

Stochastic Volatility Jump-Diffusions for European Equity Index Dynamics

Andreas Kaeck and Carol Alexander

ICMA Centre, Henley Business School at the University of Reading, UK
E-mails: a.kaeck@icmacentre.rdg.ac.uk; c.alexander@icmacentre.rdg.ac.uk

Abstract

Major research on equity index dynamics has investigated only US indices (usually the S&P 500) and has provided contradictory results. In this paper a clarification and extension of that previous research is given. We find that European equity indices have quite different dynamics from the S&P 500. Each of the European indices considered may be satisfactorily modelled using either an affine model with price and volatility jumps or a GARCH volatility process without jumps. The S&P 500 dynamics are much more difficult to capture in a jump-diffusion framework.

Keywords: *equity indices, jump-diffusions, generalised autoregressive conditional heteroscedasticity, GARCH, Markov chain Monte Carlo, MCMC*

JEL classification: C15, C32, G15.

1. Introduction

Accurate models of equity index dynamics are important for numerous applications in risk and portfolio management, including: non-vanilla option pricing; option portfolio hedging; hedging with futures; trading on equity and volatility risk premia; global equity portfolio allocation; basis arbitrage of new structured products such as variance swaps; and indeed *any* strategy for trading equity index-based products.

Motivated by some classic papers in the option pricing field – notably Heston (1993), Bates (1996) and Duffie *et al.* (2000) – state-of-the-art dynamic models feature stochastic volatility with price and volatility jumps.¹ Consequently these models have become a main topic for empirical research on equity index dynamics. The most influential articles (reviewed below) have only examined US equity indices and the vast majority of these focus exclusively on the S&P 500. Even so, many of the findings are contradictory. The only clear consensus to emerge is that the volatility of US equity indices evolves stochastically over time, it mean-reverts and is negatively correlated with the index

We would like to thank an anonymous referee for excellent comments on an earlier draft of this paper. Correspondence: Andreas Kaeck

¹ For interesting alternative ways to model option prices see e.g. Schönbucher (1999) or Skiadopoulos and Hodges (2001).

returns, and there are sudden jumps in the price process that cannot be captured by the price and volatility diffusion components.

The majority of stochastic volatility specifications will not admit even quasi-analytic solutions for vanilla option prices. However, the square-root model introduced by Heston (1993) belongs to the general class of affine models for which Fourier transform methods can provide tractable pricing solutions. Not surprisingly, therefore, most continuous-time equity index dynamics research has focused on jump extensions of this model. Apart from the Heston model, there are two other volatility specifications that have received particular attention in the literature: (a) a mean-reverting variance process with a diffusion coefficient proportional to variance raised to some exponent other than $1/2$, and (b) a mean-reverting diffusion for the log volatility. The most popular model of type (a) employs an exponent of 1 and a standard type (b) model is the log volatility diffusion introduced by Scott (1987).²

The literature on equity index dynamics has focused almost exclusively on the US. All papers reviewed in the following base their findings on two-factor continuous time models for the S&P 500 index, unless otherwise stated. Using data until the late 1990's, Andersen *et al.* (2002) tested the mean-reverting affine variance process of Heston (1993) against the type (b) alternatives above. They found that both specifications are adequate for modeling the S&P 500 dynamics and are structurally stable over time, provided they are augmented with jumps in prices. Moreover, Eraker *et al.* (2003) conclude that jumps in both volatility and price processes are necessary for the square root model, since variance can increase very rapidly – too rapidly to be captured by a square root diffusion.

Type (a) alternatives to the Heston model are tested in another strand of literature. Jones (2003) concludes that these alternatives provide more realistic dynamics, although they still fall short of explaining some features of the spot and option data. Chacko and Viceira (2003) find that the exponent on variance in the variance diffusion term is significantly different from $1/2$ (its value in the Heston model) and estimate its value to be slightly less than 1. However, the significance of this difference vanishes with the inclusion of jumps and thus the good performance of type (a) alternatives might be driven by model misspecification due to the excluded possibility of jumps. Ait-Sahalia and Kimmel (2007) also conclude that this exponent lies between $1/2$ and 1. Christoffersen *et al.* (2010) find that the GARCH diffusion stochastic volatility model also outperforms the Heston model in an option pricing framework. Alternative specifications including multi-factor volatility models are discussed in Chernov *et al.* (2003) or Fatone *et al.* (2011).

For our analysis, we select three representatives of the European equity index market, namely the Eurostoxx 50, DAX 30 and FTSE 100 indices. Eurostoxx 50 is a blue-chip index built from 50 leading European companies from twelve different Eurozone countries. DAX 30 consists of the 30 largest German enterprises as measured by order book volume and market capitalisation. The FTSE 100 includes the 100 most highly capitalised UK companies which are traded at the London Stock Exchange. Finally, we use the S&P 500 as a benchmark.

² Both these alternatives are related to popular discrete-time generalised autoregressive conditionally heteroscedastic (GARCH) models. A mean-reverting variance with diffusion coefficient proportional to variance can be regarded as the continuous limit of the symmetric GARCH process introduced by Bollerslev (1986). Similarly, the log volatility specification is a continuous-time counterpart of the discrete-time exponential GARCH process introduced by Nelson (1991).

For each index we test the specifications of twelve different continuous-time two factor models. The mean-reverting variance diffusion component can follow either the affine process of Heston (1993), the scale-invariant GARCH process of Nelson (1990), or the log volatility process of Scott (1987), and each may be augmented with price and volatility jump extensions. To the best of our knowledge no other paper has tested all three classes of diffusion and jump-diffusions against each other on a similar data set, and some of the specifications that we consider have not been studied in the literature before.

An important contribution of our paper is to fill a gap in the literature by examining the continuous time dynamics of European equity indices. Many indices in this sizable market have very actively traded futures, exchange traded funds, options and structured products such as volatility index futures, and therefore knowledge of their continuous-time dynamics is an extremely relevant research topic. Nevertheless, until now, this topic has been almost completely ignored in the literature, as the vast majority of empirical research focuses exclusively on the S&P 500. Our first goal is to see whether the ambiguous results that have been reported for US equity indices carry over to the European markets. In particular, we investigate whether a departure from the affine model class is necessary for European equity indices. For the S&P 500 index, some previous research favors non-affine specifications, but – to the best of our knowledge – there is no attempt in the literature to test similar specifications on the European equity market. Our research provides evidence that affine models with sufficiently rich jump specifications perform well for European equity indices and that similarly clear results are not apparent for the S&P500. Regarding the two alternatives to the square-root model class, our empirical results imply the superiority of GARCH alternatives which consistently outperform – especially for jump extensions – models of with a log volatility process. This finding is consistent across all indices we consider.

Our choice of models and indices allows us to address a range of additional questions concerning the jump behaviour of different equity indices. For instance, are jumps in volatility significant in the USA, but not in Europe? How important are jumps in non-affine specifications? Does the FTSE 100 index behave like the S&P 500 index, or is it more similar to the European indices? Regarding these research questions, we find that volatility jumps are far more important to add to the square-root model class than to alternative stochastic volatility models; especially the GARCH specifications can create realistic volatility dynamics without resorting to the inclusion of jump processes. Indeed, for our European indices but not for the S&P 500 a simple GARCH stochastic volatility process without jumps in either state variable already performs surprisingly well. By contrast, within the affine model class the inclusion of jumps for both state variables is essential for generating realistic dynamics. We also confirm that the three European indices have similar dynamics and these are different from the S&P 500 dynamics. The S&P 500 is definitely the most difficult index to model. Especially, modelling the skewness of returns in this market poses a very difficult challenge.

This paper further adds to the existing literature in two significant ways. We present very extensive simulation results for detecting model misspecifications which are required in order to discriminate between alternative models. We select numerous statistics from the observed equity index data and gauge the ability of alternative specifications to produce similar characteristics. Though computationally intensive, this approach provides more detailed evidence on the features of the data that a model fails to capture, and yields valuable insights regarding the adequacy of continuous time jump-diffusion models. Thus our results reach beyond the evidence currently presented

in the literature. Moreover, we employ a very large sample of stock index prices from 1987 to 2010 which includes the recent banking crisis of 2008–9. This period represents the most prolonged and excessively stressful equity markets ever experienced, and it is important that dynamic model specification tests encompass such a market regime.

We proceed as follows: Section 2 introduces the continuous-time models; Section 3 describes the data; Section 4 specifies the discrete-time counterparts for MCMC estimation; Section 5 presents the estimation results; Section 6 provides the specification tests; and Section 7 concludes.

2. Model Specification

We consider an equity index modelled by a jump-diffusion process that admits stochastic volatility and random jumps. In particular, we assume that the log index value $Y_t \equiv \log S_t$ evolves according to

$$dY_t = \mu dt + \sqrt{V_{t-}} dW_t^y + dJ_t^y,$$

where μ is the constant drift of the process and W_t^y denotes a standard Brownian motion.³ We allow the stock price variance V_t to evolve stochastically over time and sample paths for the stock price index can exhibit sudden jumps specified by the pure jump process J_t^y .

We study three different classes for the variance process, each having a mean-reverting property which prohibits variance to move too far from a long-term equilibrium value. Furthermore, we make the standard assumption that the correlation ρ between the Brownian motions driving the spot price and the variance process is constant, but need not be zero. This flexibility is important to model the well-known leverage effect.⁴

In the first class we model the variance V_t with a square-root process as in Heston (1993) and following Duffie *et al.* (2000) we extend this to accommodate jumps in variance as well as jumps in prices. Hence the general specification is

$$dV_t = \kappa(\theta - V_t) dt + \sigma \sqrt{V_{t-}} dW_t^v + dJ_t^v, \quad (\text{S})$$

where κ is the speed of mean reversion, θ determines the long-term variance level, σ is the volatility-of-variance parameter, W_t^v is a Brownian motion (which has a correlation of ρ with W_t^y) and J_t^v specifies the jump in the variance process.

Our second class is the continuous-time GARCH model of Nelson augmented with a non-zero price-variance correlation and the possibility of a jump component. Thus the general specification is

$$dV_t = \kappa(\theta - V_t) dt + \sigma V_{t-} dW_t^v + dJ_t^v, \quad (\text{G})$$

where the parameters κ , θ , σ and ρ have the same interpretation as in (S).

³ We use the shorthand notation V_{t-} for the left limit $V_{t-} = \lim_{s \uparrow t} V_s$. Furthermore, we could have included a variance risk premium into the drift term of the equity index, however for jump-diffusion models Eraker (2004) and Andersen *et al.* (2002) find no significant dependence of the drift of the process on its variance. Therefore, to keep the model as parsimonious as possible, we drop such any dependence on the variance from the drift specification.

⁴ New evidence regarding the origin of the leverage effect for the DAX is presented in Masset and Wallmeier (2010).

The third class specifies the evolution of the log of volatility as a Gaussian Ornstein-Uhlenbeck process, as in Scott (1987), but also augmented with the possibility of jumps. Denoting $v_t \equiv \log\sqrt{V_t}$, we have

$$dv_t = \kappa(\theta - v_t)dt + \sigma dW_t^v + dJ_t^v, \quad (\text{L})$$

where the parameters κ , θ , σ and ρ have a similar interpretation to above, but in relation to the log volatility rather than the variance.

Each class contains four different models depending on the assumptions on the jump distributions:⁵

1. Pure diffusion models where $dJ_t^y = 0$ and $dJ_t^v = 0$ for all t . We use the acronyms (S-SV), (G-SV) and (L-SV) respectively;
2. We include jumps in the log price process only, setting $dJ_t^v = 0$ for all t . Jump arrivals are driven by a Poisson process with intensity parameter λ_y . We assume the sizes of the jumps are normally distributed, independent over time and also independent of the Poisson process.⁶ Hence $dJ_t^y = \xi_t^y dN_t^y$, where N_t^y is a Poisson process and ξ_t^y is a normally distributed variable with mean μ_y and standard deviation σ_y . Here we use the acronyms (S-SVYJ), (G-SVYJ) and (L-SVYJ);⁷
3. These models have jumps in prices and volatility that occur simultaneously, so the same Poisson process N_t drives both jumps. We assume that their sizes are correlated, i.e. $dJ_t^y = \xi_t^y dN_t$ with normal jump size ($\xi_t^y \sim \mathcal{N}(\mu_y + \rho_J \xi_t^v, \sigma_y)$) and $dJ_t^v = \xi_t^v dN_t$ with exponentially distributed jump size ($\xi_t^v \sim \exp(\mu_v)$). Note that the parameter ρ_J determines whether the jump size in volatility influences the jump size in price. We refer to these models as (S-SVCJ), (G-SVCJ) and (L-SVCJ);
4. Finally, we allow independent jumps in both processes, i.e. $dJ_t^y = \xi_t^y dN_t^y$ where $\xi_t^y \sim \exp(\mu_y)$ and $dJ_t^v = \xi_t^v dN_t^v$ where $\xi_t^v \sim \mathcal{N}(\mu_y, \sigma_y)$. The acronyms for these models are (S-SVIJ), (G-SVIJ) and (L-SVIJ).⁸

Jump distributions for the volatility process are chosen so that they produce only upward jumps. This has the attractive feature that variance cannot jump to a negative value and the process stays positive throughout. For the log volatility model positivity of the process is not an issue and a jump distribution with support on the whole real axis could be chosen to model negative as well as positive jumps. Since sudden negative jumps in volatility appear to be of little empirical relevance, we use the exponential distribution for all models. This also facilitates the comparison of the models as they depend on the same distributional assumptions for jumps.

⁵ Our jump specifications coincide with those studied in Eraker *et al.* (2003) for the square-root variance process.

⁶ Although other distributions are possible for the jump in log prices, the vast majority of research focuses on the normal distribution.

⁷ Note that (S-SVYJ) is identical to the option pricing model derived in Bates (1996).

⁸ In an earlier draft of this paper we have also included results on a model with jumps in variance only. However, this model had similar performance as the simple (SV) model, and we omit results for brevity.

Table 1
Descriptive statistics of equity percentage log returns

This table reports descriptive statistics for the four equity indices (Eurostoxx 50, DAX 30, FTSE 100 and S&P 500) used in this study. The statistics are calculated on daily percentage log returns and the sample period is from January 1987 until April 2010.

	Eurostoxx 50	DAX 30	FTSE 100	S&P 500
Mean	0.020	0.025	0.021	0.026
Standard deviation	1.318	1.466	1.146	1.198
Skewness	-0.196	-0.301	-0.543	-1.397
Kurtosis	9.388	9.446	13.637	33.574
Largest negative return	-8.262	-13.706	-13.029	-22.900
Largest positive return	10.438	10.797	9.384	10.957

3. Data

We choose to estimate model parameters using daily return data from 1 January 1987 until 1 April 2010. This sample includes several interesting periods such as the global equity crash of 1987, the outbreak of two Gulf wars (1990–91 and 2003), the Asian currency crisis (1997), the LCTM bailout (1998), the dot-com bubble during the late 1990's and its subsequent bursting, the 9/11 terrorist attacks (2001) and most importantly the recent credit and banking crisis (2008–9). By estimating the models over a large sample including several crises we hope to distinguish well between alternative dynamics for the indices.

For all indices in this study we collect end-of-day quotes and compute percentage log returns (from henceforth just called returns). Visual inspection reveals that all indices possess similar characteristics, with common volatile periods mainly before and after the dot-com bubble and towards the end of the sample when the credit and banking crises affected economies all over the world. Descriptive statistics for the indices are reported in Table 1. Whereas all index returns exhibit strong deviations from normality, statistics are most extreme for the S&P 500 with the highest (absolute) skewness, the highest kurtosis and the largest outliers.

4. Econometric Specification

Estimation of the structural parameters and the latent state variables in the jump-diffusion models described above is a non-trivial econometric problem that may be addressed using Bayesian estimation procedures, and in particular we use a Markov-Chain-Monte-Carlo (MCMC) sampler for all models under consideration. MCMC methods for discrete-time stochastic volatility models were introduced by Jacquier *et al.* (1994) and have been subsequently applied in other contexts. For example, Eraker *et al.* (2003) use a MCMC sampler to estimate parameters of affine continuous-time jump-diffusion models for US equity indices and Li *et al.* (2008) extend their methodology to Levy jump models.⁹

⁹ Other estimation methodologies applied to affine and non-affine models include the efficient method of moments developed in Gallant and Tauchen (1996), which has been

Regarding the time discretisation of the continuous-time process, our algorithm is closely related to the ideas developed in Eraker *et al.* (2003) to which we refer for further details. Using a first-order Euler scheme, the log of the equity index value for all models under consideration can be written as

$$Y_{t+1} = Y_t + \mu \Delta + \sqrt{\Delta V_t} \varepsilon_{t+1}^y + \xi_{t+1}^y N_{t+1}^y, \tag{1}$$

where ε_t^y is a standard normal variate and Δ denotes the discretisation step. Changes in the Poisson process are discretised by a sequence of independent Bernoulli variates N_t^y , where the event $N_t^y = 1$ occurs with probability λ_y .¹⁰ The approximation of the volatility processes is analogous, for instance in the log model we obtain:

$$v_{t+1} = v_t + \kappa(\theta - v_t)\Delta + \sigma\sqrt{\Delta} \varepsilon_{t+1}^v + \xi_{t+1}^v N_{t+1}^v \tag{2}$$

where the jump part of the process is again approximated by a Bernoulli variate N_t^v , and ε_t^v is a second standard normal variable with $\text{Corr}(\varepsilon_t^y, \varepsilon_t^v) = \rho$. Throughout the remainder of this study we work with daily return data and set $\Delta \equiv 1$. Simulation experiments in Eraker *et al.* (2003) confirm that at this observation frequency the discretisation bias is negligible.

In Bayesian statistics, inference about unknown parameters and latent state variables is based on the distribution of all unknown quantities given the observed data $\mathbf{Y} = \{Y_t\}_{1:T}$, which is referred to as the *posterior density*. For instance for the log volatility models the posterior can be written as

$$p(\mathbf{v}, \boldsymbol{\xi}^y, N^y, \boldsymbol{\xi}^v, N^v, \boldsymbol{\Theta} \mid \mathbf{Y}) \propto p(\mathbf{Y}, \mathbf{v} \mid \boldsymbol{\xi}^y, N^y, \boldsymbol{\xi}^v, N^v, \boldsymbol{\Theta}) \times p(\boldsymbol{\xi}^y \mid N^y, \boldsymbol{\xi}^v, \boldsymbol{\Theta}) \\ \times p(\boldsymbol{\xi}^v \mid N^v, \boldsymbol{\Theta}) \times p(N^v \mid \boldsymbol{\Theta}) \times p(N^y \mid \boldsymbol{\Theta}) \times p(\boldsymbol{\Theta}),$$

where $\boldsymbol{\Theta} = \{\mu, \kappa, \theta, \sigma, \rho, \lambda_y, \mu_y, \sigma_y, \lambda_v, \mu_v\}$ is the unknown parameter vector, $p(\boldsymbol{\Theta})$ is the prior density that reflects any beliefs of the researcher regarding the unknown structural parameters and latent state variables are collected in vectors where the same notation applies as for \mathbf{Y} , for example $\boldsymbol{\xi}^y = \{\xi_t^y\}_{1:T}$. Eraker *et al.* (2003) point out that the likelihood function can be unbounded in a jump-diffusion framework and this complicates likelihood-based inference without prior information. On the other hand including subjective prior information yields results that are not universally applicable, and for this reason we choose priors that are identical or very similar to the uninformative priors in Eraker *et al.* (2003).

The dimension of the posterior density is several times the sample size and this complicates the direct analytical use of the posterior. We therefore apply the Gibbs sampler to reduce the dimensionality of the problem and to obtain information about the posterior density by simulation. Although this requires the derivation of complete conditional distributions this practice has become mainstream in the Bayesian literature. Using standard conjugate priors for most of the structural parameters these distributions are easy to derive. The only parameters that lead to non-standard densities are σ and ρ . For these two parameters we use the re-parametrization suggested in Jacquier *et al.* (2004) as

applied to continuous-time financial models in Andersen *et al.* (2002) and Chernov *et al.* (2003).

¹⁰ This is a slight abuse of notation because N_t^y was the Poisson process in the continuous-time process and represents the change in this process in the discrete-time version. To avoid introducing further variables, we follow the literature and use this slightly inconsistent notation.

it circumvents the implementation of Metropolis steps. In the sampling of the complete conditional distributions for the latent state variables the only complicated step arises for the variance vector. Since it is not possible to draw \mathbf{v} or V as a block we cycle through the variance vector one by one using the ARMS Metropolis algorithm of Gilks *et al.* (1995). Li *et al.* (2008) report that ARMS has superior updating performance compared with the simpler random walk Metropolis algorithm used in Eraker *et al.* (2003). To mitigate the effect of starting values and to insure that the chain has converged, we discard the first 30,000 runs of the sampler (which are commonly referred to as the ‘burn-in’) and summarise the posterior with the ensuing 100,000 draws.¹¹

5. Estimation Results

This section provides our estimation results. We first present MCMC estimates for the European indices and subsequently compare them with the S&P 500.

5.1. European stock indices

Results for the FTSE 100, DAX 30 and Eurostoxx 50 are presented in Tables 2–4. We begin with the interpretation of the estimated parameters in the square-root models. For all indices, our estimates for κ deviate only marginally from each other with values between 0.016 and 0.02. Given the standard error of 0.003 in all models there is no significant difference between the mean-reversion speeds of the indices. The other two variance parameters θ and σ show more substantial differences: in line with the observed standard deviation of the returns (Table 1) $\hat{\theta}$ is smallest for FTSE 100 (1.165), followed by Eurostoxx 50 (1.502) and DAX 30 (1.869), estimates that imply long-term volatility levels of 17% to 22%.¹² A similar comment applies to $\hat{\sigma}$ (0.14 for FTSE 100, 0.181 for Eurostoxx 50 and 0.205 for DAX 30) and hence Eurostoxx 50 and DAX 30 have the most erratic variance paths. The correlations between log price and variance innovations are very similar in all three indices with values around -50% . The estimated drift $\hat{\mu}$ is similar to the mean reported in Table 1.¹³

As expected, adding price jumps to the Heston model (S-SVYJ) mainly affects our parameter estimates for the vol-of-variance parameter as the inclusion of jumps reduces the need for the variance process to create large sudden movements. The characteristics of the jump part in the (S-SVYJ) are specific to each index. The lowest jump frequencies are estimated for the FTSE 100 where $\hat{\lambda}_y = 0.003$ gives about 0.75 jumps per year. Jumps in the DAX 30 and Eurostoxx 50 are more than twice as likely, with $\hat{\lambda}_y = 0.007$ and $\hat{\lambda}_y = 0.008$. The occurrence of jumps in the FTSE 100 index are not only less likely, they also have the smallest impact with an average jump size of -1.749% . The DAX 30 and Eurostoxx 50 have only slightly larger jump sizes (-1.962% and -2.717% respectively), but these estimates are statistically indistinguishable. The standard deviation of the jumps

¹¹ Models with independent jumps in returns and variance converge more slowly so we use 300,000 draws after burn-in for these.

¹² This is to be expected as the more diverse the stocks in index the lower its volatility, *ceteris paribus*.

¹³ To obtain the expected return of the process for the jump models, μ has to be adjusted by the estimated contribution of the jump part and thus this parameter is not directly comparable across models.

Table 2
MCMC estimates for the FTSE 100

This table reports the estimates of the structural parameters for all models introduced in Section 2 based on the mean of the posterior distributions. Standard deviations of the posterior are given in parenthesis. The parameter estimates correspond to daily log returns of the equity index values. One can easily obtain annual decimals by scaling some of the parameters. For example, assuming 252 trading days a year, in the square-root model class κ and λ have to be scaled by 252, σ by 2.52, $\sqrt{252\theta}/100$ provides the mean volatility and $\sqrt{252\eta_v}/100$ the mean jump in volatility. Similar scaling applies to the other model classes.

	SV		SVYJ		SVCJ		SVIJ	
<i>FTSE 100 – square-root models</i>								
μ	0.025	(0.011)	0.025	(0.011)	0.033	(0.011)	0.031	(0.011)
κ	0.016	(0.003)	0.014	(0.002)	0.029	(0.003)	0.026	(0.003)
θ	1.165	(0.12)	1.130	(0.124)	0.621	(0.059)	0.631	(0.058)
σ	0.140	(0.009)	0.125	(0.008)	0.097	(0.009)	0.096	(0.008)
ρ	-0.511	(0.045)	-0.549	(0.047)	-0.545	(0.056)	-0.567	(0.055)
λ_y			0.003	(0.001)	0.006	(0.002)	0.001	(0.001)
μ_y			-1.749	(1.721)	0.517	(0.84)	-10.662	(3.48)
ρ_J					-0.762	(0.2)		
σ_y			4.825	(1.443)	1.746	(0.373)	2.307	(1.635)
λ_v							0.006	(0.002)
η_v					2.778	(0.731)	2.723	(0.73)
<i>FTSE 100 – GARCH models</i>								
μ	0.028	(0.01)	0.037	(0.012)	0.029	(0.01)	0.035	(0.012)
κ	0.009	(0.003)	0.008	(0.003)	0.012	(0.003)	0.013	(0.003)
θ	1.342	(0.342)	1.348	(0.367)	0.859	(0.179)	0.802	(0.167)
σ	0.147	(0.01)	0.144	(0.01)	0.124	(0.009)	0.123	(0.011)
ρ	-0.542	(0.047)	-0.573	(0.047)	-0.618	(0.047)	-0.622	(0.048)
λ_y			0.023	(0.018)	0.004	(0.002)	0.013	(0.013)
μ_y			-0.536	(0.37)	0.973	(1.306)	-1.284	(1.548)
ρ_J					-2.239	(0.699)		
σ_y			1.133	(0.307)	1.972	(0.619)	2.203	(1.726)
λ_v							0.003	(0.001)
η_v					1.394	(0.637)	2.812	(1.396)
<i>FTSE 100 – log volatility models</i>								
μ	0.030	(0.01)	0.040	(0.014)	0.036	(0.011)	0.037	(0.012)
κ	0.015	(0.003)	0.014	(0.002)	0.015	(0.002)	0.015	(0.002)
$\exp(\theta)$	0.880	(0.057)	0.865	(0.06)	0.718	(0.062)	0.696	(0.067)
σ	0.073	(0.005)	0.072	(0.005)	0.059	(0.005)	0.060	(0.005)
ρ	-0.534	(0.05)	-0.570	(0.048)	-0.596	(0.057)	-0.629	(0.061)
λ_y			0.030	(0.027)	0.013	(0.007)	0.021	(0.02)
μ_y			-0.500	(0.382)	0.140	(0.513)	-0.202	(0.862)
ρ_J					-4.728	(0.932)		
σ_y			1.087	(0.339)	1.437	(0.335)	1.273	(0.41)
λ_v							0.016	(0.009)
η_v					0.253	(0.073)	0.247	(0.082)

Table 3
MCMC estimates for the DAX 30

This table reports the estimates of the structural parameters for all models introduced in Section 2 based on the mean of the posterior distributions. Standard deviations of the posterior are given in parenthesis. The parameter estimates correspond to daily log returns of the equity index values. One can easily obtain annual decimals by scaling some of the parameters. For example, assuming 252 trading days a year, in the square-root model class κ and λ have to be scaled by 252, σ by 2.52, $\sqrt{252\theta}/100$ provides the mean volatility and $\sqrt{252\eta_v}/100$ the mean jump in volatility. Similar scaling applies to the other model classes.

	SV		SVYJ		SVCJ		SVIJ	
<i>DAX 30 – square-root models</i>								
μ	0.044	(0.013)	0.046	(0.013)	0.051	(0.013)	0.055	(0.014)
κ	0.020	(0.003)	0.015	(0.003)	0.022	(0.003)	0.022	(0.003)
θ	1.869	(0.173)	1.799	(0.189)	1.038	(0.105)	0.960	(0.108)
σ	0.205	(0.012)	0.174	(0.012)	0.132	(0.011)	0.121	(0.011)
ρ	-0.505	(0.037)	-0.543	(0.041)	-0.581	(0.05)	-0.594	(0.052)
λ_y			0.007	(0.003)	0.006	(0.002)	0.010	(0.007)
μ_y			-1.962	(1.542)	-1.923	(1.564)	-1.103	(0.753)
ρ_J					-0.211	(0.271)		
σ_y			3.586	(0.955)	3.590	(0.622)	2.853	(0.824)
λ_v							0.006	(0.002)
η_v					3.805	(1.155)	4.096	(1.333)
<i>DAX 30 – GARCH models</i>								
μ	0.053	(0.013)	0.059	(0.013)	0.054	(0.013)	0.059	(0.014)
κ	0.011	(0.003)	0.008	(0.002)	0.013	(0.003)	0.014	(0.003)
θ	2.115	(0.421)	2.028	(0.447)	1.133	(0.238)	0.949	(0.201)
σ	0.178	(0.011)	0.150	(0.01)	0.132	(0.01)	0.127	(0.01)
ρ	-0.507	(0.04)	-0.559	(0.041)	-0.623	(0.047)	-0.672	(0.044)
λ_y			0.011	(0.006)	0.008	(0.003)	0.014	(0.008)
μ_y			-1.090	(0.652)	-0.117	(0.975)	-1.242	(0.649)
ρ_J					-1.342	(0.766)		
σ_y			2.776	(0.696)	2.833	(0.661)	2.234	(0.727)
λ_v							0.009	(0.004)
η_v					1.514	(0.643)	1.492	(0.482)
<i>DAX 30 – log volatility models</i>								
μ	0.048	(0.013)	0.053	(0.013)	0.055	(0.013)	0.059	(0.015)
κ	0.021	(0.003)	0.015	(0.003)	0.014	(0.002)	0.015	(0.002)
$\exp(\theta)$	1.111	(0.065)	1.091	(0.075)	0.897	(0.086)	0.813	(0.106)
σ	0.091	(0.006)	0.077	(0.005)	0.065	(0.005)	0.066	(0.006)
ρ	-0.493	(0.043)	-0.568	(0.045)	-0.604	(0.05)	-0.633	(0.054)
λ_y			0.013	(0.006)	0.012	(0.006)	0.020	(0.013)
μ_y			-0.981	(0.537)	-1.047	(0.833)	-0.983	(0.507)
ρ_J					-1.983	(2.329)		
σ_y			2.540	(0.641)	2.790	(0.645)	1.962	(0.724)
λ_v							0.024	(0.016)
η_v					0.251	(0.077)	0.223	(0.09)

Table 4
MCMC estimates for the Eurostoxx 50

This table reports the estimates of the structural parameters for all models introduced in Section 2 based on the mean of the posterior distributions. Standard deviations of the posterior are given in parenthesis. The parameter estimates correspond to daily log returns of the equity index values. One can easily obtain annual decimals by scaling some of the parameters. For example, assuming 252 trading days a year, in the square-root model class κ and λ have to be scaled by 252, σ by 2.52, $\sqrt{252\theta}/100$ provides the mean volatility and $\sqrt{252\eta_v}/100$ the mean jump in volatility. Similar scaling applies to the other model classes.

	SV		SVYJ		SVCJ		SVIJ	
<i>Eurostoxx 50 – square-root models</i>								
μ	0.039	(0.011)	0.044	(0.011)	0.049	(0.011)	0.054	(0.011)
κ	0.017	(0.003)	0.014	(0.002)	0.023	(0.003)	0.024	(0.003)
θ	1.502	(0.157)	1.473	(0.17)	0.741	(0.089)	0.700	(0.081)
σ	0.181	(0.009)	0.160	(0.01)	0.118	(0.01)	0.113	(0.01)
ρ	-0.481	(0.036)	-0.527	(0.039)	-0.533	(0.049)	-0.560	(0.051)
λ_y			0.008	(0.006)	0.007	(0.002)	0.012	(0.006)
μ_y			-2.717	(1.815)	-2.263	(1.028)	-1.228	(0.656)
ρ_J					-0.127	(0.253)		
σ_y			2.014	(0.548)	2.372	(0.568)	1.924	(0.421)
λ_v							0.006	(0.002)
η_v					2.945	(0.661)	3.486	(0.849)
<i>Eurostoxx 50 – GARCH models</i>								
μ	0.046	(0.011)	0.057	(0.011)	0.050	(0.011)	0.055	(0.011)
κ	0.009	(0.003)	0.007	(0.002)	0.010	(0.003)	0.012	(0.004)
θ	1.832	(0.456)	1.824	(0.496)	1.107	(0.356)	0.871	(0.265)
σ	0.186	(0.011)	0.161	(0.01)	0.150	(0.01)	0.143	(0.011)
ρ	-0.509	(0.04)	-0.572	(0.038)	-0.584	(0.045)	-0.649	(0.048)
λ_y			0.017	(0.008)	0.009	(0.004)	0.017	(0.007)
μ_y			-1.263	(0.49)	-1.548	(1.087)	-1.142	(0.41)
ρ_J					-1.346	(1.419)		
σ_y			1.682	(0.318)	1.769	(0.477)	1.613	(0.279)
λ_v							0.008	(0.004)
η_v					0.881	(0.422)	1.302	(0.529)
<i>Eurostoxx 50 – log volatility models</i>								
μ	0.043	(0.011)	0.054	(0.011)	0.056	(0.012)	0.054	(0.011)
κ	0.019	(0.003)	0.014	(0.002)	0.014	(0.002)	0.014	(0.002)
$\exp(\theta)$	0.957	(0.062)	0.940	(0.073)	0.731	(0.086)	0.686	(0.107)
σ	0.094	(0.006)	0.080	(0.005)	0.072	(0.006)	0.071	(0.006)
ρ	-0.490	(0.042)	-0.576	(0.043)	-0.566	(0.051)	-0.624	(0.049)
λ_y			0.018	(0.007)	0.019	(0.008)	0.019	(0.009)
μ_y			-1.172	(0.412)	-0.680	(0.521)	-1.143	(0.49)
ρ_J					-4.709	(1.698)		
σ_y			1.637	(0.285)	1.312	(0.313)	1.658	(0.29)
λ_v							0.030	(0.02)
η_v					0.191	(0.046)	0.167	(0.057)

in the FTSE 100 is the highest among all indices, at about 5%, yet the DAX 30 and Eurostoxx 50 have a lower jump standard deviation with 3.6% and 2.0% respectively. Although there is some variability in the point estimates of the jump size distribution across the indices, the fact that jumps are extremely rare events makes it very difficult to distinguish between the effect of jumps on the European indices. When models allow both state variables to jump, our estimates imply a variance jump between 2.7 (FTSE) and 4.1 (DAX). The differences, however, as also found for the price jumps, are not significant. Interestingly, the estimate for the jump correlation is only significant in the FTSE 100.

The parameter estimates for the GARCH models are reported in the middle section of Tables 2, 3 and 4. The estimate for σ in the pure diffusion model (G-SV) for all indices is similar to the parameter in the square-root models, but note this is not directly comparable with the parameter in (S-SV). Yet the other parameter estimates also deviate from their square-root counterparts: $\hat{\rho}$ is more negative; $\hat{\kappa}$ for most models is only about half the size of the estimate in (S-SV); and $\hat{\theta}$ also exhibits higher point estimates compared with the square-root specification. These differences are highly consistent across all the three indices and four different models, yet statistical significance is difficult to obtain as most parameters exhibit high standard errors.

There is also a striking difference between the jump parameter estimates in GARCH models, compared with the equivalent parameter estimates when jumps augment a square-root model: in GARCH models the jump occurrence is more frequent and their impact is much lower. Jump sizes are on average smaller with point estimates around zero (and also small standard deviations of around 2%), but they occur far more frequently than in the (S) specifications, although the significance of these differences is again low. A possible explanation for the more frequent but smaller jumps in GARCH specifications is as follows: because the variance diffusion in GARCH specifications can change more rapidly than in the square-root diffusion there is less pressure on the jump part to produce large positive and negative returns. With one exception, jumps in variance are also of considerably smaller magnitude than they are in the square-root process, with values typically between 1 and 2.

Figure 1 depicts the evolution of volatility in the (G-SVCJ) models for all three indices around the time of the crash of 1987 (left) and credit and banking crisis (right). The 1987 crash appears to come more as a complete surprise, as volatility in all indices jumps from levels around 15% to almost 60% in the space of a few days. The more recent crisis also leads to jumps in volatilities but the increase in variance is less sudden. It is interesting to note that the estimated variance paths for the three indices (indeed all four indices) stay extremely close during these crash events. It is well known that returns of equity indices become more highly correlated during volatile periods, and our results suggest that their volatilities might also be driven by a common factor.

The log volatility model parameter estimates are more difficult to compare with the other two classes as some of the parameters refer to the log volatility rather than the variance. The estimates for κ and ρ are similar to those for the square-root process. Consistent with our findings from the GARCH models, price jumps occur more often than in the (S) class, yet their impact is rather small.¹⁴ The estimate for σ in the GARCH diffusion is almost exactly half the size of its log volatility counterpart. This is

¹⁴ For example, the FTSE 100 (L-SVYJ) estimates imply jumps with mean -0.5% and slightly more than 1% standard deviation. Jumps in volatility are of similar magnitudes in

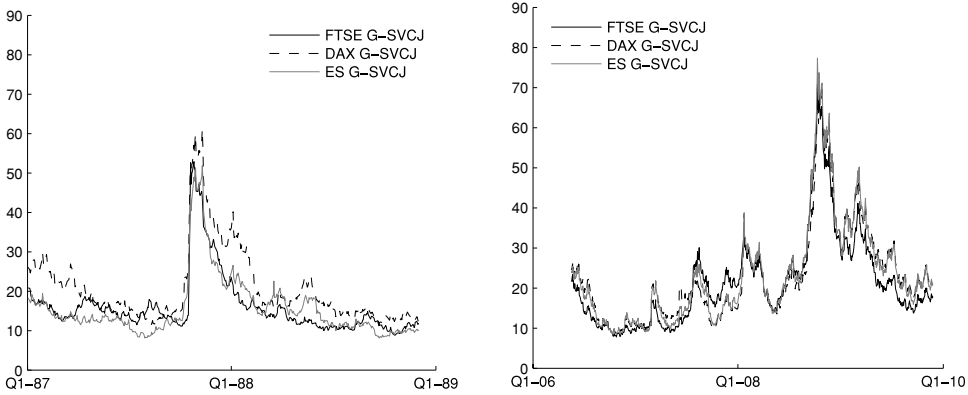


Fig. 1. Volatility for European GARCH models.

This figure depicts the estimated yearly volatility (in %, left axis) of FTSE 100 (abbreviated FTSE), DAX 30 (abbreviated DAX) and Eurostoxx 50 (abbreviated ES) around the market crash of 1987 and the recent credit and banking crisis of 2008–9; the volatility is derived from a GARCH diffusion model augmented with price and variance jumps (denoted G-SVCJ).

theoretically unsurprising, as an application of Ito’s Lemma for jump-diffusion models to (L) yields:

$$dV_t = 2 V_t[\kappa(\theta - \log\sqrt{V_t}) + \sigma^2] dt + 2\sigma V_{t-} dW_t^v + V_{t-}[\exp(2\xi_t^v) - 1] dN_t^v.$$

The diffusion part is hence expected to be similar, and the only difference between the GARCH and the log-volatility diffusion models springs from the drift specification. Another fundamental difference is that the importance of jump sizes in the GARCH model fades away with increasing volatility because jump sizes are independent of the variance level. In log-volatility models, jumps are relative to the level of the volatility.

Table 5 compares the in-sample fit of the competing models. Here we report the DIC (deviance information criterion) developed as a generalisation of the Akaike information criterion (AIC), which provides our first indication of the relative performance of alternative specifications. Note that a smaller DIC value is preferred.¹⁵ As a caveat, in this context Bayesian fit statistics are not as developed as they are in frequentist econometrics. Hence we provide more detailed results on model selection in Section 6. The DIC fit statistics for DAX 30, FTSE 100 and Eurostoxx 50 are presented in Table 5. The GARCH model with correlated jumps in price and variance outperforms all other model specifications, for all three indices. Whatever the diffusion specification it is important to include jumps and contemporaneous price and volatility jumps provide the best fit. As noted before, we shed more light to this question in subsequent sections. Note that (G-SVYJ) outperforms (S-SVCJ) for the FTSE 100, so whether a jump in volatility is needed is not clear at this stage.

(L-SVCJ) for all indices and these estimates change only marginally, under the (L-SVIJ) only.

¹⁵ DIC adjusts for the complexity (the effective number of parameters) of the model and thus allows one to compare nested and non-nested models.

Table 5
Model fit for FTSE 100, DAX 30 and Eurostoxx 50

Entries in this table are the estimates of the DIC in-sample fit statistic for the volatility specification indicated by the row and the jump augmentation specified by the column. Lower values of the DIC statistic indicate a superior fit by the model. As usual results are presented in three separate groups, according to the equity index being modelled. Results may be compared within a group but not across groups; e.g. considering the square-root model without jumps, the DIC for FTSE (15,369) is lower than the DIC for DAX (18,223) but this does not indicate that the model fits to FTSE better than the DAX. However, the DIC for the GARCH model with correlated jumps when estimated on the FTSE is 14,009; this indicates that the GARCH model with correlated jumps fits the FTSE better than the square-root model without jumps.

	SV	SVYJ	SVCJ	SVIJ
<i>FTSE 100</i>				
Square-root model	15369	15117	14964	15124
GARCH model	15212	14740	14009	15009
Log volatility model	15323	14772	14365	15090
<i>DAX 30</i>				
Square-root model	18223	17856	17371	18133
GARCH model	18211	17712	16930	17850
Log volatility model	18343	17710	17231	18130
<i>Eurostoxx 50</i>				
Square-root model	16514	16252	15755	16472
GARCH model	16419	15854	15298	16187
Log volatility model	16585	15885	15560	16553

5.2. S&P 500

In this section we briefly discuss our results for the S&P 500. As this index has been subject to intensive empirical research we only provide a short outline of our empirical results and use these mainly to benchmark our findings in the next section. Note that Andersen *et al.* (2002), Eraker *et al.* (2003), Eraker (2004) or Li *et al.* (2008) also provide estimations of some of the proposed model specifications.

Our parameter estimates in Table 6 are in line with previous research for the square root model class, although point estimates differ due to our extended sample covering the recent crisis. Without jumps $\hat{\theta} = 1.254$ implies an annual long-term volatility level of 17.8% which is slightly higher than the estimate found by Eraker *et al.* and Andersen *et al.* so the addition of data from 2000 to 2010 has a clear impact. This is also true of the other parameter estimates. In particular $\hat{\sigma} = 0.171$ exceeds the values in Eraker *et al.* (0.1434), Eraker (0.108) and Andersen *et al.* (0.0771). Furthermore, it is well known that the correlation between returns and variance is more pronounced during periods of crisis and $\hat{\rho} = -0.598$ (compared with -0.3974 , -0.373 and -0.3799 in the three previous studies) confirms this. However the mean reversion estimate $\hat{\kappa} = 0.019$ is similar to those found in previous research. Also our estimates for the jump parameters are comparable with the results in Eraker

Table 6
MCMC estimates for the S&P 500

This table reports the estimates of the structural parameters for all models introduced in Section 2 based on the mean of the posterior distributions. Standard deviations of the posterior are given in parenthesis. The parameter estimates correspond to daily log returns of the equity index values. One can easily obtain annual decimals by scaling some of the parameters. For example, assuming 252 trading days a year, in the square-root model class κ and λ have to be scaled by 252, σ by 2.52, $\sqrt{252\theta}/100$ provides the mean volatility and $\sqrt{252\eta_v}/100$ the mean jump in volatility. Similar scaling applies to the other model classes.

	SV		SVYJ		SVCJ		SVIJ	
<i>S&P 500 – square-root models</i>								
μ	0.025	(0.01)	0.028	(0.01)	0.031	(0.01)	0.034	(0.01)
κ	0.019	(0.003)	0.014	(0.002)	0.020	(0.003)	0.022	(0.004)
θ	1.254	(0.128)	1.243	(0.145)	0.850	(0.1)	0.856	(0.092)
σ	0.171	(0.01)	0.146	(0.009)	0.135	(0.01)	0.136	(0.009)
ρ	-0.598	(0.035)	-0.666	(0.032)	-0.676	(0.034)	-0.674	(0.034)
λ_y			0.005	(0.002)	0.006	(0.002)	0.009	(0.005)
μ_y			-3.215	(1.299)	-1.810	(0.964)	-1.822	(0.79)
ρ_J					-1.656	(0.673)		
σ_y			4.004	(1.069)	2.040	(0.81)	1.864	(0.575)
λ_v							0.002	(0.001)
η_v					1.621	(0.554)	7.114	(5.413)
<i>S&P 500 – GARCH models</i>								
μ	0.034	(0.01)	0.043	(0.01)	0.038	(0.01)	0.039	(0.01)
κ	0.008	(0.003)	0.006	(0.002)	0.008	(0.003)	0.010	(0.004)
θ	1.692	(0.47)	1.769	(0.513)	1.114	(0.389)	1.004	(0.387)
σ	0.181	(0.011)	0.165	(0.01)	0.153	(0.01)	0.163	(0.011)
ρ	-0.627	(0.035)	-0.690	(0.032)	-0.721	(0.035)	-0.755	(0.042)
λ_y			0.014	(0.008)	0.009	(0.004)	0.013	(0.007)
μ_y			-1.405	(0.642)	-0.591	(0.776)	-1.490	(0.635)
ρ_J					-3.197	(0.767)		
σ_y			1.702	(0.443)	1.402	(0.322)	1.619	(0.39)
λ_v							0.011	(0.007)
η_v					0.707	(0.255)	0.692	(0.407)
<i>S&P 500 – log volatility models</i>								
μ	0.036	(0.01)	0.043	(0.01)	0.044	(0.01)	0.044	(0.01)
κ	0.019	(0.003)	0.015	(0.002)	0.015	(0.002)	0.015	(0.002)
$\exp(\theta)$	0.859	(0.055)	0.848	(0.063)	0.698	(0.07)	0.688	(0.093)
σ	0.091	(0.006)	0.084	(0.005)	0.074	(0.006)	0.077	(0.006)
ρ	-0.592	(0.04)	-0.680	(0.038)	-0.700	(0.037)	-0.718	(0.04)
λ_y			0.017	(0.009)	0.015	(0.007)	0.018	(0.011)
μ_y			-1.264	(0.522)	-1.158	(0.656)	-1.231	(0.539)
ρ_J					-2.357	(1.598)		
σ_y			1.562	(0.358)	1.465	(0.294)	1.484	(0.285)
λ_v							0.026	(0.022)
η_v					0.229	(0.058)	0.166	(0.081)

Table 7
Model fit for the S&P 500

Entries in this table are the estimates of the DIC in-sample fit statistic for the volatility specification indicated by the row and the jump augmentation specified by the column. Lower values of the DIC statistic indicate a superior fit by the model.

	SV	SVYJ	SVCJ	SVIJ
Square-root model	14738	14126	13799	14565
GARCH model	14590	13790	12928	14197
Log volatility model	14901	13939	13518	14697

et al. but those in Eraker imply fewer jumps with greater impact (although our estimates are not significantly different). Compared to the existing literature, we obtain a considerably larger variance jump size in the (S-SVIJ) model (our estimate is 7.114, whereas Eraker *et al.* (2003) find 1.798), where our estimate would cause very large but rare volatility jumps. Yet again, the standard deviation of this estimate is too high for differences to be statistically significant.

For the S&P 500 index the overall best performing model in each class is the (SVCJ) (see Table 7). Price jumps lead to an improvement in the fit but the independent jump models tend to overfit the data and these underperform all other jump models. Among the three volatility specifications we find, consistent with our findings for the European indices, that GARCH models perform best, with substantially lower DIC values. Note that the DIC values for (G-SVYJ) are even better than those for (S-SVCJ). Therefore, when the restriction that the model be affine is dropped a more parsimonious specification without volatility jumps might suffice.

6. Specification Tests and Model Comparison

This section provides specification tests for all competing model classes. First we provide an analysis of the residual errors and then we present extensive simulation results.

6.1. Residual error analysis

The estimated residuals ε_t^y and ε_t^v (as in equations (1) and (2)) should follow standard normal distributions,¹⁶ so any systematic deviation from normality indicates model misspecification. We test for normality by applying a standard Bayesian procedure. In every (after-burnin) run of the Markov chain we calculate the skewness and the kurtosis of the residual vector for log returns and variances (or log volatilities). These estimates allow one to obtain a distribution for the skewness and kurtosis of the log return and variance (or log volatility) equation errors, for every model and every index. We report the mean of these distributions as point estimates for the skewness and kurtosis and the 1 and 99% posterior intervals to obtain a probabilistic statement of the range of values for residual skewness and kurtosis generated by each model. Misspecified models will

¹⁶ Whereas this distributional assumption holds exactly in the discretised model, it holds only approximately for the continuous-time processes.

produce skewness and kurtosis statistics significantly different from 0 and 3 respectively. Results for the residuals of the log return equation for all indices are reported in Table 8.¹⁷

Considering the results for the European indices, none of the models with price jumps produce kurtosis statistics that are significantly different from 3 at the 1% level. Lower kurtosis levels are found in GARCH and log volatility specifications. These results confirm that the square-root process requires a jump in variance as well as price, whereas the other two volatility specifications are fine with just a price jump. The results regarding the skewness are however less encouraging. Given there is negative skewness even at the 0.01 percentile for all indices, we conclude that all models produce a significantly longer left tail than they should.

There is a stronger misspecification in all the models for the S&P 500 index, especially for models without jumps, and especially in the square-root class. The kurtosis is significantly greater than 3 in all models, so the residual vector also contains more extreme outliers than the normal distribution can produce. The skewness is also still significantly different from 0 for all models. Therefore, the dynamics of the European equity indices are easier to capture with the proposed models.

In order to quantify whether our results are robust to changing the sampling frequency of the data, we also re-estimated all models for all indices on weekly return observations.¹⁸ The conclusions drawn from this set of estimations are similar to the ones presented above. In particular, the skewness of the residuals still poses a severe challenge for the models. For brevity we do not detail the empirical results here, but they are available from the authors upon request.

6.2. Simulation study

If a model is a realistic description of the evolution of an equity index then repeated simulations should produce trajectories with characteristics similar to those of the observed time series. So in this sub-section we test whether the competing models could have produced the observed data. For instance, the DAX 30 sample kurtosis is 9.4 and if a model can capture this feature we would expect each simulated paths to exhibit a similar level of kurtosis. That is, 9.4 should not be located in the far tails of the model's kurtosis distribution. This idea is formalised by the concept of *posterior predictive p-values* introduced by Rubin (1984).

Consider the distribution of a statistic S under model M after observing the data \mathbf{Y} . This distribution for the statistic S is given by

$$p(S | \mathbf{Y}, M) = \int p(S | \Theta, M) p(\Theta | \mathbf{Y}, M) d\Theta, \quad (3)$$

where Θ is a general notation for the parameter vector of the model. The predictive p -value locates the observed $S(\mathbf{Y})$ in this distribution and high (close to one) or low (close to zero) p -values indicate that the model is not capable of producing the magnitudes of S that were observed in the actual data.

¹⁷ The corresponding statistics for the variance vector carry little useful information to distinguish between the competing models and thus here we only report and interpret results for the log return residuals. The results are available from the authors on request.

¹⁸ We thank an anonymous referee for suggesting this.

Table 8
Specification tests

Entries in the table summarise the distribution of skewness and kurtosis in the residuals ε_t^y (see equation (1)). If the model is well-specified, the skewness should be insignificantly different from zero, and the kurtosis should be insignificantly different from 3. The point estimates of these statistics are provided by the mean of the distribution, and the 1% and 99% percentiles indicate how variable the skewness and kurtosis estimates were about this point estimate.

		SV		SVYJ		SVCJ		SVIJ	
		Skew	Kurt	Skew	Kurt	Skew	Kurt	Skew	Kurt
<i>FTSE 100</i>									
Square-root models	mean	-0.170	3.293	-0.148	3.072	-0.120	3.027	-0.114	3.083
	1% percentile	-0.219	3.121	-0.196	2.961	-0.161	2.928	-0.156	2.947
	99% percentile	-0.123	3.525	-0.098	3.198	-0.070	3.172	-0.072	3.225
GARCH models	mean	-0.156	3.121	-0.124	3.028	-0.139	3.016	-0.124	2.996
	1% percentile	-0.203	2.991	-0.183	2.892	-0.185	2.907	-0.182	2.865
	99% percentile	-0.109	3.276	-0.062	3.188	-0.088	3.153	-0.055	3.139
Log volatility models	mean	-0.155	3.136	-0.120	3.033	-0.110	2.973	-0.112	2.979
	1% percentile	-0.203	3.007	-0.190	2.892	-0.163	2.858	-0.177	2.854
	99% percentile	-0.107	3.299	-0.054	3.196	-0.052	3.120	-0.046	3.127
<i>DAX 30</i>									
Square-root models	mean	-0.190	3.452	-0.115	3.109	-0.103	3.078	-0.097	3.038
	1% percentile	-0.241	3.252	-0.168	2.975	-0.149	2.971	-0.149	2.923
	99% percentile	-0.141	3.709	-0.059	3.263	-0.059	3.202	-0.044	3.178
GARCH models	mean	-0.168	3.254	-0.111	3.035	-0.116	3.034	-0.104	3.022
	1% percentile	-0.221	3.085	-0.166	2.912	-0.164	2.929	-0.158	2.909
	99% percentile	-0.116	3.480	-0.056	3.178	-0.068	3.158	-0.050	3.164
Log volatility models	mean	-0.167	3.245	-0.113	3.036	-0.101	3.004	-0.100	3.014
	1% percentile	-0.222	3.081	-0.168	2.915	-0.152	2.897	-0.156	2.898
	99% percentile	-0.114	3.463	-0.056	3.178	-0.048	3.135	-0.041	3.160
<i>Eurostoxx 50</i>									
Square-root models	mean	-0.202	3.446	-0.111	3.129	-0.102	3.100	-0.085	3.053
	1% percentile	-0.253	3.253	-0.167	2.992	-0.156	2.977	-0.142	2.927
	99% percentile	-0.152	3.696	-0.055	3.284	-0.052	3.236	-0.028	3.197
GARCH models	mean	-0.193	3.289	-0.103	3.045	-0.111	3.049	-0.099	3.030
	1% percentile	-0.247	3.114	-0.158	2.919	-0.164	2.933	-0.155	2.911
	99% percentile	-0.140	3.526	-0.047	3.187	-0.059	3.184	-0.043	3.164
Log volatility models	mean	-0.192	3.292	-0.102	3.047	-0.084	3.018	-0.098	3.027
	1% percentile	-0.248	3.117	-0.159	2.923	-0.142	2.899	-0.153	2.905
	99% percentile	-0.138	3.525	-0.047	3.193	-0.026	3.148	-0.042	3.176
<i>S&P 500</i>									
Square-root models	mean	-0.250	4.018	-0.089	3.239	-0.080	3.187	-0.080	3.253
	1% percentile	-0.313	3.697	-0.140	3.118	-0.127	3.082	-0.136	3.111
	99% percentile	-0.190	4.437	-0.038	3.382	-0.030	3.308	-0.026	3.444
GARCH models	mean	-0.203	3.544	-0.099	3.262	-0.085	3.170	-0.102	3.254
	1% percentile	-0.258	3.357	-0.157	3.120	-0.135	3.070	-0.157	3.111
	99% percentile	-0.150	3.789	-0.042	3.472	-0.034	3.283	-0.046	3.455
Log volatility models	mean	-0.189	3.539	-0.092	3.284	-0.070	3.196	-0.078	3.214
	1% percentile	-0.249	3.350	-0.152	3.136	-0.125	3.086	-0.138	3.096
	99% percentile	-0.132	3.787	-0.033	3.495	-0.013	3.329	-0.020	3.367

Note that the calculation of predictive p -values is easy to implement once a MCMC sampler has been derived, since we can approximate the integral with the outcome of the MCMC runs and simulations of the data-generating process. Furthermore, this approach also takes into account the uncertainty in the estimated parameters and hence accounts for estimation risk.

For the calculation of the integral above, we can further condition on so-called auxiliary statistics. In the repeated experiments these statistics are kept constant. In our case we only fix the sample size such that it coincides with the sample size of the observed time series and start the simulations at the long-term volatility level implied by the model. We use 100,000 simulated paths for each of the 48 (model, index) pairs.

The selection of relevant statistics S is crucial to the problem at hand, as their careful choice will affect whether inconsistencies in the models are detected. The statistics that we deem important for modeling equity indices and that can potentially help to distinguish between models are:

- The sample statistics from Table 1 except the unconditional mean, i.e. standard deviation (*stdev*), skewness (*skew*) and kurtosis (*kurt*), and the minimum (*min*) and maximum (*max*) of the returns;
- Further statistics linked to extreme behaviour, i.e. the average over the 10 largest positive jumps (*avgmax10*) and the average over the 10 largest negative jumps (*avgmin10*);
- Indications of outlier clustering: we record the highest and lowest sum of absolute returns (*absmax20* and *absmin20*) observed in a period of 20 trading days;
- Percentiles of the estimated unconditional distribution of the index returns, *percNUM* where NUM indicates the percent.

Results for the European indices in Tables 9 and 10 are encouraging. We start by interpreting the results for the FTSE. The Heston (S-SV) model is clearly misspecified, as shown by many of the statistics. In particular neither the skewness nor the kurtosis can be replicated (confirming our results from the previous section). The only moment that can be reproduced by the simulations is the standard deviation (with a p -value of 0.818), hence although simulations imply lower standard deviation values on average, we cannot reject the model using this statistic. The pure SV model does not capture large jumps in price: it is not surprising that the high negative jumps are impossible for the (S-SV) to generate, yet it also fails to produce jumps of considerable positive size. The inclusion of return jumps (S-SVYJ) into the Heston model improve the p -values for most statistics, but the model is still rejected. Jumps in variance are required for the FTSE in the square-root model class, where none of the statistics indicate significant model misspecification at the 5% level. Note that also the skewness of the returns is well captured by these models, although the residual error analysis pointed towards some weaknesses of the proposed models to capture this feature.

The GARCH model simulations convey a very different picture. Even the pure diffusion model can handle large returns much better than its square-root counterpart. The model creates realistic values for high jumps and both positive and negative jumps are frequent enough. Extending the (G-SV) model by jumps in state variables has surprisingly little effect on the results. The p -values for all statistics stay extremely close to each other with no discernible improvement from the inclusion of price jumps. The most difficult characteristic to capture is the skewness, but this is the only characteristic where the GARCH models without jumps in the variance remain unsuccessful. As for

Table 9
Simulation study: S&P 500 and FTSE 100 indices

This table reports a selection of statistics designed to test the specification of each model. The idea is to locate the sample statistic derived directly from the data in the simulated distribution of the statistic. A value of 0.80 for example indicates that in 80% of 100,000 simulation runs, the simulated statistic was smaller than the observed statistic. Values close to zero and one therefore indicate model misspecification. For each index the three blocks of four columns refer to the base diffusion model and the specific column refers to the jump augmentation specification. Abbreviations of the statistics are as follows: *stddev*: standard deviation, *skew*: skewness, *kurt*: kurtosis, *perc#*: #% percentile, *min*: minimum, *max*: maximum, *avgmax10* (*avgmin10*) is the average of the ten largest positive (negative) returns, *absmax20* (*absmin20*) reports the highest (lowest) sum of absolute changes that was observed over a 20 trading day period during the sample.

Data	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	<i>FTSE 100 – log volatility models</i>			
													FTSE 100 – square-root models	FTSE 100 – GARCH models	FTSE 100 – log volatility models	FTSE 100 – log volatility models
stdev	1.146	0.818	0.752	0.628	0.536	0.679	0.664	0.566	0.547	0.820	0.796	0.999	0.999	0.999	0.999	
skew	-0.543	0.000	0.168	0.102	0.559	0.021	0.021	0.208	0.064	0.002	0.003	0.140	0.001	0.001	0.001	
kurt	13.637	1.000	0.862	0.938	0.541	0.868	0.861	0.776	0.854	0.996	0.995	0.979	1.000	1.000	1.000	
avgmax10	6.884	1.000	0.975	0.940	0.934	0.726	0.716	0.774	0.668	0.972	0.966	1.000	1.000	1.000	1.000	
avgmin10	-7.902	0.000	0.087	0.121	0.475	0.155	0.163	0.272	0.204	0.008	0.011	0.001	0.000	0.000	0.000	
perc1	-3.206	0.042	0.055	0.313	0.249	0.276	0.303	0.382	0.420	0.111	0.145	0.001	0.001	0.001	0.001	
perc5	-1.672	0.656	0.617	0.570	0.560	0.393	0.432	0.523	0.529	0.429	0.486	0.009	0.009	0.009	0.009	
perc10	-1.181	0.775	0.734	0.625	0.628	0.409	0.428	0.534	0.514	0.475	0.511	0.017	0.015	0.015	0.015	
perc90	1.186	0.069	0.106	0.167	0.159	0.410	0.432	0.297	0.324	0.316	0.336	0.977	0.981	0.977	0.981	
perc95	1.650	0.129	0.182	0.245	0.234	0.444	0.453	0.331	0.336	0.358	0.368	0.987	0.989	0.987	0.989	
perc99	2.895	0.656	0.688	0.410	0.412	0.483	0.478	0.393	0.332	0.584	0.575	0.997	0.997	0.997	0.997	
absmax20	86.617	1.000	1.000	0.946	0.976	0.776	0.769	0.828	0.773	0.977	0.972	1.000	1.000	1.000	1.000	
absmin20	5.723	0.994	0.992	0.932	0.944	0.924	0.930	0.900	0.887	0.989	0.988	0.996	0.998	0.996	0.998	
max	9.384	0.999	0.799	0.889	0.874	0.682	0.674	0.737	0.631	0.931	0.924	0.998	0.998	0.998	0.998	
min	-13.029	0.000	0.189	0.125	0.558	0.102	0.109	0.283	0.136	0.005	0.007	0.010	0.000	0.010	0.000	

Table 9 (Continued)

Data	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	
	<i>Eurostoxx 50 – square-root models</i>						<i>Eurostoxx 50 – GARCH models</i>						<i>Eurostoxx 50 – log volatility models</i>				
stdev	0.855	0.810	0.598	0.649	0.706	0.705	0.602	0.689	0.822	0.773	0.996	0.997	0.822	0.773	0.996	0.997	
skew	0.007	0.296	0.506	0.193	0.150	0.253	0.394	0.196	0.086	0.223	0.799	0.620	0.086	0.223	0.799	0.620	
kurt	1.000	0.992	0.933	0.941	0.451	0.509	0.565	0.698	0.889	0.886	0.856	0.893	0.889	0.886	0.856	0.893	
avgmax10	1.000	1.000	0.933	0.882	0.564	0.609	0.625	0.705	0.895	0.884	0.997	0.997	0.895	0.884	0.997	0.997	
avgmin10	0.001	0.052	0.309	0.174	0.414	0.385	0.464	0.305	0.127	0.151	0.015	0.006	0.127	0.151	0.015	0.006	
perc1	0.005	0.015	0.167	0.128	0.180	0.220	0.295	0.198	0.043	0.092	0.002	0.001	0.043	0.092	0.002	0.001	
perc5	0.312	0.366	0.441	0.373	0.134	0.191	0.336	0.279	0.186	0.282	0.006	0.005	0.186	0.282	0.006	0.005	
perc10	0.514	0.533	0.508	0.417	0.102	0.140	0.299	0.249	0.230	0.299	0.008	0.006	0.230	0.299	0.008	0.006	
perc90	0.051	0.077	0.096	0.137	0.578	0.624	0.371	0.416	0.293	0.350	0.968	0.976	0.293	0.350	0.968	0.976	
perc95	0.164	0.213	0.201	0.245	0.624	0.649	0.433	0.477	0.399	0.441	0.981	0.985	0.399	0.441	0.981	0.985	
perc99	0.896	0.898	0.628	0.583	0.630	0.650	0.561	0.617	0.751	0.742	0.996	0.997	0.751	0.742	0.996	0.997	
absmax20	1.000	0.999	0.940	0.913	0.551	0.601	0.629	0.706	0.875	0.865	0.990	0.992	0.875	0.865	0.990	0.992	
absmin20	0.790	0.816	0.587	0.514	0.157	0.227	0.245	0.200	0.797	0.780	0.924	0.945	0.797	0.780	0.924	0.945	
max	0.996	0.996	0.898	0.852	0.513	0.566	0.593	0.664	0.825	0.827	0.984	0.989	0.825	0.827	0.984	0.989	
min	0.063	0.352	0.706	0.544	0.654	0.614	0.741	0.562	0.438	0.448	0.217	0.074	0.438	0.448	0.217	0.074	

Table 10
Simulation study: Eurostoxx 50 and DAX 30 indices

This table reports a selection of statistics designed to test the specification of each model. The idea is to locate the sample statistic derived directly from the data in the simulated distribution of the statistic. A value of 0.80 for example indicates that in 80% of 100,000 simulation runs, the simulated statistic was smaller than the observed statistic. Values close to zero and one therefore indicate model misspecification. For each index the three blocks of four columns refer to the base diffusion model and the specific column refers to the jump augmentation specification. Abbreviations of the statistics are as follows: *stadev*: standard deviation, *skew*: skewness, *kurt*: kurtosis, *perc#*: #% percentile, *min*: minimum, *max*: maximum, *avgmax10* (*avgmin10*) is the average of the ten largest positive (negative) returns, *absmax20* (*absmin20*) reports the highest (lowest) sum of absolute changes that was observed over a 20 trading day period during the sample.

Data	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	
	DAX 30 – square-root models				DAX 30 – GARCH models				DAX 30 – log volatility models				
<i>stadev</i>	1.466	0.870	0.827	0.617	0.661	0.723	0.730	0.643	0.742	0.812	0.783	0.997	0.999
<i>skew</i>	-0.301	0.000	0.189	0.357	0.074	0.081	0.144	0.310	0.082	0.023	0.074	0.520	0.249
<i>kurt</i>	9.446	1.000	0.946	0.848	0.944	0.584	0.614	0.657	0.840	0.950	0.938	0.832	0.927
<i>avgmax10</i>	8.103	1.000	0.995	0.871	0.833	0.576	0.634	0.697	0.772	0.892	0.893	0.998	0.999
<i>avgmin10</i>	-8.498	0.000	0.091	0.374	0.135	0.322	0.310	0.412	0.178	0.065	0.089	0.037	0.002
<i>perc1</i>	-4.505	0.002	0.006	0.121	0.096	0.143	0.175	0.205	0.124	0.029	0.057	0.001	0.000
<i>perc5</i>	-2.284	0.262	0.270	0.338	0.313	0.127	0.159	0.272	0.221	0.187	0.232	0.004	0.001
<i>perc10</i>	-1.567	0.507	0.479	0.454	0.408	0.123	0.142	0.284	0.233	0.277	0.299	0.008	0.003
<i>perc90</i>	1.533	0.105	0.175	0.207	0.237	0.629	0.670	0.453	0.518	0.345	0.414	0.976	0.992
<i>perc95</i>	2.130	0.167	0.250	0.252	0.271	0.590	0.628	0.433	0.490	0.349	0.414	0.980	0.994
<i>perc99</i>	3.835	0.816	0.829	0.538	0.496	0.574	0.604	0.515	0.584	0.627	0.647	0.995	0.998
<i>absmax20</i>	87.833	0.999	0.999	0.824	0.822	0.506	0.583	0.651	0.723	0.827	0.838	0.989	0.994
<i>absmin20</i>	4.631	0.454	0.545	0.284	0.251	0.028	0.046	0.042	0.048	0.488	0.482	0.589	0.719
<i>max</i>	10.797	0.992	0.939	0.804	0.779	0.503	0.572	0.643	0.711	0.803	0.815	0.975	0.987
<i>min</i>	-13.706	0.000	0.091	0.182	0.050	0.234	0.200	0.254	0.093	0.039	0.041	0.028	0.003

Table 10 (Continued)

Data	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SVCJ	SVIJ	SV	SVYJ	SV	SVYJ	SVCJ	SVIJ
	S&P 500 – square-root models				S&P 500 – GARCH models				S&P 500 – log volatility models					
stdev	1.198	0.842	0.662	0.561	0.606	0.640	0.609	0.389	0.560	0.875	0.787	0.997	0.997	0.991
skew	-1.397	0.000	0.059	0.088	0.001	0.008	0.006	0.017	0.004	0.000	0.000	0.001	0.001	0.002
kurt	33.574	1.000	0.984	0.978	0.996	0.921	0.929	0.954	0.959	0.999	0.999	1.000	1.000	1.000
avgmax10	7.287	1.000	0.995	0.996	0.762	0.566	0.574	0.536	0.609	0.959	0.931	0.999	0.997	0.997
avgmin10	-9.560	0.000	0.096	0.186	0.053	0.182	0.189	0.272	0.147	0.003	0.008	0.000	0.000	0.000
perc1	-3.173	0.158	0.340	0.528	0.532	0.440	0.545	0.764	0.613	0.207	0.409	0.022	0.022	0.037
perc5	-1.766	0.566	0.616	0.662	0.544	0.281	0.367	0.630	0.498	0.259	0.406	0.016	0.016	0.035
perc10	-1.177	0.855	0.848	0.867	0.755	0.349	0.402	0.687	0.571	0.442	0.529	0.042	0.042	0.072
perc90	1.201	0.049	0.072	0.065	0.162	0.522	0.534	0.232	0.351	0.376	0.390	0.966	0.966	0.932
perc95	1.685	0.091	0.120	0.114	0.225	0.476	0.480	0.210	0.316	0.394	0.383	0.972	0.972	0.937
perc99	3.241	0.881	0.856	0.806	0.618	0.547	0.538	0.368	0.483	0.811	0.743	0.997	0.997	0.990
absmax20	89.360	1.000	1.000	0.983	0.803	0.609	0.620	0.607	0.664	0.953	0.930	0.996	0.996	0.993
absmin20	3.969	0.732	0.789	0.687	0.563	0.157	0.205	0.252	0.232	0.764	0.783	0.861	0.861	0.889
max	10.957	0.999	0.961	0.993	0.809	0.598	0.611	0.600	0.652	0.942	0.923	0.995	0.995	0.991
min	-22.900	0.000	0.006	0.029	0.004	0.051	0.053	0.051	0.033	0.000	0.000	0.000	0.000	0.000

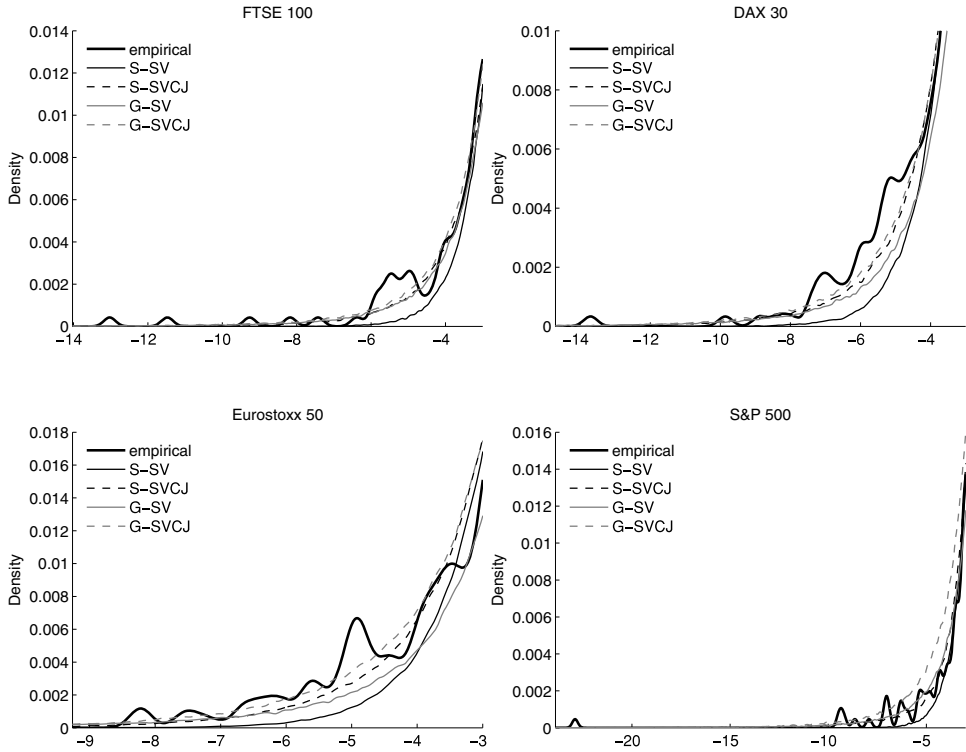


Fig. 2. Unconditional densities (left tail).

This figure depicts the tails of the kernel density fitted to observations on daily log returns over the entire sample from January 1987 to April 2010 (solid line, labelled empirical). It compares this tail with the tails of the densities generated by (S-SV), (S-SVCJ), (G-SV) and (G-SVCJ). The model densities are obtained by Monte-Carlo simulation.

the square-root model, we find that the SVCJ cannot be rejected, but this time even at a significance level of 10%.

The pure diffusion log model fails to capture many of the characteristics of the original data. The inclusion of jumps into the price process proves the most fruitful improvement, but compared with the GARCH specification, its performance is rather weak. Especially the inclusion of jumps in the volatility process may now even deteriorate the simulation results.

The results for the other two European indices follow very similar patterns. Both indices require jumps in both state variables for the square-root model class, which cannot be rejected for both indices at the 5% level. For the Eurostoxx, the simple (G-SV) model however outperforms the complex (S-SVCJ) model, which confirms earlier findings that the GARCH model class can capture many of the features of the European indices even without resorting to complex jump specifications. For the DAX, the (G-SV) also performs extremely well and only the *absmi20* statistic has a low p -value of 0.028. Log volatility models are – for both indices – no improvement over the square-root model class.

The lower half of Table 10 provides the simulation results for the S&P 500 index. In terms of the relative performance of the various models there is little difference to the

European indices. In absolute terms, however, the results are very interesting, because all of the proposed models can be rejected at high significance levels. The more pronounced skewness is the major source of misspecification where all models have p -values of less than 2%. Other characteristics are well captured in the (G) models which confirms their overall superiority over all other model classes considered in this paper.

It is instructive to inspect the tail behaviour of the models in more detail. In Figure 2 we depict the left tails of the empirical densities for all four equity indices and compare them with the densities generated by the point estimates reported in Tables 2, 3, 4 and 6.¹⁹ For expositional clarity, we focus on (S-SV), (S-SVCJ), (G-SV) and (G-SVCJ). The figure confirms that the extreme behaviour of all four equity indices is very poorly represented by the Heston model dynamics. The density of this model converges far more quickly towards zero in the left tail compared with the other models shown here.

7. Summary and Conclusion

We have used daily log returns on four major European and US equity indices between 1987 and 2010 to study the adequacy of twelve different continuous-time jump-diffusion models for capturing the dynamics of the data-generating process. Our model choice includes the popular square-root diffusion model and related specifications where both state variables (returns and variance) are augmented by possibly simultaneous jumps. In addition, we study the same jump extensions for the GARCH diffusion and the log volatility model. Relative performance was assessed according to in-sample fit, residual error analysis and an extensive simulation of posterior predictive p -values. These last results in particular provide vital information on whether the models can produce dynamics that are similar to the observed time-series observations.

In contrast to Andersen *et al.* (2002) and Eraker *et al.* (2003) we find that for the S&P 500 square-root models even with jumps in returns and/or variances are severely misspecified, and this finding is supported by all diagnostic tools we use. One of our main concerns is that a large negative skewness cannot be captured by these models. Log volatility diffusion models improve on the square-root model, but specifications with jumps in price and/or volatility appear to be overspecified and show no overall improvement over the square-root model class. Pure in-sample fit statistics point towards the inclusion of simultaneous price and variance jumps, yet both the analysis of residuals and our simulation study indicate that the simple GARCH diffusion without jumps performs just as well!

The dynamics of European indices are easier to capture than those for the S&P 500. Even square-root models perform quite well for European indices, provided they have jumps in both state variables: they create a realistic number of large negative and positive jumps, with realistic size, and the unconditional distribution generated by the model closely resembles empirical observations. GARCH models improve on the square-root class both in-sample and in simulation experiments, even without resorting to a jump component in the variance. The use of this specification is therefore advantageous in applications that require no (quasi) closed-form of the transition densities.

Our results have important implications for option pricing applications using European equity indices. Since option pricing models are often difficult to distinguish on pure in-sample fit statistics, studying the dynamic behaviour of the underlying process provides

¹⁹ We thank an anonymous referee for suggesting this.

valuable information regarding which option pricing models are empirically relevant. Our results also motivate further empirical research using both data from both the underlying equity index and its options. Since the direct use of option data in models without analytic solutions to European vanilla options is extremely time-consuming, at least over a long sample period such as ours, it might be fruitful to add the term structure of volatility indices (i.e. VIX for the S&P 500 or VDAX for the DAX) into the estimation procedure. The construction methodology of all major volatility indices allows one to derive closed-form solutions even for some non-affine specifications. Although the use of the VIX for estimation purposes is not new, adding the whole term structure rather than a single index might stabilise the estimation of risk premia (and especially the volatility risk premia). This way, using a similar MCMC procedure as in this paper, future research could provide more insights into the structure of risk premia and the ability of the proposed models to explain both the underlying price process and the dynamics of index derivatives.

References

- Ait-Sahalia, Y. and Kimmel, R., 'Maximum likelihood estimation of stochastic volatility models', *Journal of Financial Economics*, Vol. 83, 2007, pp. 413–52.
- Andersen, T.G., Benzoni, L. and Lund, J., 'An empirical investigation of continuous-time equity return models', *The Journal of Finance*, Vol. 57(3), 2002, pp. 1239–84.
- Bates, D.S., 'Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options', *Review of Financial Studies*, Vol. 9(1), 1996, pp. 69–107.
- Bollerslev, T., 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics*, Vol. 31, 1986, pp. 307–27.
- Chacko, G. and Viceira, L.M., 'Spectral GMM estimation of continuous-time processes', *Journal of Econometrics*, Vol. 116(1–2), 2003, pp. 259–92.
- Chernov, M., Gallant, A.R., Ghysels, E. and Tauchen, G., 'Alternative models for stock price dynamics', *Journal of Econometrics*, Vol. 116(1–2), 2003, pp. 225–57.
- Christoffersen, P., Jacobs, K., and Mimouni, K., 'Volatility dynamics for the S&P 500: evidence from realized volatility, daily returns, and option prices.' *Review of Financial Studies*, Vol. 23(8), 2010, pp. 3141–3189.
- Duffie, D., Pan, J. and Singleton, K., 'Transform analysis and asset pricing for affine jump-diffusions', *Econometrica*, Vol. 68(6), 2000, pp. 1343–76.
- Eraker, B., 'Do stock prices and volatility jump? Reconciling evidence from spot and option prices', *Journal of Finance*, Vol. 59(3), 2004, pp. 1367–1404.
- Eraker, B., Johannes, M. and Polson, N., 'The impact of jumps in volatility and returns', *Journal of Finance*, Vol. 53(3), 2003, pp. 1269–1300.
- Fatone, L., Mariani, F., Recchioni, M.C., and Zirilli, F., 'The analysis of real data using a multiscale stochastic volatility model', *European Financial Management*, forthcoming 2011, doi: 10.1111/j.1468-036X.2010.00584.x.
- Gallant, A.R., and Tauchen, G., 'Which moments to match?' *Econometric Theory*, Vol. 12(04), 1996, pp. 657–81.
- Gilks, W.R., Best, N.G. and Tan, K.K., 'Adaptive rejection Metropolis sampling within Gibbs sampling', *Applied Statistics*, Vol. 44(4), 1995, pp. 455–72.
- Heston, S.L., 'A closed-form solution for options with stochastic volatility with applications to bond and currency options', *Review of Financial Studies*, Vol. 6(2), 1993, pp. 327–43.
- Jacquier, E., Polson, N.G. and Rossi, P.E., 'Bayesian analysis of stochastic volatility models', *Journal of Business & Economic Statistics*, Vol. 12(4), 1994, pp. 371–89.
- Jacquier, E., Polson, N.G. and Rossi, P.E., 'Bayesian analysis of stochastic volatility models with fat-tails and correlated errors', *Journal of Econometrics*, Vol. 122(1), 2004, pp. 185–212.

- Jones, C.S., 'The dynamics of stochastic volatility: evidence from underlying and options markets', *Journal of Econometrics*, Vol. 116, 2003, pp. 181–224.
- Li, H., Wells, M.T. and Yu, C.L., 'A Bayesian analysis of return dynamics with Lévy jumps', *Review of Financial Studies*, Vol. 21(5), 2008, pp. 2345–78.
- Masset, P. and Wallmeier, M., 'A high-frequency investigation of the interaction between volatility and DAX returns', *European Financial Management*, Vol. 16(3), 2010, pp. 327–44.
- Nelson, D.B., 'ARCH models as diffusion approximations', *Journal of Econometrics*, Vol. 45(1–2), 1990, pp. 7–38.
- Nelson, D.B., 'Conditional heteroskedasticity in asset returns: a new approach', *Econometrica*, Vol. 59, 1991, pp. 347–70.
- Rubin, D.B., 'Bayesianly justifiable and relevant frequency calculations for the applied statistician', *Ann. Statist.*, Vol. 12, 1984, pp. 1151–72.
- Schönbucher, P.J., 'A market model for stochastic implied volatility', *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, Vol. 357(1758), 1999, pp. 2071–92. Mathematics of Finance.
- Scott, L.O., 'Option pricing when the variance changes randomly: theory, estimation, and an application', *The Journal of Financial and Quantitative Analysis*, Vol. 22(4), 1987, pp. 419–38.
- Skiadopoulos, G. and Hodges, S., 'Simulating the evolution of the implied distribution', *European Financial Management*, Vol. 7(4), 2001, p. 497.