi = [-0.5, 2, -0.5]$ and I_3 is a 3×3 identity matrix. In line with the existing literature (see for example Feunou and Okou, 2018; Fulop and Li, 2019), we truncate the domain of the prior such that $\phi_1^v \in (-\infty, 0)$, $\phi_2^v \in (0, \infty)$ and $\phi_\lambda \in (-\infty, 0)$. We find the following parameter estimates, which are statistically significant at the 5% level (standard error in parenthesis): $\phi_1^v = -0.1373\ (0.0221)$, $\phi_2^v = 1.2522\ (0.3779)$ and $\phi_\lambda = -0.0929\ (0.0466)$. Given the estimated ϕ_1^v, ϕ_2^v, ϕ_λ, we compute the implied risk-neutral and historical variance, skewness, and kurtosis, and take their difference (compare with Feunou and Okou, 2018, Figure 7). In this way, we obtain a dynamic representation of the risk premia implied by higher order moments. In Figure 9, we show the time series of variance, skewness, and kurtosis risk premia for $\tau = 1$ month, defined as the difference between the risk-neutral and the historical series. We find that the implied variance risk premium is pos-

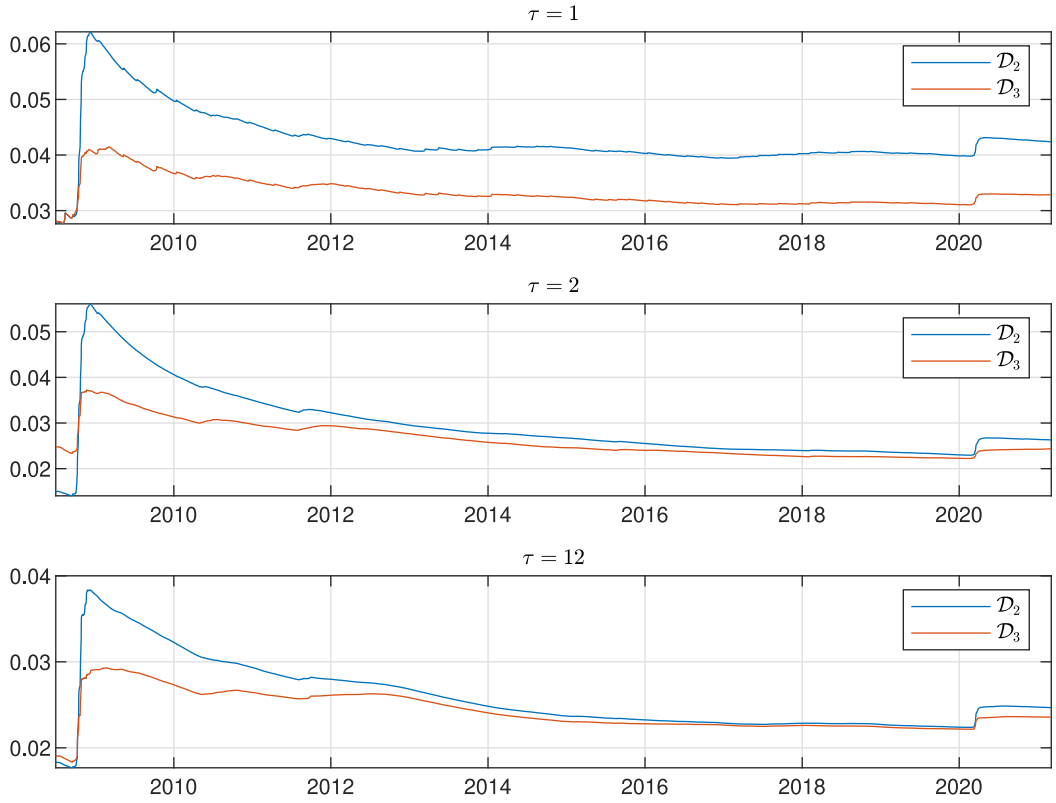

Fig. 8. Running mean of the RMSPE for three selected maturities calculated using $\mathcal{D}_2$ (blue line) and $\mathcal{D}_3$ (orange line).

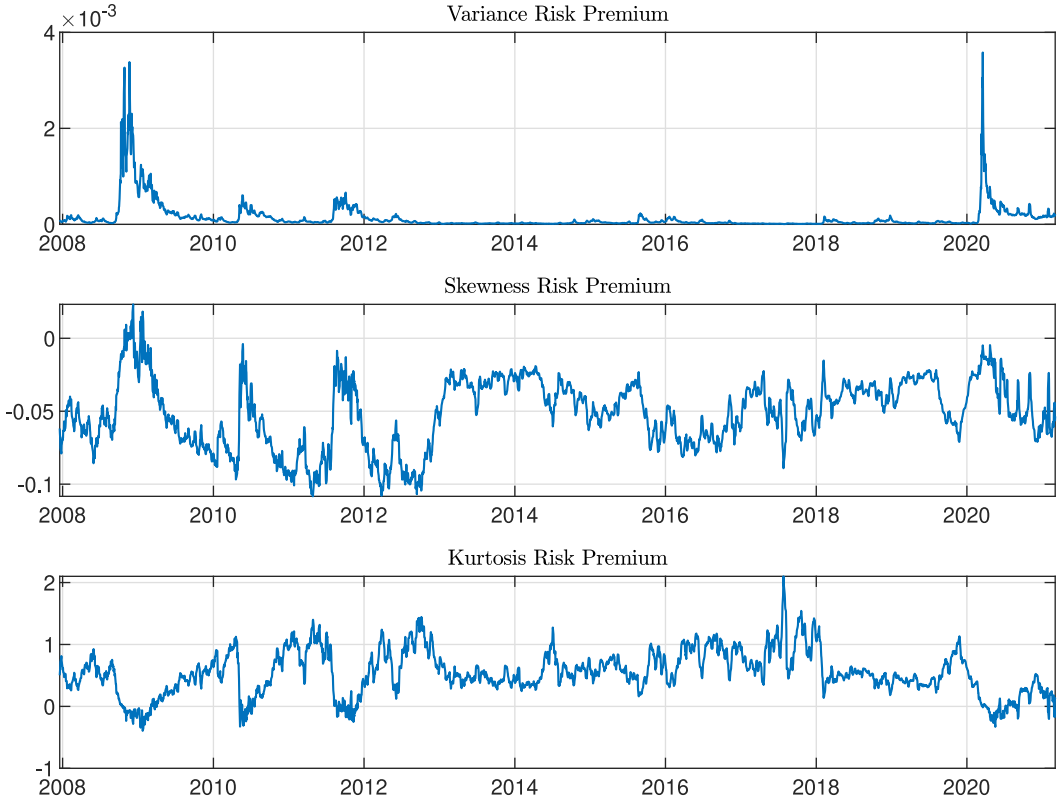

Fig. 9. Model-implied moment risk premia.

itive. This is in line with the results in Bollerslev et al. (2009), Feunou and Okou (2018) and Li and Zinna (2018). The skewness risk premia is negative, confirming the empirical evidence in Kozhan et al. (2013); Feunou et al. (2017); Feunou and Okou (2018). Finally, we obtain that the kurtosis risk premium is mostly positive. This result differs somewhat from Feunou and Okou (2018), who find more mixed evidence. However, kurtosis risk premium should be positive, as discussed in Rauch and Alexander (2016). This gives us confidence that our estimation approach can be exploited to adequately analyze risk premia and their financial implications.

6. Conclusions

In this paper, we propose a new Quasi-Bayesian method based on a modified Kalman filter and a density tempered SMC sampler for estimating affine option pricing models using informative portfolios of weighted options (i.e., the risk-neutral cumulants of log-returns). Through extensive Monte Carlo studies, we show that the new approach can be viewed as an efficient global optimizer that makes the estimation of complex state-space models relatively simple and fast. The method provides accurate parameter estimates thanks to the efficiency of density tempered SMC samplers. We apply our methodology to estimate a state-of-the-art affine option pricing model on real data. Numerical results confirm the accuracy and computational efficiency of the method. We compare our method to maximum likelihood estimation approaches and find that it is superior in terms of pricing errors. Moreover, our empirical results confirm the importance of a second variance factor in accurately fitting the time series of higher order risk-neutral cumulants. Finally, we find that the inclusion of kurtosis allows for a more accurate estimation of the model, for both real and simulated data, than when only variance and skewness are used. In addition, this allows to reduce option pricing errors, especially during periods of market stress.

Data Availability

The authors do not have permission to share data.

CRediT authorship contribution statement

Riccardo Brignone: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization. **Luca Gonzato:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization, Supervision. **Eva Lütkebohmert:** Conceptualization, Formal analysis, Writing – review & editing, Supervision, Funding acquisition.

Acknowledgments

We thank the editor and two anonymous reviewers for various constructive remarks that helped to significantly improve our paper. Moreover, we are indebted to Andras Fulop for fruitful comments and discussions that helped broaden the scope of this research and improved the content at various levels. Earlier versions of this work have been presented at the 1st Florence-Paris Workshop on Mathematical Finance, the 10th AMaMeF Conference, the 27th International Conference on Computing in Economics and Finance (CEF 2021) and at the Econometric Society Australasian Meeting (ESAM 2021), we thank all participants for useful feedbacks. This work was supported by the German Research Foundation (DFG) through the grant LU 1186/4-1. Financial support is gratefully acknowledged by the first and the third author.

Appendix A. Moment Generating Function

We define the MGF of log-returns as follows

$$\Psi(u, X_t, V_{1t}, V_{2t}, \lambda_t, t, T) = \mathbb{E}^{\mathbb{Q}}[e^{uX_T}|\mathcal{F}_t]. \tag{A.1}$$

By exploiting the Feynman-Kac theorem we get the following partial differential equation for $\tau = T - t$

$$\begin{aligned}
&-\Psi_\tau + \Big(r - \frac{1}{2}(V_{1t}+V_{2t}) - \lambda_t\mu^\star\Big)\Psi_x + \frac{1}{2}(V_{1t}+V_{2t})\Psi_{xx} + k_1(\theta_1 - V_{1t})\Psi_{v_1} + \\
&+ \frac{1}{2}\sigma_1^2 V_{1t}\Psi_{v_1v_1} + \rho_1\sigma_1 V_{1t}\Psi_{xv_1} + k_2(\theta_2 - V_{2t})\Psi_{v_2} + \frac{1}{2}\sigma_2^2 V_{2t}\Psi_{v_2v_2} \\
&+ \rho_2\sigma_2 V_{2t}\Psi_{xv_2} + k_\lambda(\theta_\lambda - \lambda_t)\Psi_\lambda + \frac{1}{2}\sigma_\lambda^2\lambda_t\Psi_{\lambda\lambda} \\
&+ \lambda_t\int_{-\infty}^{\infty}\int_0^{\infty}[\Psi(u, X_t + J_x, V_{1t}, V_{2t} + J_v, \lambda_t, \tau) \\
&-\Psi(u, X_t, V_{1t}, V_{2t}, \lambda_t, \tau)]\nu(dJ_x, dJ_v) = 0.
\end{aligned} \tag{A.2}$$

Since the model is affine we can guess a solution of the form

$$\begin{aligned}
&\Psi(u, X_t, V_{1t}, V_{2t}, \lambda_t, \tau) \\
&= \exp\Big(uX_t + A(u,\tau) + B(u,\tau)V_{1t} + C(u,\tau)V_{2t} + D(u,\tau)\lambda_t\Big).
\end{aligned}$$

For the jump transform we guess

$$\begin{aligned}
&\Psi(u, X_t + J_x, V_{1t}, V_{2t} + J_v, \lambda_t, \tau) - \Psi(u, X_t, V_{1t}, V_{2t}, \lambda_t, \tau) = \\
&= \Psi(u, X_t, V_{1t}, V_{2t}, \lambda_t, \tau)\Big[e^{uJ_x + J_vC(u,\tau)} - 1\Big].
\end{aligned}$$

Now, we need the partial derivatives of Ψ

$$\begin{aligned}
&\Psi_\tau = \Psi(A_\tau(u,\tau) + B_\tau(u,\tau)V_{1t} + C_\tau(u,\tau)V_{2t} + D_\tau(u,\tau)\lambda_t), \\
&\Psi_x = \Psi u, \quad \Psi_{v_1} = \Psi B(u,\tau), \Psi_{v_1v_1} = \Psi B(u,\tau)^2, \\
&\Psi_{xv_1} = \Psi uB(u,\tau), \quad \Psi_{xx} = \Psi u^2, \quad \Psi_{v_2} = \Psi C(u,\tau), \\
&\Psi_{v_2v_2} = \Psi C(u,\tau)^2, \quad \Psi_{xv_2} = \Psi uC(u,\tau), \quad \Psi_\lambda = \Psi D(u,\tau), \\
&\Psi_{\lambda\lambda} = \Psi D(u,\tau)^2.
\end{aligned}$$

By substituting the partial derivatives into (A.2) we obtain the following expression

$$\begin{aligned}
&-(A_\tau(u,\tau) + B_\tau(u,\tau)V_{1t} + C_\tau(u,\tau)V_{2t} + D_\tau(u,\tau)\lambda_t) \\
&+ (r - 0.5(V_{1t}+V_{2t}) - \lambda_t\mu^\star)u \\
&+ \frac{1}{2}V_{1t}u^2 + k_1(\theta_1 - V_{1t})B(u,\tau) + \frac{1}{2}\sigma_1^2 V_{1t}B(u,\tau)^2 + \rho_1\sigma_1 V_{1t}B(u,\tau)u \\
&+ \frac{1}{2}V_{2t}u^2 + k_2(\theta_2 - V_{2t})C(u,\tau) + \frac{1}{2}\sigma_2^2 V_{2t}C(u,\tau)^2 \\
&+ \rho_2\sigma_2 V_{2t}C(u,\tau)u + k_\lambda(\theta_\lambda - \lambda_t)D(u,\tau) \\
&+ \frac{1}{2}\sigma_\lambda^2\lambda_t D(u,\tau)^2 + \lambda_t\int_{-\infty}^{\infty}\int_0^{\infty}\Big[e^{uJ_x + J_vC(u,\tau)} - 1\Big]\nu(dJ_x, dJ_v) = 0,
\end{aligned} \tag{A.3}$$

where $\int_{-\infty}^{\infty}\int_0^{\infty}\Big[e^{uJ_x + J_vC(u,\tau)} - 1\Big]\nu(dJ_x, dJ_v) = \frac{e^{u\mu_J + u^2\sigma_J^2/2}}{1 - C(u,\tau)\mu_v} - 1$. Finally, separation of variables leads to the ODE system (6).

Appendix B. Analytic Solution for (11): Matlab® Codes

In what follows we present a Matlab® code for computing the cumulants in (11) for the model (1)–(4). By simply modifying the expression of the functional characteristics (i.e. the right hand side of the equations in (6), see Hubalek et al., 2017 and the references therein) it is possible to get the expression for cumulants in other 4-factor affine models.

```
% define functional characteristics
syms F(u,tau) R(u,tau) Z(u,tau) Y(u,tau)
% define auxiliary functions
syms B(u,tau) C(u,tau) D(u,tau)
% define symbolic variables
syms k1 th1 sg1 rho1 k2 th2 sg2 rho2 muv kL thL sgL muJ sgJ V10 V20 L0
syms t
F(u, tau) = k1*th1*B(u,tau) + k2*th2*C(u, tau) + kL*thL*D(u, tau) ;
R(u, tau) = -1/2*(u - u2) - (k1 - rho1*sg1*u)*B(u, tau) + 1/2*sg12*B(u, tau)2;
Z(u, tau) = -1/2*(u - u2) - (k2 - rho2*sg2*u)*C(u, tau) + 1/2*sg22*C(u, tau)2;
Y(u, tau) = -kL*D(u, tau) + 1/2*sgL2*D(u, tau)2 + (exp(u*muJ + u2*sgJ2/2)/...
(1-C(u,tau)*muv)-1) - (exp(muJ + sgJ2/2) - 1)*u;
% Number of needed cumulants
M = 3;
% define auxiliary functions for solving ODEs
syms x(t) y(t) z(t) q(t)
for i=1:M
   % Compute i - th derivative of R with respect to u
   Rprime = diff(R(u,tau),u,i);
   % Prepare the output for the ODE solution setting i - th derivative
   % value equal to generic x(t)
   Rprime = subs(Rprime, diff(B(u, tau), u, i), x(t));
   % Substitute all the precedently computed derivatives with respect to u
   for j=1:i-1
      Rprime = subs(Rprime, diff(B(u, tau), u, j), DB(j));
   end
   % Substitute B(0, tau) ->0
   Rprime = subs(Rprime, B(u, tau),0);
   % and u ->0
   Rprime = subs(Rprime, u, 0);
   % Solve the ODE analytically
   eqn = diff(x,t) == Rprime; % define equation
   cond = x(0) == 0; % initial condition
   DB(i) = dsolve(eqn,cond,'MaxDegree',2); % solve ODE
   % Compute i - th derivative of R with respect to u
   Zprime = diff(Z(u,tau),u, i);
   % Prepare the output for the ODE solution setting i - th derivative
   % value equal to generic y(t)
   Zprime = subs(Zprime, diff(C(u, tau), u, i), y(t));
   % Substitute i-th derivative value with its value computed before
   Zprime = subs(Zprime, diff(B(u, tau), u, i), DB(i));
   % Substitute all the precedently computed derivatives with respect to u
   for j=1:i-1
      Zprime = subs(Zprime, diff(B(u, tau), u, j), DB(j));
      Zprime = subs(Zprime, diff(C(u, tau), u, j), DC(j));
   end
   % Substitute B(0, tau) ->0
   Zprime = subs(Zprime, B(u, tau),0);
   % Substitute C(0, tau) ->0
   Zprime = subs(Zprime, C(u, tau),0);
   % and u ->0
   Zprime = subs(Zprime, u, 0);
   % Solve the ODE analytically
   eqn = diff(y,t) == Zprime; % define equation
   cond = y(0) == 0; % initial condition
   DC(i) = dsolve(eqn,cond,'MaxDegree',2); % solve ODE
   % Compute i - th derivative of R with respect to u
   Yprime = diff(Y(u,tau),u, i);
   % Prepare the output for the ODE solution setting i - th derivative
   % value equal to generic y(t)
   Yprime = subs(Yprime, diff(D(u, tau), u, i), q(t));
   % Substitute i-th derivative value with its value computed before
   Yprime = subs(Yprime, diff(C(u, tau), u, i), DC(i));
   % Substitute all the precedently computed derivatives with respect to u
   for j=1:i-1
      Yprime = subs(Yprime, diff(B(u, tau), u, j), DB(j));
      Yprime = subs(Yprime, diff(C(u, tau), u, j), DC(j));
      Yprime = subs(Yprime, diff(D(u, tau), u, j), DD(j));
   end
   % Substitute B(0, tau) ->0
   Yprime = subs(Yprime, B(u, tau),0);
   % Substitute C(0, tau) ->0
   Yprime = subs(Yprime, C(u, tau),0);
   % Substitute D(0, tau) ->0
   Yprime = subs(Yprime, D(u, tau),0);
   % and u ->0
   Yprime = subs(Yprime, u, 0);
   % Solve the ODE analytically
   eqn = diff(q,t) == Yprime; % define equation
   cond = q(0) == 0; % initial condition
   DD(i) = dsolve(eqn,cond,'MaxDegree',2); % solve ODE
   % Compute i - th derivative of F with respect to u
   Fprime = diff(F(u,tau), u, i);
   % Substitute i-th derivative of B and C with its value computed before
   Fprime = subs(Fprime, diff(D(u, tau), u, i), DD(i));
   Fprime = subs(Fprime, diff(C(u, tau), u, i), DC(i));
   Fprime = subs(Fprime, diff(B(u, tau), u, i), DB(i));
   % Solve ODE analytically
   eqn = diff(z,t) == Fprime; % define equation
   cond = z(0) == 0; % initial condition
   DA(i) = dsolve(eqn,cond); % solve ODE
end
```

Appendix C. Computation of Risk-Neutral Cumulants with Self-exciting Jumps

We tested our proposed econometric approach on many different models. Obviously, all the affine option pricing models nested in our formulation (1)–(4) can be estimated through our QB-SMC algorithm (e.g. the Double Heston model in Christoffersen et al., 2009 or Model IV in Fulop and Li, 2019). However, we encountered some technical difficulty when applying our methodology to models which allow for self-excitation, an important stylized feature of financial returns. More specifically, we tried to estimate Model III in Fulop and Li (2019), which is a two-factor model allowing for a self-exciting jump intensity and co-jumps between price and volatility. For this model (using the same notation as Fulop and Li, 2019) the related functional characteristics[7] are given by:

$$\begin{aligned}
\frac{\partial A(u,\tau)}{\partial \tau} &= ru + k_\lambda \theta_\lambda C(u,\tau) + k_v \theta_v B(u,\tau),\\
\frac{\partial B(u,\tau)}{\partial \tau} &= -0.5(u-u^2) - (k_v - \rho\sigma_v u)B(u,\tau) + 0.5\sigma_v^2 B(u,\tau)^2,\\
\frac{\partial C(u,\tau)}{\partial \tau} &= -k_\lambda C(u,\tau) + 0.5\sigma_\lambda^2 C(u,\tau)^2 + (\exp(u\mu_J + u^2\sigma_J^2/2 + \\
&\quad C(u,\tau)\beta)/(1 - B(u,\tau)\mu_v) - 1) - (\exp(\mu_J + \sigma_J^2/2) - 1)u,
\end{aligned} \tag{C.1}$$

where $\beta > 0$ controls the level of self-excitation. In this case, we were not able to obtain the symbolic closed form expression for $\frac{\partial^3 A(0,\tau)}{\partial u^3}$. Indeed, Matlab® failed in solving analytically the last ordinary differential equation due to memory problems. This is primarily just a numerical issue, because the calculation should be theoretically possible (Feunou and Okou, 2018, Propositions 1–2). The problem is that the mathematical expressions needed to solve the ODE

$$\frac{\partial\left[\frac{\partial^3 A(0,\tau)}{\partial u^3}\right]}{\partial \tau} = k_\lambda \theta_\lambda \frac{\partial^3 C(0,\tau)}{\partial u^3} + k_v \theta_v \frac{\partial^3 B(0,\tau)}{\partial u^3} \tag{C.2}$$

become extremely long and the computation failed on our standard PCs. The difficulty arises from the term $\exp(u\mu_J + u^2\sigma_J^2/2 + C(u,\tau)\beta)/(1 - B(u,\tau)\mu_v)$ in (C.1), which makes the ODE for (C.2) very cumbersome and complicated to solve. We do not report the above mentioned expression here as it is illegibly long, but it is available upon request. To investigate this aspect further, we considered the special case $\mu_v = 0$ and found that the analytic expressions for all necessary quantities $\frac{\partial^n A(u,\tau)}{\partial u^n}\Big|_{u=0}$, $\frac{\partial^n B(u,\tau)}{\partial u^n}\Big|_{u=0}$, $\frac{\partial^n C(u,\tau)}{\partial u^n}\Big|_{u=0}$ for $n = \{1,2,3\}$ can be computed in approximately 2 minutes. Thus, to avoid the computational problems we encountered when estimating Model III in Fulop and Li (2019), we decided to specify another model without self-excitation, but which

[7] See for example Hubalek et al. (2017) for a definition of functional characteristics.

is still flexible enough and which allows an easy calculation of risk-neutral cumulants. Therefore, instead of self-excitation we added another volatility factor that allows for jumps. However, as explained above, we stress that the main problem is not directly the self-excitation property, but the simultaneous presence of self-excitation and co-jumps, which complicates the symbolic computations. In simpler models with self-excitation (e.g. Hainaut and Moraux, 2018; Bernis et al., 2021), we were able to obtain analytic expressions for slopes coefficients, as also found in Brignone and Sgarra (2020, Section 3.2).

Appendix D. Accuracy of the Filtering Method

To assess the accuracy of the modified Kalman filter, we perform the following experiment. First, using the parameters estimated on real data, we conducted a single experiment on simulated data (15 years of daily data) to show that solving the filtering problem with this modified Kalman filter results in negligible bias. More precisely, in Figure D.10 we compare the true (simulated) trajectories of the total variance and the jump intensity with their filtered counterparts. The two trajectories are nearly indistinguishable, suggesting that using this approximated technique together with risk-neutral cumulants yields very good results, even in the presence of very pronounced jumps. Second, following Du and Luo (2019) we computed the RMSE between the true and filtered states by running 100 repetitions of the algorithm on 100 different synthetic datasets and we calculated the mean and standard deviation of the RMSE. We conclude that the bias induced by the modified Kalman filter is very small. Indeed, for the total variance process the RMSE mean is 8.55×10^{-4} and the standard deviation is 3.11×10^{-4}, while the RMSE mean for the jump intensity is equal to 0.0047 with a standard deviation of 0.0012. Overall, we found that their magnitude is comparable to some other results reported in the literature, e.g. Du and Luo (2019, Appendix C).

Appendix E. Additional Results

In this section, we report some additional results for Sections 4 and 5. Figures E.11 and E.12 complement Figures 3 and 4 showing, respectively, the bridge between priors and quasi posteriors and the comparison between full sample quasi posteriors and the prior distribution for the remaining parameters. Also in this case, the results show good convergence for all parameters. It is also clear from the prior/posterior comparison that the prior is not informative for the final output.

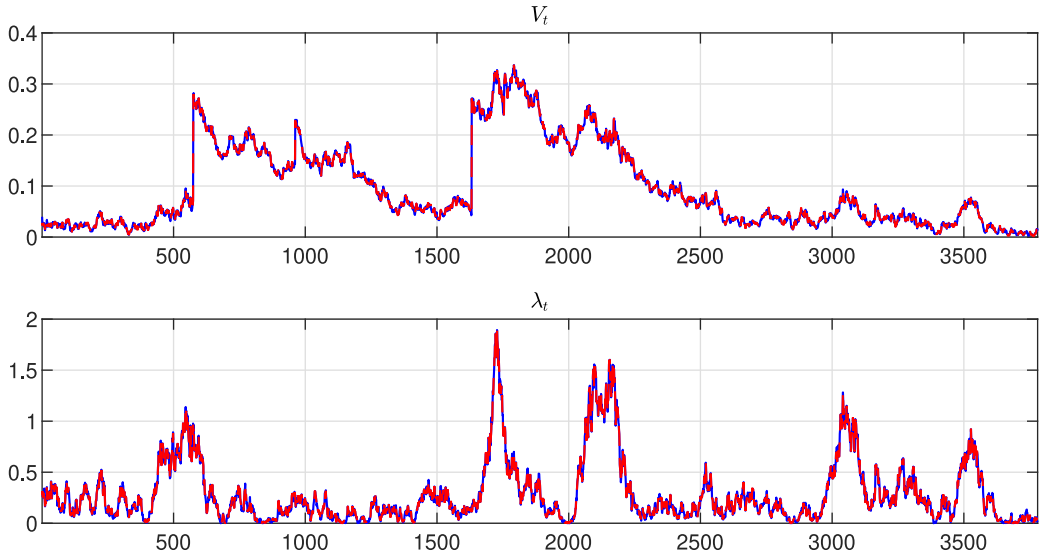

Fig. D1. Latent states on simulated data. Parameters as those in Table 6. The solid blue and dashed red lines represent, respectively, the filtered and true latent states.

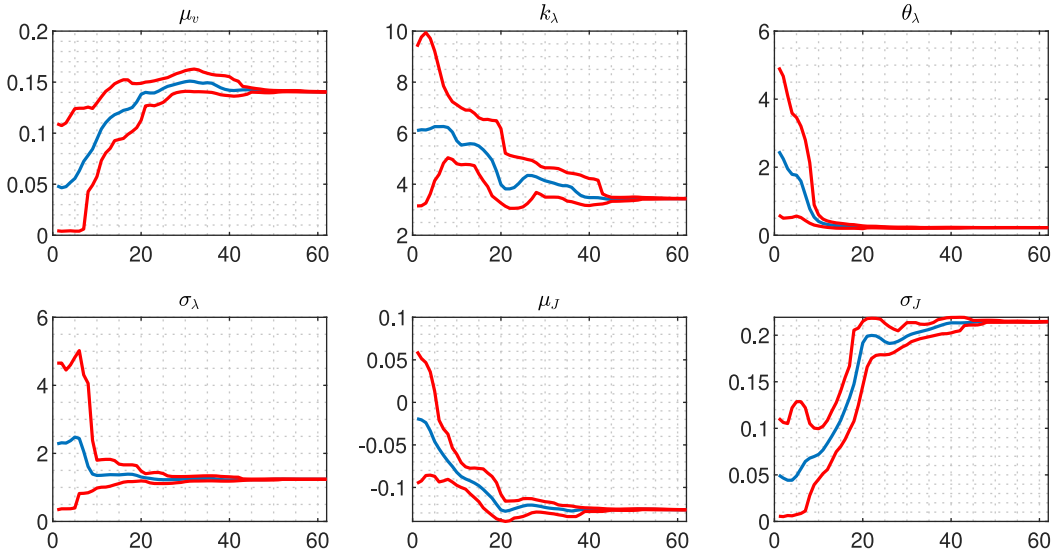

Fig. E1. Bridging the priors and the quasi-posteriors: mean (blue line) and (5, 95)% quantiles (red lines) of some selected parameters at each tempering step.

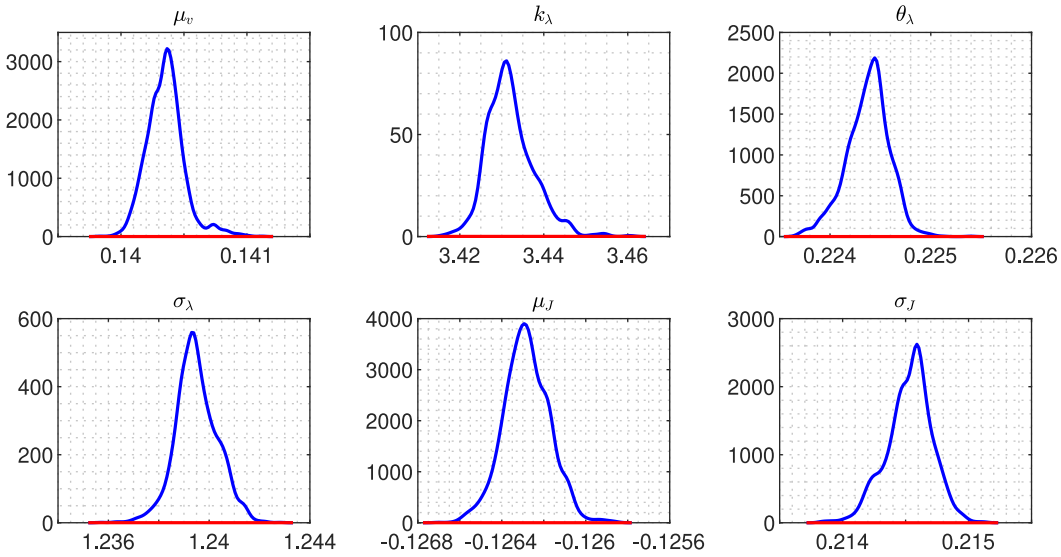

Fig. E2. Full sample quasi-posterior distributions (blue line), obtained with a kernel density estimator, and prior distributions (red line) of some selected parameters.

References

Andersen, T., Fusari, N., Todorov, V., 2015. The risk premia embedded in index options. Journal of Financial Economics 117 (3), 558–584.

Andrieu, C., Doucet, A., Holenstein, R., 2010. Particle Markov chain Monte Carlo. Journal of the Royal Statistical Society. Series B 72 (3), 269–342.

Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws and the differential pricing of individual equity options. Review of Financial Studies 16, 101–143.

Bardgett, C., Gourier, E., Leippold, M., 2019. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. Journal of Financial Economics 131 (3), 593–618.

Barndorff-Nielsen, O., 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 (2), 167–241.

Bates, D.S., 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. Review of Financial Studies 9 (1), 69–107.

Bernis, G., Brignone, R., Scotti, S., Sgarra, C., 2021. A Gamma Ornstein-Uhlenbeck model driven by a Hawkes process. Mathematics and Financial Economics 15, 747–773.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.

Brignone, R., Sgarra, C., 2020. Asian options pricing in Hawkes-type jump-diffusion models. Annals of Finance 16 (1), 101–119.

Brix, A.F., Lunde, A., Wei, W., 2018. A generalized Schwarz model for energy spot prices-Estimation using a particle MCMC method. Energy Economics 72, 560–582.

Callegaro, G., Gaïgi, M., Scotti, S., Sgarra, C., 2017. Optimal investment in markets with over and under-reaction to information. Mathematics and Financial Economics 11, 299–322.

Carr, P., Geman, H., Madan, D., Yor, M., 2003. Stochastic volatility for Lévy processes. Mathematical Finance 13, 345–382.

Chernozhukov, V., Hong, H., 2003. An MCMC approach to classical estimation. Journal of Econometrics 115, 293–346.

Chopin, N., Jacob, P.E., Papaspiliopoulos, O., 2013. SMC^2: an efficient algorithm for sequential analysis of state-space models. Journal of the Royal Statistical Society. Series B. 75 (3), 397–426.

Christoffersen, P., Heston, S., Jacobs, K., 2009. The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. Management Science 55 (12), 1914–1932.

Cox, J., Ingersoll, J., Ross, S., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–407.

Del Moral, P., Doucet, A., Jasra, A., 2006. Sequential Monte Carlo samplers. Journal of the Royal Statistical Society: Series B 68 (3), 411–436.

Du, D., Luo, D., 2019. The pricing of jump propagation: Evidence from spot and options markets. Management Science 65 (5), 2360–2387.

Duan, J.C., Fulop, A., 2015. Density-tempered marginalized sequential Monte Carlo samplers. Journal of Business and Economic Statistics 33, 192–202.

Dufays, A., Jacobs, K., Liu, Y., Rombouts, J., 2022. Fast filtering with large option panels: Implications for asset pricing. Working paper, Available at SSRN: https://ssrn.com/abstract=4066005 or https://doi.org/10.2139/ssrn.4066005.

Duffie, D., Filipović, D., Schachermayer, W., 2003. Affine processes and applications in finance. The Annals of Applied Probability 13 (3), 984–1053.

Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68 (6), 1343–1376.

Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance 59, 1367–1403.

Fang, F., Oosterlee, C., 2008. A novel pricing method for European options based on Fourier-cosine series expansions. SIAM Journal on Scientific Computing 31 (2), 826–848.

Feunou, B., Jahan-Parvar, M.R., Okou, C., 2017. Downside variance risk premium. Journal of Financial Econometrics 16 (3), 341–383.

Feunou, B., Okou, C., 2018. Risk-neutral moment-based estimation of affine option pricing models. Journal of Applied Econometrics 33, 1007–1025.

Fulop, A., Li, J., 2013. Efficient learning via simulation: a marginalized resample–move approach. Journal of Econometrics 176, 146–161.

Fulop, A., Li, J., 2019. Bayesian estimation of dynamic asset pricing models with informative observations. Journal of Econometrics 209 (1), 114–138.

Fulop, A., Li, J., Yu, J., 2015. Self-exciting jumps, learning, and asset pricing implications. The Review of Financial Studies 28 (3), 876–912.

Gonzato, L., Sgarra, C., 2021. Self-exciting jumps in the oil market: Bayesian estimation and dynamic hedging. Energy Economics 99, 105279.

Hainaut, D., Moraux, F., 2018. Hedging of options in the presence of jump clustering. Journal of Computational Finance 22 (3), 1–35.

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.

Hubalek, F., Keller-Ressel, M., Sgarra, C., 2017. Geometric Asian option pricing in general affine stochastic volatility models with jumps. Quantitative Finance 17 (6), 873–888.

Hurn, A.S., Lindsay, K.A., McClelland, A.J., 2015. Estimating the parameters of stochastic volatility models using option price data. Journal of Business and Economic Statistics 33, 579–594.

Kozhan, R., Neuberger, A., Schneider, P., 2013. The skew risk premium in the equity index market. The Review of Financial Studies 26 (9), 2174–2203.

Li, J., Yin, W., 2014. Macroeconomic fundamentals and the exchange rate dynamics: A no-arbitrage macro-finance approach. Journal of International Money and Finance 41, 46–64.

Li, J., Zinna, G., 2018. The variance risk premium: Components, term structures and stock return predictability. Journal of Business and Economic Statistics 36 (3), 411–425.

Monfort, A., Pegoraro, F., Renne, J.P., Rousselet, G., 2017. Staying at zero with affine processes: An application to term structure modelling. Journal of Econometrics 201 (2), 348–366.

Orlowski, P., 2021. Informative option portfolios in filter design for option pricing models. Quantitative Finance 21 (6), 945–965.

Rauch, J., Alexander, C., 2016. Tail risk premia for long-term equity investors. Working Paper.

Rosenthal, J.S., 2011. Optimal proposal distributions and adaptive MCMC. Handbook of Markov Chain Monte Carlo.

Ruge-Murcia, F., 2012. Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles. Journal of Economic Dynamics and Control 36 (6), 914–938.

Todorov, V., 2011. Econometric analysis of jump-driven stochastic volatility models. Journal of Econometrics 160 (1), 12–21.

Wachter, J.A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? The Journal of Finance 68 (3), 987–1035.

Yeap, C., Kwok, S., Choy, S., 2018. A flexible generalized hyperbolic option pricing model and its special case. Journal of Financial Econometrics 16 (3), 425–460.