Volatility Modelling with Jumps: An Empirical Approach

Angela Black

University of Aberdeen, Business School

Jing Chen

Swansea University, School of Business and Economics

Oleg Gustap

University of Aberdeen, Business School

Julian Williams*

University of Aberdeen, Business School

Abstract

This paper presents a practical approach for calibrating jump models using historical data for the underlying instrument and forward looking information from the derivatives market. Our investigations reveal that the AR-Jump model outperforms most other models when applied to historical data. Additionally, a test data-set of options contracts indicates that, during 2008, jump models are most appropriate in calibrating the volatility forecasts to the actual prices of options contracts. A distinctive finding is that, in general, the best fitting model for historical data is not necessarily the model that best calibrates to the forward looking options data.

Keywords: Autoregressive Jump Models, NGARCH, Volatility

*Corresponding Author. We would like to thank Yuzhi Cai, Allan Walkins and Dylan Thomas for constructive comments on this initial draft.

Email addresses: angela.j.black@abdn.ac.uk (Angela Black), jing.chen@swansea.ac.uk (Jing Chen), O.Gustap@abdn.ac.uk (Oleg Gustap), julian.williams@abdn.ac.uk (Julian Williams)

1. Introduction and Literature

This paper demonstrates and calibrates an econometric model of time varying volatility with jumps, mixing historical observations with forward looking options data. We then use a dataset of historical equity index data combined with intraday index options prices to demonstrate the change in jump intensities during 2008.

Since Bollerslev (1986) generalised the Autoregressive Conditional Heteroskedasticity (ARCH) model, proposed by Engle (1982), a number of different approaches have improved the explanatory power and the out-of-sample performance of volatility clustering models. However, the traditional GARCH model has significant limitations for option valuation: first, it fails to account for a potential leverage effect and second, spot volatility is usually found to be highly persistent. The former arises because negative news may have an additional adverse impact on prices through the decrease of the debt-to-equity ratio thereby complicating the valuation of deep out-of-the-money options. The latter point reduces the performance of the model for close to maturity options. These shortcomings have been addressed through a wide range of extensions, including Nelson (1991) (EGARCH), Glosten et al. (1993) (GJR-GARCH) and Heston (1993) (NGARCH), where all of these models generate negative implied skewness of the asset return distribution, which in turn leads to higher prices of the out-of-the-money put options when compared to the Black and Scholes formula under time invariant volatility. More recently, several empirical studies have demonstrated the importance of time-varying volatility and negative skewness in reducing observed anomalies associated with the Black and Scholes pricing formula when compared to actual market data, see for instance, Bakshi et al. (1997), Bates (2000), Broadie et al. (2007)) and, for an example using non-normal innovations, see Christoffersen et al. (2006)).

Modelling higher levels of persistence in time varying volatility models has been the aim of a number of discrete and continous time models. For example, Jorion (1998) proposes a model with normally distributed Poisson jumps highlighting the implications for option pricing. Additionally, strong evidence of jump components in foreign exchange markets has been found by Jorion (1998) and Palm and Vlaar (1993), while Chan and Maheu (2002)

and Maheu and McCurdy (2004) investigate a Poisson-normal jump model on individual stocks and broad indices. The two latterly mentioned papers model time-varying conditional jump intensity as the simple time-series process driven by the autoregressive component and the ex-post jump probability. Further observations in the literature reveal that modelling Poisson-normal jumps in returns and variance along with the conventional GARCH processes could lead to significant improvement in capturing the underlying process, see for instance Maheu and McCurdy (2004), and Christoffersen et al. (2008)).

Following the analysis in Christoffersen et al. (2008) this paper presents the component GARCH specification of the Heston and Nandi (2000) NGARCH model. A closed-form solution exists for the component GARCH models, hence option pricing is relatively straightforward. The data covers 1988-2010 which provides an additional challenge for modelling index volatility due to the highly volatile period during the 2008 financial crisis. We find that the long-run volatility component is highly persistent, hence we study a special case, where the long-run component is modelled to be fully persistent. Our models generate richer autocorrelation structures and volatility of variance paths in comparison to the benchmark NGARCH model. However, based on the likelihood criterion the persistent specification is found inferior not only for the component model but also for the benchmark NGARCH for all the indices studied in this paper.

The NGARCH model with Poisson-normal jumps in returns and variance, which can potentially alleviate some biases associated with the Heston and Nandi (2000) model, is used as an additional benchmark for the component GARCH model in Christoffersen et al. (2008). It appears that modelling conditional non-normalities in the return distribution improves the fit of the model for the historical returns series. On tbe other hand, the option valuation comparison provides strong evidence in favour of the component GARCH model. Poisson-normal jumps improve the NGARCH model's performance particularly in valuing the long-maturity and deeply out-of-the-money options, at the same time they worsen pricing of the short-maturity options.

Evidence that jumps tend to cluster together has been found in the continuous-time and discrete-time literature. We choose a functional form for the conditional jump intensity

process similar to the one used in Chan and Maheu (2002) and in Maheu and McCurdy (2004). The model is applied to data on daily total returns from the following: the Standard and Poor's 500 (S&P 500), the Dow Jones Industrial Average (DJIA), the Deutscher Aktien Index 30 (DAX 30) and the FTSE 100 indices. The maximum likelihood analysis suggests that the best fit of historical return dynamics is achieved by the AR-Jump model, followed by the Constant-Jump model with the constant jump intensity and the component GARCH model being the worst alternative among these three models.

Option valuation performance for a complete cross section of index options for the AR-Jump model is encouraging. First, the component NGARCH outperforms the component GARCH model using the root mean squared error (RMSE), which is 1% lower for the AR-Jump model. Second, the RMSE value corresponding to the shortest maturity options is 23% lower for the AR-Jump model. In addition, a comparison of the implied volatility biases suggests that the AR-Jump model is able to generate an implied volatility path which better matches the spot volatility path from market prices.

The remainder of this paper is structured as follows: §(2) provides the derivation and analysis of the component GARCH model by modelling short-run and long-run volatility components. §(3) discusses models with Poisson-normal jumps and introduces the AR-Jump model. §(4) examines the option valuation performance of the component GARCH and AR-Jump models through calibration to SPX option data, and finally, §(5) concludes.

2. Modelling short-run and long-run volatility components

This section follows the analysis conducted in Christoffersen et al. (2008) and derives the component GARCH model with the special case of extreme persistence called the persistent GARCH model. Both models stem from the Heston and Nandi (2000) NGARCH model.

Construction of the component GARCH model

Heston and Nandi (2000) show that the following NGARCH specification allows for a closed-form solution for the price of a European call option

$$R_{t+1} = \ln\left(S_{t+1}/S_t\right) = r + \lambda h_{t+1} + z_{t+1}\sqrt{h_{t+1}} \tag{1}$$

$$h_{t+1} = w + bh_t + a\left(z_t - c\sqrt{h_t}\right)^2 \tag{2}$$

where S_{t+1} denotes the asset price; R_{t+1}, the logarithmic return; r, the risk-free rate; λ, the risk premium; h_{t+1}, the daily variance; z_{t+1}, the independently identically distributed (i.i.d.) normal $\mathcal{N}(0,1)$ shock. This model stems from the conventional GARCH(1,1) specification and includes a leverage effect, captured by the parameter c. The unconditional variance, is therefore

$$\mathbb{E}\left[h_{t+1}\right] = \sigma^2 = w + b\sigma^2 + a + ac^2\sigma^2 \tag{3}$$

$$\sigma^2 = \frac{w+a}{1-b-ac^2} \tag{4}$$

substituting w into the original equation one reveals

$$h_{t+1} = \sigma^2 + b\left(h_t - \sigma^2\right) + \alpha\left(\left(z_t - c\sqrt{h_t}\right)^2 - \left(1 + c^2\sigma^2\right)\right) \tag{5}$$

Now, assuming the unconditional mean varies with time and denoting it by q_{t+1}, the equation may be rewritten as

$$h_{t+1} = q_{t+1} + \beta\left(h_t - q_t\right) + \alpha\left(\left(z_t - \gamma_1\sqrt{h_t}\right)^2 - \left(1 + {\gamma_1}^2 q_t\right)\right) \tag{6}$$

This model stems from the Heston and Nandi (2000) NGARCH model and is functional for a model that allows for the closed-form solution for the price of European call option. Following Christoffersen et al. (2008), we denote two components of the volatility process, a long-run component q_{t+1} and a short-run mean-zero deviations $h_{t+1} - q_{t+1}$. The model

could be rewritten as:

$$h_{t+1} = q_{t+1} + \left(\alpha {\gamma_1}^2 + \beta\right)(h_t - q_t) + \alpha\left(\left(z_t - \gamma_1\sqrt{h_t}\right)^2 - \left(1 + {\gamma_1}^2 h_t\right)\right) \tag{7}$$

Where $\left(z_t - \gamma_1\sqrt{h_t}\right)^2 - (1 + {\gamma_1}^2 h_t) = (z_t - 1) - 2\gamma_1 z_t\sqrt{h_t} = v_{1,t}$, is a mean-zero innovation. Denoting $\tilde{\beta} = \alpha{\gamma_1}^2 + \beta$ and specifying the functional form of the long run component q_{t+1} , the model appears as

$$\begin{aligned} h_{t+1} &= q_{t+1} + \tilde{\beta}(h_t - q_t) + \alpha v_{1,t} && (8) \\ q_{t+1} &= \omega + \rho q_t + \varphi v_{1,t} && (9) \end{aligned}$$

where

$$v_{i,t} = \left({z_t}^2 - 1\right) - 2{\gamma_i}^2 z_t\sqrt{h_t}, \quad \text{for } i = 1, 2. \tag{10}$$

This model is now referred to as the component GARCH model or simply as the component model. Note that if $\rho < 1$ then the unconditional variance is equal to the unconditional expectation of the long-run component, that is $\mathbb{E}[h_{t+1}] = \mathbb{E}[q_{t+1}] = \omega/(1-\rho)$. In addition, both innovation terms are mean-zero $\mathbb{E}[v_{i,t}] = \mathbb{E}_{t-1}[v_{i,t}]$ for $i = 1, 2$. As a result, the component model has eight parameters λ, $\tilde{\beta}$, α, γ_1, ω, ρ, *varphi* and γ_2, while the benchmark Heston and Nandi (2000) NGARCH model has five.

The Persistent GARCH model

Conditional variance in the component model is a mean-reverting process if $\rho < 1$, therefore the fully persistent long-run component is non-degenerate when $\rho = 1$,

$$\begin{aligned} h_{t+1} &= q_{t+1} + \tilde{\beta}(h_t - q_t) + \alpha v_{1,t} && (11) \\ q_{t+1} &= \omega + q_t + \varphi v_{2,t}, && (12) \end{aligned}$$

where $v_{i,t}$ for $i = 1, 2$ are the mean-zero innovations, with functional forms identical to those of the component model. We refer to this model as the persistent GARCH model or

simply as the persistent model. This model contains an additional two additional parameters compared to the benchmark model Heston and Nandi (2000) .

In this model, the long-run volatility component is a unit-root process, hence shocks to this component never die out. As in the case of the component model we refer to the q_{t+1} as the long-run component and to $h_{t+1} - q_{t+1}$ as the short-run component. It is possible to refer to these components as the permanent and transitory, respectively, as the former is the unit-root process and later reverts to zero. It is expected that the component model, which nests the permanent component model, is superior to the latter in-sample, while out-of-sample the permanent component alternative may perform better. The fact that the permanent component model represents a unit-root process may play a crucial role out-of-sample, as it is able to capture structural break effects and adjust for it, while the component model is not.

Empirical results from the S&P 500 returns

The parameters of the NGARCH, component and persistent models can be easily estimated using the maximum likelihood (ML) method. Table 1 presents the ML estimates of these models, obtained from the daily total return data on the Standard and Poor's 500 (S&P 500) index over the period 1988-2010. Most of the estimates are significantly different from zero at 5% significance level. The risk premium is highly significant in all models. Based on the log-likelihood criterion, the component model performs the best followed by the NGARCH model, which is followed, in turn, by the persistent model. Surprisingly, the persistent model performs worse than the NGARCH benchmark. This highlights the importance of modelling the long-run variance persistence lower than 1.

Properties of the return process

Figure 1 presents an analysis of the improvement of the component model fit over the benchmark NGARCH model. It displays the sample path of the spot variance in the NGARCH model (Panel A), as well as in the component model (Panel B), and the variance components in the component model over the period 2006-2010. Panel D presents the short-run mean-zero variance component, which adds high-frequency noise to the long-run

component (Panel C). As a result, the spot variance path of the component model (Panel B) appears to be noisier than the path generated by NGARCH model (Panel A). This lead to a higher volatility of the variance in the component model.

Figure 2 further analyses the performance of the component GARCH model in comparison to the persistence GARCH model. The three left panels depict the 2005-2010 sample path of spot variance (Panel A), the sample path of long-run variance component (Panel C) and the sample path of short-run variance component (Panel E) for the component GARCH model. The three right panels present the similar sample paths for the persistent GARCH model. Clearly, the persistent model's performance is affected by the presence of a unit-root in the long-run variance component, particularly during the time of financial crisis 2008-2009. Due to extreme persistence of the variance process, the persistent model is unable to adjust for the spike in volatility as fast as the component model. Furthermore, the effect of this shock decays at a slow pace, causing serious disturbances in the short-run variance component, which is found to be negative over the most part of 2009. Note that the long-run variance component in the component model returns to its 2007 level by the end of 2009, whereas in the persistence model the variance level observed in the first half of 2008 is reached by the end of 2010. To conclude, in-sample comparison of the component GARCH and the persistent GARCH models highlights the necessity of modelling the long-run variance persistence different from one.

In addition to the conventional log-likelihood comparison, the three models can be compared by assessing their variance persistence properties. The variance persistence of the NGARCH model is defined as $b + ac^2$ of the component GARCH as $\rho + \tilde{\beta}\,(h_t - q_t)$ (see Christoffersen et al. (2008)). As expected, the variance persistence of the component GARCH model 0.9954 is almost equal to 1, which is the persistent GARCH case. However, it appears that modelling persistence as the component GARCH model is important for the performance. The difference between the likelihoods can be caused by the difference in the short-run and the long-run persistence of the two models. Short-run persistence of the component GARCH model is equal to 0.6673 versus 0.8852 for the persistent GARCH, whereas the persistence of the long-run component is 0.9861 in contrast to 1. The persistence

of the NGARCH model is 0.9698, which is higher then the average of estimates in Heston and Nandi (2000), the estimate in Christoffersen et al. (2008) and slightly higher than the estimate in Christoffersen et al. (2006).

[INSERT A SINGLE PANEL OF FIGURES 1 AND 2 ABOUT HERE]

The volatility of the conditional variance and the covariance and correlation between returns and variance play a crucial role in option valuation. Option prices depend on the underlying volatility, thus variation of the option price is a consequence of the volatility of variance. Christoffersen et al. (2008) show that the volatility of the conditional variance is closely related to kurtosis and the ability to replicate option prices volatility. Therefore, the higher values of conditional variance indicate an ability to generate richer kurtosis and better fit the evolution of option prices.

The functional form of the NGARCH model generates the following function of the conditional variance of variance

$$\mathrm{Var}_t\left(h_{t+2}\right) = 2a^2 + 4a^2c^2h_{t+1} \tag{13}$$

For the component and persistent models, the conditional variance of variance is equal to

$$\mathrm{Var}_t\left(h_{t+2}\right) = 2(a+\omega)^2 + 4(a\gamma_1 + \omega\gamma_2)^2 h_{t+1} \tag{14}$$

The three left panels in Figure 3 present the paths of the conditional standard deviation of variance in the three models during the period 2006-2010. It is clear that the standard deviation of variance in the component model is generally higher and more volatile than in the NGARCH model. Table 1 reports the average volatility of variance during 2006-2010; the component GARCH result is almost two times higher than the NGARCH one.

[INSERT FIGURE 3 ABOUT HERE]

There is a link between the volatility of variance and kurtosis. Figure 3, therefore, shows that component GARCH model is able to generate richer kurtosis and potentially richer option prices. Heston and Nandi (2000) show that the conditional covariance between the variance and return in the NGARCH model is given by

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = \mathbb{E}_t\left[\left(R_{t+1} - \mathbb{E}_t\left[R_{t+1}\right]\right)\left(h_{t+2} - \mathbb{E}_t\left[h_{t+2}\right]\right)\right] = -2ach_{t+1} \tag{15}$$

Consequently, the conditional correlation in the NGARCH model is defined as follows

$$\mathrm{Corr}_t\left[h_{t+1}, h_{t+2}\right] = \frac{-2ch_{t+1}}{\sqrt{2 + 4c^2 h_{t+1}}} \tag{16}$$

The conditional covariance in the component GARCH model is given by

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = -2\left(a\gamma_1 + \varphi\gamma_2\right)h_{t+1} \tag{17}$$

thus, conditional correlation in the component GARCH model is

$$\mathrm{Corr}_t\left[h_{t+1}, h_{t+2}\right] = \frac{-2\left(a\gamma_1 + \varphi\gamma_2\right)h_{t+1}}{\sqrt{2(a+\varphi)^2 + 4(a\gamma_1 + \varphi\gamma_2)^2 h_{t+1}}} \tag{18}$$

The three right panels in the Figure 3 show the conditional correlation between the return and the conditional variance in the three models. In general, conditional correlation in the component and persistent models is lower than in the NGARCH model. In the component and persistent models conditional correlation drops almost to -1 in the first quarter of 2007 and stays extremely low till the end of 2009 in the component model and till the end of the sample (end of 2010) in the persistent model. Table 1 reports the average correlations during 2006-2010, -92.74% in the NGARCH model, -97.41% in the component GARCH and -96.71% in the persistent GARCH.

[INSERT TABLE 1 ABOUT HERE]

3. Poisson-normal jumps and an Autoregressive (AR) jumps model

Our motivation for contributing a deeper understanding of the jump process is fundamentally rooted in the findings from previous empirical studies which appear to confirm the importance of including jump effects in both returns and volatility. For example, Jorion (1998) and Palm and Vlaar (1993) show that modelling constant intensity jumps in the return dynamics can improve the fit of the model while Andersen et al. (2002) demonstrates the apparent reduction in biases for models such as the Heston (1993) model, when accounting for jumps. Others, such as Chan and Maheu (2002), Jorion (1998), Maheu and McCurdy (2004) and Palm and Vlaar (1993) show the relevance of jump effects in return dynamics for discrete-time models, while evidence confirming the importance of Poisson-normal jumps in returns and volatility for continuous time modelling can be found in Andersen et al. (2002), Bakshi et al. (1997), Bates (1996, 2006), Chernov et al. (2003), Eraker et al. (2003), Eraker (2004) and Pan (2002). Perfectly correlated jumps are examined in greater depth by Broadie et al. (2007), Eraker (2004), and Eraker et al. (2003) while Christoffersen et al. (2008) and Elkamhi and Ornthanalai (2010) demonstrate improvement in estimating option valuation when using a discrete-time Poisson-normal jump model. This section outlines Poisson-normal jumps and provides the building blocks for the AR jump model which we will compare with the Poisson-normal jump model in §4 for pricing options.

Christoffersen et al. (2008) augment the Heston and Nandi (2000) NGARCH model with perfectly correlated Poisson-normal jumps in returns and variance to provide a more challenging benchmark for the component model. Their GARCH(1,1)-Jump model significantly improves the fit of the S&P 500 daily return series. However, while the component model is superior, when compared on the basis of option pricing the GARCH(1,1)-Jump model performs well in valuing deep out-of-the-money and long-maturity options, although it fails to surpass NGARCH benchmark in pricing short-maturity options. The better performance of the GARCH(1,1)-Jump model for long-maturity options could be linked to the functional form of variance, which yields non-normal behaviour of the conditional distribution at longer horizon.

Poisson-normal jumps in returns and volatility

The Constant-Jump model can be described in the following way: Consider the following return R_t generating process

$$R_{t+1} = r + \lambda_z h_{t+1} + \lambda_y \eta + z_{t+1}\sqrt{h_{t+1}} + y_{t+1} - \mu\eta \tag{19}$$

$$h_{t+1} = w + bh_t + a\left(z_t + \frac{y_t}{\sqrt{h_t}} - c\sqrt{h_t}\right)^2 \tag{20}$$

where λ_z is the diffusive risk premium, λ_y is the jump risk premium, z_{t+1} is the diffusive innovation and y_{t+1} is the jump innovation. The functional form of the model insures that the expected log excess returns are equal to the sum of diffusive and jump risk premiums. Diffusive innovation z_{t+1} is assumed to be independently and identically distributed $\mathcal{N}(0,1)$, while the jump innovation y_{t+1} is a simple compound Poisson process of the following form

$$y_{t+1} = \sum_{j=1}^{N_{t+1}} X_{t+1}^{\{j\}}, \tag{21}$$

with

$$X_{t+1}^{\{j\}} \sim \mathcal{N}\left(\mu, \tau^2\right), \quad \text{for } j = 1, 2, ..., N_{t+1} \tag{22}$$

where N_{t+1} is a Poisson random variable with the constant intensity η and conditional density

$$\mathbb{P}\left(N_{t+1} = j\right) = \frac{e^{-\eta}\eta^j}{j!} \tag{23}$$

There is growing evidence that jumps tend to cluster together in a similar manner to the conditional volatility processes. Logically, we can think of sustained episodes of abnormal volatility and market crashes which can be realized in a series of jumps over a short period of time. For example, using option pricing data, Bates (1991) finds that the number of expected jumps around the 1987 crash behaved in a systematic fashion i.e. a piecewise constant intensity around the crash. This finding is reinforced by the continuous-time literature where evidence of clustering of jump intensities in equity returns is apparent, see for instance Andersen et al. (2002), Bates (2000), Chernov et al. (2003), and Pan (2002).

To shed further light on the appropriate jumps process process we offer additional insight to the GARCH(1,1)-Jump model of Christoffersen et al. (2008) by modelling jump arrival intensity in a time-varying framework. We take as our starting point the model presented in Chan and Maheu (2002) and Maheu and McCurdy (2004) where the conditional jump intensity is modelled as an ARMA process, which is determined by a one-lag autoregressive component and the ex-post assessment of the jump probability. This model has a similar functional form to the Constant-Jump model with the added dimenstion that the jump intensity is modelled as a time-varying process. We refer to this model as the AR-Jump model.

Intuitively, the clustering of jumps is interpreted as a function of the clustering of news innovations. Diffusion and jump components, in turn, are the normal and the highly unusual, high impact, news events, respectively. Conventionally, the former is modelled as the normal innovation process and the latter is modelled as the compound jump Poisson process. It is, therefore, assumed that the normal news cause smoothly evolving fluctuations of the conditional variance, while the jump process causes exceptionally large movements of returns and variance. For the individual stocks, unusual news represents unexpected announcements about prospective cash flows and for the stock indices, they represent unexpected information, which affect a wide range of companies.

Building the AR-Jump model

Denoting information available at time t by Φ_t, the AR-Jump model for returns is defined as follows

$$R_{t+1} = r + \lambda_z h_{t+1} + \lambda_y \eta_{t+1} + z_{t+1}\sqrt{h_{t+1}} + y_{t+1} - \mu\eta_{t+1} \tag{24}$$

with volatility dynamics

$$h_{t+1} = w + bh_t + a\left(z_t + \frac{y_t}{\sqrt{h_t}} - c\sqrt{h_t}\right)^2 \tag{25}$$

and jump intensity

$$\eta_{t+1} = \eta_0 + \rho\eta_t + \gamma\xi_t \tag{26}$$

$$y_{t+1} = \sum_{j=1}^{N_{t+1}} X_{t+1}^{\{j\}} \tag{27}$$

with

$$X_{t+1}^{\{j\}} \sim \mathcal{N}\left(\mu, \tau^2\right) \tag{28}$$

for $j = 1, 2, ..., N_{t+1}$, and

$$\mathbb{P}\left(N_{t+1} = j \,|\Phi_t\right) = \frac{e^{\eta_{t+1}}\eta_{t+1}^j}{j!} \tag{29}$$

where ξ_t is the intensity residual and η_t is the ex ante assessment of the number of jumps to occur in period t. Jorion (1998), Christoffersen et al. (2008) and Constant-Jump models represent the special case of the above specification with ρ and γ restricted to zero. The intensity residual has the following form

$$\xi_t = \mathbb{E}\left[N_t \,|\Phi_t\right] - \eta_t = \sum_{j=0}^{\infty} j\mathbb{P}\left(N_t = j \,|\Phi_t\right) - \eta_t \tag{30}$$

The probability $\mathbb{P}\left(N_t = j \,|\Phi_t\right)$ is the ex post inference on the number of jumps occurred in the period t, given information available at time t. $\mathbb{E}\left[N_t \,|\Phi_t\right]$ is the ex post assessment of the number of jumps that occurred at time t and η_t is the ex ante expectation. In other words, intensity residual ξ_t represents the change of the conditional expectation as the information set updates

$$\xi_t = \mathbb{E}\left[N_t \,|\Phi_t\right] - \mathbb{E}\left[N_t \,|\Phi_{t-1}\right]. \tag{31}$$

Note, first, that by construction conditional and unconditional expectations of the intensity residual are both equal zero

$$\mathbb{E}\left[\xi_t \,|\Phi_{t-1}\right] = \mathbb{E}\left[\xi_t\right] = 0 \tag{32}$$

Second, if jump intensity is a stationary process ($\rho < 1$), then the unconditional jump intensity is equal to

$$\mathbb{E}\left[\eta_t\right] = \frac{\eta_0}{1-\rho} \tag{33}$$

which is time invariant.

Conditional moments of returns

Both models generate similar conditional moments of returns. The derivation of the first four conditional moments of returns can be found in Devroye (1982). Denoting conditional variance, skewness and kurtosis of the returns as $\mathrm{Var}\left(R_t\left|\Phi_{t-1}\right.\right)$, $Sk\left(R_t\left|\Phi_{t-1}\right.\right)$ and $Kur\left(R_t\left|\Phi_{t-1}\right.\right)$, respectively:

$$\begin{aligned}
\mathbb{E}\left(R_t\left|\Phi_{t-1}\right.\right) &= r + \lambda_z h_t + \lambda_y \eta_t && (34)\\
\mathrm{Var}\left(R_t\left|\Phi_{t-1}\right.\right) &= h_t + \left(\tau^2 + \mu^2\right)\eta_t && (35)\\
Sk\left(R_t\left|\Phi_{t-1}\right.\right) &= \frac{\eta_t\left(\mu^3 + 3\mu\tau^2\right)}{\left(h_t + \eta_t\tau^2 + \eta_t\mu^2\right)^{3/2}} && (36)\\
Kur\left(R_t\left|\Phi_{t-1}\right.\right) &= 3 + \frac{\eta_t\left(\mu^4 + 6\mu^2\tau^2 + 3\tau^4\right)}{\left(h_t + \eta_t\tau^2 + \eta_t\mu^2\right)^2} && (37)
\end{aligned}$$

In the case of the Constant-Jump model, conditional jump intensity is a constant. All moments are effected by the conditional jump intensity η_t and the conditional variance h_t. The sign of the conditional skewness, in turn, depends only on the sign of the mean of the jumps' distribution.

Log likelihood function

The log-likelihood function for the mixture of the normal and Poisson processes represents the sum of conditional Normal densities of returns, given the number of jumps occurring

$$\mathbb{P}\left(R_t\right|N_t = j, \Phi_{t-1}) = \frac{1}{\sqrt{2\pi\left(h_t + j\tau^2\right)}}\exp\left(-\frac{\left(R_t - r - \lambda_z h_t - \lambda_y\eta_t - \mu j + \mu\eta_t\right)^2}{2\left(h_t + j\tau^2\right)}\right) \tag{38}$$

Thus, log-likelihood LL has the following functional form

$$LL = \sum_{t=1}^{r} \log \left(\sum_{j=0}^{\infty} \mathbb{P}\left(R_t | \, N_t = j, \Phi_{t-1} \right) \mathbb{P}\left(N_t = j | \Phi_{t-1} \right) \right) \tag{39}$$

The jump intensity function of the AR-Jump is defined as above, with the special case of constant jump intensity for the Constant-Jump model. The ex post inference on the number of jumps occurred at time t, based on the information available at this time is defined as

$$\mathbb{P}\left(N_t = j \, | \Phi_t \right) = \frac{\mathbb{P}\left(R_t | \, N_t = j, \Phi_{t-1} \right) \mathbb{P}\left(N_t = j | \Phi_{t-1} \right)}{\sum\limits_{j=0}^{\infty} \mathbb{P}\left(R_t | \, N_t = j, \Phi_{t-1} \right) \mathbb{P}\left(N_t = j | \Phi_{t-1} \right)} \tag{40}$$

and represents the ex post distribution for the number of jumps, N_t. Note, that the terms in the log-likelihood function and ex post distribution include infinite summation. In order to numerically optimise these functions, the summation has to be truncated after some value. Similar to Jorion (1998) we truncate the summation in the Constant-Jump model at 10. For the AR-Jump model the satisfactory accuracy is achieved when the summation is truncated at 40. Maheu and McCurdy (2004) truncate it at 25, however our model is estimated on a much more volatile data and needs to capture up to 20 jumps per day.

The S&P 500 model

Table 2 presents the MLE estimates of the Constant-Jump and the AR-Jump models' parameters for the total returns data on the S&P 500 index over the period 1988-2010. Seven out of nine parameter estimates for the Constant-Jump model are significantly different from zero at 5% significance level. Since the jump intensity, the average jump size and the standard deviation of jumps are found significant, there is a strong jump effect in returns and volatility. Eight out of eleven parameters are significantly different from zero at a 5% significance level. All the parameters of the Poisson-normal jumps are found significant as in the Constant-Jump model. The significance of the conditional jump arrival intensity parameters confirms that the jump arrival process tends to significantly deviate from its unconditional mean. The difference between likelihood values (19,675 in the Constant-Jump

model versus 19,783 in the AR-Jump model) suggests that the AR-Jump model is able to better fit historical returns.

[INSERT TABLE 2 ABOUT HERE]

Interestingly, the estimate for the jump risk premium is not statistically significant suggesting that this variable is less important than we may have expected. The diffusive risk premium is found to be significant in both models and the estimate in the AR-Jump model is similar to the estimate in Christoffersen et al. (2008). The average jump size of the Constant-Jump model is estimated at -1.2%, similar to the estimates in Christoffersen et al. (2008) and Chernov et al. (2003), and lower than the estimates in Eraker et al. (2003) and Eraker (2004). The jump arrival intensity in the Constant-Jump model is equal to 0.025, which is equivalent to approximately 6.3 jumps per year. Both the mean and the standard deviation of the jump size distribution in AR-Jump model are significant at the 5% significance level. The lower average jump size in the AR-Jump model allows for a higher flexibility in the number of jumps that occur in each period. Negative signs of the estimates in the Constant-Jump and the AR-Jump models ensure that the conditional skewness of returns is negative in both models. Overall, the significance of the jump size distributions' parameters in both models suggests that the jump dynamics affect return distributions through the conditional variance, conditional skewness and conditional kurtosis.

The average jump intensity (mean of η_t) is equal to 0.84, while unconditional expectation of the number of jumps ($\frac{\eta_0}{1-\rho}$)is estimated at 0.93, this does not make the later an unbiased forecast of the former. The significant difference between the unconditional and realized mean arises due to the fact that the effect of the previous period jump intensity is neutralised by the intensity residual component, as a result, the conditional jump intensity process bears only 10% of information.

Figure 4 depicts the conditional variance of returns of the Constant-Jump (Panel A) and the AR-Jump (Panel B) models, and the conditional jump intensity (Panel C), which denotes conditional ex ante expectation of the number of jumps to occur in the next period. It is clear

that the AR-Jump model is able to generate more volatile return and conditional variance dynamics and absorb shocks during the volatile period of 2008-2009. The functional form of the time-varying jump intensity process enables it to adjust for the increased volatility in 2008-2009 and the ex ante expectation of the number of jumps increases from nearly zero in 2006 to 19 by the end of 2008.

[INSERT A SINGLE PANEL OF FIGURES 4 AND 5 ABOUT HERE]

It is also possible to reckon the ex post probability of at least one jump in a period as $\mathbb{P}(N_t \geqslant j \,|\Phi_t) = 1 - \mathbb{P}(N_t = 0 \,|\Phi_t)$. Figure 5 displays the ex ante jump intensity expectation η_t in Panel A, the realized number of jumps $\mathbb{E}(R_t \,|\Phi_t)$ in Panel B, the ex ante jump probabilities $\mathbb{P}(N_t \geqslant j \,|\Phi_{t-1})$ in Panel C and the ex post jump probabilities $\mathbb{P}(N_t \geqslant j \,|\Phi_t)$ in Panel D. As expected, realized number of jumps appears to be more volatile than the conditional jump intensity. The average ex ante probability of jump is equal to 36.9% versus the average ex post probability of jump 37.0%, which makes the former to be an unbiased forecast of the later. In the most volatile period of the second half 2008 and the first half of 2009 both probabilities are close to 1, as a result of sharp increase in the conditional jump intensity and the realized number of jumps.

The conditional variance of returns, which is given above, represents a combination of the conventional GARCH component h_t and the jump variance component $(\tau^2 + \mu^2)\,\eta_t$. Figure 6 shows the conditional variance decomposition during 2006-2011. Panel A displays the total conditional variance and Panel B displays individual variance components. It is clear that the GARCH component generates a smoothly evolving path of the conditional variance, while the noise and large deviations are generated by the Jump component. Jump component accounts for 33% of the spot variance during the whole sample and for 47% during second half of 2008 and first half of 2009. Hence, although Poisson-normal jumps play significant role in the return distribution, more than 50% of disturbances are explained by the smooth normal diffusion even in periods of extreme volatility.

[INSERT FIGURE 6 ABOUT HERE]

Comparison with other indices

In addition to the previous analysis of S&P 500 total return index data we investigate the performance of the discussed models on different sets of data. We study daily total return data for three indices: the Dow Jones Industrial Average (DJIA), the FTSE 100 and the Deutscher Aktien Index (DAX 30) from January 4, 1988 to December 31, 2010. Our findings indicate that the component model is superior to the NGARCH and persistence models, while the AR-Jump model outperforms the Constant-Jump model on any set of data. Therefore, these inferior models are excluded from the analysis.

Maximum likelihood estimates of the parameters for the indices are presented in Table 3. In general, they are consistent with the estimates for the S&P 500 total return index reported in Table 1. All indices data exhibit high persistence of variance, with the highest level of 0.9994 persistence of the FTSE 100 return. Persistence of the long-run component is close to 1 in all series and varies from 0.979 in the DAX 30 return to 0.991 in the FTSE return, while short-run persistence is between 0.6 and 0.7 in the DAX 30 and DJIA returns and 0.94 in the FTSE 100 return.

Table 4 reports maximum likelihood estimates of parameters for the three indices. In all cases the AR-Jump model significantly outperforms the component GARCH model basing on the likelihood criterion. The diffusive risk premium λ_z estimate is significant at 5% significance level, whereas the jump risk premium λ_y is found insignificant for the three indices. Compared to the results in Table 2, jump mean sizes are significantly negative and smaller than the estimate for the S&P 500 total return index. All the results suggest that modelling time-varying jump arrival intensity may improve the fit of historical data. The autoregressive parameters ρ are significant for all series and stay in line with the estimate in Table 2.

[INSERT TABLE 4 ABOUT HERE]

The parameter γ, which determines the effect of the intensity residual, is relatively low for the DAX 30 returns and is almost equal 1 for the DJIA and FTSE 100, which is higher than estimate for the S&P 500 data, though the latter is insignificant. As a result, the unexpected component has a much higher influence on the conditional jump intensity in DJIA and FTSE 100 returns than in DAX 30 data. The unconditional mean of the jump intensity is estimated at 0.42 for the DAX 30 returns, 0.82 for the DJIA returns and 0.94 for the FTSE 100 returns, while the observed mean are equal to 0.96, 0.74 and 0.92 respectively. Hence, the unconditional mean is not an unbiased forecast of the jump arrival intensity for the DAX 30 and DJIA returns, presumably due to relatively short samples.

Figure 7 plots the conditional variance and its constituent components during 2006-2011. The GARCH component refers to the variance of the normal diffusion and the Jump component refers to the variance of the compensated Poisson-normal jump process. Although GARCH components account for the 60%-70% on average, it is clear that during the volatility at the end of 2008 stems from the Jump component. The flexibility of the AR-Jump model enables it to better capture the structure of the historical data in comparison with the Constant-Jump and the component GARCH models. However, improved fit of the historical return series does not necessary lead to the better performance in pricing of options contracts, as this is dependent on a model 'beauty contest' conducted by the market participants.

[INSERT FIGURE 7 ABOUT HERE]

4. Calibration to SPX Option Data

In this section we will now compare the performance of the component GARCH model and the AR-Jump model for pricing traded call option contracts on the Standard and Poor's 500 index.

Option valuation formula

Given the risk-neutral dynamic, at time t, European call option with the exercise price K that expires at time T is priced at:

$$\text{Call price} = e^{-r(T-t)}\mathbb{E}_t^{\psi}\left[\max\left(S_T - K, 0\right)\right], \tag{41}$$

where $\mathbb{E}_t^{\psi}$ denotes the conditional expectation under the risk-neutral dynamics. This expectation might be estimated by the Monte Carlo simulation or, in case of the NGARCH and component models, by the closed-form solution formulae provided in Heston and Nandi (2000) and Christoffersen et al. (2008). Denoting a conditional moment generating function by $f\left(t, T : \phi\right)$, Christoffersen et al. (2008) show that it can be written as:

$$\begin{aligned} f\left(t, T : \phi\right) &= \mathbb{E}_t\left[\exp\left(\varphi \ln\left(S_T\right)\right)\right] && (42) \\ &= S_t^{\phi} \exp\left(A_t + B_{1,t}\left(h_{t+1} - q_{t+1}\right) + B_{2,t} q_{t+1}\right), && (43) \end{aligned}$$

with coefficients

$$A_t = A_{t+1} + r\phi - \left(\alpha B_{1,t+1} + \varphi B_{2,t+1}\right) - 0.5 \ln\left(1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}\right) + \omega B_{2,t+1}, \tag{44}$$

where

$$B_{1,t} = \tilde{\beta} B_{1,t+1} + \lambda\phi + 2\frac{\left(\alpha\gamma_1 B_{1,t+1} + \varphi\gamma_2 B_{2,t+1} - 0.5\phi\right)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}, \tag{45}$$

and

$$B_{2,t} = \rho B_{2,t+1} + \lambda\phi + 2\frac{\left(\alpha\gamma_1 B_{1,t+1} + \varphi\gamma_2 B_{2,t+1} - 0.5\phi\right)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}. \tag{46}$$

The terminal conditions are therefore

$$A_T = B_{1,T} = B_{2,T} = 0. \tag{47}$$

Therefore, $f^{\psi}(t, T : i\phi)$ is the conditional characteristic function under the risk-neutral measure. European call option price, therefore, is the solution of the following equation,

$$\begin{aligned} \text{Call price} &= \frac{1}{2}S_t + \frac{e^{-r(T-t)}}{\pi}\int_0^{\infty} \text{Re}\left[\frac{K^{-i\varphi} f^{\psi}(t, T : i\varphi + 1)}{i\varphi}\right] \times \\ & d\varphi - Ke^{-r(T-1)}\left(\frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} \text{Re}\left[\frac{K^{-i\varphi} f^{\psi}(t, T : i\varphi)}{i\varphi}\right] d\varphi\right) \end{aligned} \tag{48}$$

where Re denotes the real part of a complex number. Note that there is no closed form solution for the option prices for the models with jump components, therefore, options must be priced by Monte Carlo simulation. We shall now consider the risk-neutral dynamics of this price process.

Risk-neutral processes are widely used in the financial literature for valuing contingent claims. The discounted spot asset dynamics under risk-neutral measure represent a martingale, as the expected return is equal to the risk-free rate. Thus, the parameters of the NGARCH and component models could be transformed to the risk-neutral parameters by equalising the expected return to the risk-free rate

$$\mathbb{E}_t^{\psi}\left[S_{t+1}/S_t\right] = e^r, \tag{49}$$

Detailed derivations of risk-neutral parameters of the NGARCH and component models are provided in Heston and Nandi (2000) and Christoffersen et al. (2008) respectively. To summarise, risk-neutral dynamics for the NGARCH model are given by

$$R_{t+1} = r - 0.5h_{t+1} + z_{t+1}^{\psi}\sqrt{h_{t+1}} \tag{50}$$

where

$$h_{t+1} = w + bh_t + \alpha\left(z_t^{\psi} - c^{\psi}\sqrt{h_t}\right)^2, \tag{51}$$

with $c^{\psi} = c + \lambda + 0.5$ and $z_{t+1}^{\psi} \sim \mathcal{N}(0, 1)$ and risk-neutral dynamics for the component model

$$R_{t+1} = r - 0.5h_{t+1} + z_{t+1}^{\psi}\sqrt{h_{t+1}} \tag{52}$$

and

$$h_{t+1} = q_{t+1} + \tilde{\beta}\left(h_t - q_t\right) + \alpha\left(z_{t+1}^{\psi 2} - 1 - 2\gamma_1^{\psi} z_t^{\psi}\sqrt{h_t}\right) \tag{53}$$

$$q_{t+1} = \omega + \rho^{\psi} q_t + \varphi\left(z_t^{\psi 2} - 1 - 2\gamma_2^{\psi} z_t^{\psi}\sqrt{h_t}\right) \tag{54}$$

with contemporaneous terms

$$\tilde{\beta}^{\psi} = \tilde{\beta} + \alpha\left(\gamma_1^{\psi 2} - \gamma_1^2\right) + \varphi\left(\gamma_2^{\psi 2} - \gamma_2^2\right) \tag{55}$$

$$\rho^{\psi} = \rho + \alpha\left(\gamma_1^{\psi 2} - \gamma_1^2\right) + \varphi\left(\gamma_2^{\psi 2} - \gamma_2^2\right) \tag{56}$$

$$\gamma_i^{\psi} = \gamma_i + \lambda + 0.5, \quad i = 1, 2. \tag{57}$$

All models with Poisson-normal jumps there are three stochastic processes in the dynamics. Deriving the risk-neutral forms is often somewhat problematic as noted in Christoffersen et al. (2008). It may be shown that risk-neutral return and variance dynamics for the simplest Constant-Jump model are defined as

$$R_{t+1} = r - 0.5h_{t+1} + \left(\lambda_y^{\psi} - \mu^{\psi}\right)\eta^{\psi} + z_{t+1}^{\psi}\sqrt{h_{t+1}} + y_{t+1}^{\psi}, \tag{58}$$

and

$$h_{t+1} = w + bh_t + \alpha\left(z_t^{\psi} + \frac{y_t^{\psi}}{\sqrt{h_t}} - c^{\psi}\sqrt{h_t}\right)^2, \tag{59}$$

where $z_t^{\psi} \sim \mathcal{N}(0,1)$ and $y_t^{\psi} \sim \text{CPoisson}\left(\eta^{\psi}, \mu^{\psi}, \tau^2\right)$, with $c^{\psi} = c + \lambda_t + 0.5$, $\mu^{\psi} = \mu + \pi_y\tau^2$, $(\lambda_y^{\psi} - \mu^{\psi})\eta^{\psi} = (\lambda_y - \mu)\eta$ and $\eta^{\psi} = \eta\exp(\pi_y\mu + 0.5\pi_y^2\tau^2)$. Elkamhi and Ornthanalai (2010) show that the market price of jump risk π_y can be solved for numerically from the equation

$$\lambda_y = \zeta(1) + \zeta(-\pi_y) - \zeta(1 - \pi_y), \tag{60}$$

where λ_y is a ML estimate of the physical process and $\zeta(\phi) = \exp(\phi\mu + 0.5\phi^2\tau^2) - 1$. As in case of the NGARCH and component models, risk-neutral dynamics are achieved by setting the diffusive risk premium λ_z equal to -0.5 and the jump risk premium λ_y equal to

$\mu - \exp(\mu + 0.5\tau^2) + 1$ (see Appendix for the formal derivation). Risk-neutral parameters of the AR-Jump model are defined in the same way as of the Constant-Jump model except for conditional jump intensity η_t^{ψ}, which for the sake of simplicity is assumed to be equal to the product of conditional jump intensity and expected jump size from the physical process $\eta_t \exp(\pi_y \mu + 0.5\pi_y^2 \tau^2)$.

Option prices for the component GARCH model are taken from the closed-form solution provided above. Monte Carlo simulation is used to find option values suggested by the AR-Jump model. Due to computational burden of Monte Carlo simulation with compound Poisson-normal process, it is impossible to generate the ideal number of iterations. The balance between accuracy and required time is achieved with 100,000 iterations.[1] However, several previous studies have suggested that the higher number of iterations can significantly improve the accuracy of estimation. We experimented with single contracts up to 1,000,000 iterations and have found convergence usually occurs at between 80,000 and 120,000.

Data

Prices for European call options on the Standard and Poor's 500 index (the SPX index option) are collected from the Thomson-Reuters EIKON data service. The data represents Thursday closing prices during the 2008 crisis period and consists of 2,134 contracts. We follow the logic presented in Christoffersen et al. (2008), Dumas et al. (1998) and Heston and Nandi (2000) and use Thursday data because (a) it is less likely to be a holiday and (b) Thursday is less likely to be affected by day-of-the-week effects.

The SPX index option contract is the second most active index options market in the United States and is regularly interpreted, for that reason, as the most appropriate market for examining option valuation performance (see for example, Heston and Nandi (2000). Rubinstein (1994) indeed argues that it is the best market for examining option valuation performance. We select option prices that are taken with up to 10 weeks to maturity in order to test our models performance in valuing short maturity speculative positions.

[1] It takes up to 1.5 hour to estimate one month of option prices by Monte Carlo simulation with 100,000 iterations

Descriptive statistics for the options data for the 2008 are presented in Table 5 (Panel A). Option contracts are divided by moneyness and time to maturity. More than half of the data represent deep-in-the-money call options and one quarter are deep out-of-the-money contracts. Therefore, models' performance for these kinds of contracts can be assessed with a reasonable degree of reliability. It is noted that at-the-money contracts which comprise the main source of analysis represents 10% of the sample.

Empirical results from the call option valuation

The results on option valuation by moneyness and maturity are presented in Table 6. The reported root mean squared errors (RMSE) are obtained from the closed-form solution for the component GARCH and Monte Carlo simulation for the AR-Jump model as described above. Risk-neutral parameters are calculated from ML estimates for the two models reported in Tables 1 and 2 respectively. In the component GARCH model the risk-neutral dynamic is modelled by using $\lambda = -0.5$. In the AR-Jump model this is achieved by taking $\lambda = -0.5$ and $\lambda_y = \mu - \exp(\mu + 0.5\tau^2) + 1 = -3.313E - 5$.

Panel B in Table 5 reports the RMSE of the component GARCH model. The results are potentially affected by the above noted shortcomings in the data with the resulting RMSE values being too large to make any confident conclusions. All the same, we can learn some lessons from the patterns that are evident. The component GARCH appears to perform better in valuing the longest maturity options (more than 45 days to maturity). As expected, the lowest value of the RMSE corresponds to deeply out-of-the-money contracts. The RMSE for at-the-money options, which are defined hereafter as contracts with the strike prices within 2.5% of the underlying index values, are high, which might be caused by the relatively small sample.

[INSERT TABLES 5 AND 6 ABOUT HERE]

Panel C displays the ratio of the AR-Jump RMSE to the component GARCH results. The total RMSE values do not differ significantly although the AR-Jump model does appear

to perform well in valuing short maturity options compared to the component GARCH. We note that the AR-Jump model performs best in valuing deeply out-of-the-money and at-the-money options with less than 25 days to maturity. Furthermore, results for the 25 to 45 days to maturity options are better than the component GARCH model. The RMSE for the AR-Jump model for the long maturity options are clearly inferior and higher than the component GARCH results at all moneyness levels. Comparing the results by moneyness, the AR-Jump model surpasses the component GARCH model in pricing at-the-money contracts.

The performance of the models can also be assessed by comparing the ability of the model(s) to match the market volatility pattern. To analyse this we impute the Black and Scholes implied volatility for at-the-money options. The implied volatility is calculated for each week's market price, the component GARCH price and the AR-Jump price. Figure 8 depicts the average weekly implied volatility bias for the two models, which is the difference between average market implied volatility and average model implied volatility for at-the-money call contracts during 2008. Parameter estimates are obtained from Tables 1 and 2. During 2008 the average value of at-the-money option-implied volatility index (VIX) is 22.3% over the first eight months and 48.0% for the rest of the year. Therefore, we refer to the period from January to August as a low volatility period and the next four months as a high volatility period.

[INSERT FIGURE 8 ABOUT HERE]

Both models show significant underpricing (positive bias) in the first fifteen weeks. Thereafter, the volatility implied by the AR-Jump model stays close to the market-implied volatility, while the component GARCH model shows overpricing (negative bias) during the summer months with an extreme drop in the middle, which may be caused by the data shortcomings discussed above. Overall, both models perform in a somewhat similar way during the low volatility period. On the other hand, the models generate quite different paths during the high volatility period models and it is clear that the AR-Jump model better matches the spot volatility. As in Christoffersen et al. (2008) the component GARCH

model shows prolonged underpricing during a high volatility period. Notably, however, the implied volatility bias from the AR-Jump model is generally lower than the one implied by component GARCH and even drops below zero in two weeks.

5. Conclusion

This work analyses the component GARCH model proposed by Christoffersen et al. (2008) and presents an extension to their GARCH(1,1)-Jump model with Poisson-normal jumps in returns and variance, and time-varying jump arrival intensity. The component GARCH is found to outperform the Heston and Nandi (2000) NGARCH model and the fully persistent in modelling daily total returns on four indices during the 1988-2010 period. Although variance persistence in the component GARCH model is close to 1, it is essential that it is not modelled as fully persistent. Anecdotal evidence on poor performance of the persistent GARCH modification is often attributed to the lack of flexibility in the intensity dynamics, which significantly worsen its results during the period of extreme volatility in 2008-2009. The AR-Jump model is found to achieve the best fit to the daily historical returns, however this has eleven parameters versus the component GARCH's eight. Strong evidence of a jump effect is found in all four series. Moreover, the impressive performance of the AR-Jump model strongly suggests that jumps tend to cluster together. During the latter part of 2008 we observe extreme clustering of jumps after the Lehmen Brothers bankruptcy. The model adapts to this exited state and is able to capture the change in market conditions in a remarkably adept manner.

References

Andersen, T. G., L. Benzoni, and J. Lund (2002). Empirical investigation of continuous-time equity return models. *The Journal of Finance 57*(3), 1239–1284.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *The Journal of Finance 52*(5), 2003–2049.

Bates, D. S. (1991). The crash of '87: Was it expected? The evidence from the options markets. *The Journal of Finance 46*(3), 1009–1044.

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

Bates, D. S. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics 94*(1-2), 181–238.

Bates, D. S. (2006). Maximum likelihood estimation of latent affine processes. *The Review of Financial Studies 19*(3), 909–965.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Broadie, M., M. Chernov, and M. Johannes (2007). Model specification and risk premiums: evidence from futures options. *The Journal of Finance 62*(3), 1453–1490.

Chan, W. H. and J. M. Maheu (2002). Conditional jump dynamics in the stock market returns. *Journal of Business and Economic Statistics 20*(3), 377–389.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative models for stock price dynamics. *Journal of Econometrics 116*(1-2), 225–257.

Christoffersen, P., S. Heston, and K. Jacobs (2006). Option valuation with conditional skewness. *Journal of Econometrics 131*(1-2), 253–284.

Christoffersen, P., K. Jacobs, C. Ornthanalai, and Y. Wang (2008). Option valuation with long-run and short-run volatility components. *Journal of Financial Economics 90*(3), 272–297.

Devroye, L. (1982). A log log law for maximal uniform spacings. *The Annals of Probability 10*(3), 863–868.

Dumas, B., J. Fleming, and R. E. Whaley (1998). Implied volatility functions: empirical tests. *The Journal of Finance 53*(6), 2059–2106.

Elkamhi, R. and C. Ornthanalai (2010). Market jump risk and the price structure of individual equity options. Technical report, WFA 2010 Victoria meetings.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of uk inflations. *Econometrica 50*(4), 987–1008.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *The Journal of Finance 59*(3), 1367–1403.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. *The Journal of Finance 53*(3), 1269–1300.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance 48*(5), 1779–1801.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*(2), 327–43.

Heston, S. L. and S. Nandi (2000). A closed-form GARCH option valuation model. *Review of Financial Studies 13*(3), 585–625.

Jorion, P. (1998). On jump processes in the foreign exchange and stock markets. *Review of Financial Studies 1*(4), 427–445.

Maheu, J. M. and T. H. McCurdy (2004). News arrival, jump dynamics, and volatility components for individual stock returns. *The Journal of Finance 20*(2), 755–793.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(1), 347–370.

Palm, F. C. and P. J. G. Vlaar (1993). The message in weekly exchange rates in the European monetary system: Mean reversion, conditional heteroscedasticity, and jumps. *Journal of Business & Economic Statistics 11*(3), 351–360.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics 63*(1), 3–50.

Rubinstein, M. (1994). Implied binomial trees. *The Journal of Finance 49*(3), 778–818.

Tables and Figures

Table 1: This table gives Maximum likelihood (ML) estimates of the three models, using daily total return data from 1988-2010. Parameters are estimated for the Standard and Poors 500 index from January 4, 1988 to December 31, 2010. Estimate that are significant at the 5% significance level are reported with a *. Persistence indicates the persistence of the conditional variance. Average annual volatility refers to the mean annualised daily standard deviation during 1988-2010. Average volatility of variance refers to the average standard deviation of the conditional variance during 2005-2010. Average correlation indicates the average correlation between the return and the conditional variance during 2005 to 2010.

NGARCH			Component NGARCH			Persistent NGARCH		
Parameter	Estimate	Standard error	Parameter	Estimate	Standard error	Parameter	Estimate	Standard error
λ	1.747*	0.003317	λ	1.59*	0.01351	λ	1.762*	0.001027
w	1.039E-11	0.7939	$\tilde{\beta}$	0.6673*	0.002311	$\tilde{\beta}$	0.8852*	0.001798
b	0.8569*	0.001308	α	3.271E-07	0.0149	α	6.652E-07	0.004269
a	3.203E-06	0.001174	γ_1	2729*	0.0007812	γ_1	1332*	0.00009094
c	187.7*	0.0007022	ω	1.516E-06	0.001497	ω	7.643E-08	0.005384
			ρ	0.9861*	0.001554	ρ	1	
			φ	2.803E-06	0.001062	φ	2.014E-06	0.0004017
			γ_2	103.8*	0.000342	γ_2	1.012E-11	0.00009596
Ln likelihood	19481		19566			19357		
Persistence	0.9698		0.9954			1		
Average annual volatility	0.1562		0.1569			0.169		
Average volatility of variance	0.00001288		0.000024			0.00001404		
Average correlation	-0.9174		-0.9741			-0.9671		

Table 2: This table shows maximum likelihood estimates of the Constant-Jump and AR-Jump models. Estimates of the parameters are obtained using daily total returns on the Standard and Poor's 500 index from January 4, 1988 to December 31, 2010 from the Dow Jones Industrial Average (DJIA), the Deutscher Aktien Index (DAX 30) and the FTSE 100 indices. Parameters that are significant at the 5% significance level are reported with a *.

	Constant-Jump			AR-Jump	
Parameter	Estimate	Standard error	Parameter	Estimate	Standard error
λ_z	5.929*	0.009218	λ_z	3.083*	0.005526
λ_y	-0.01547*	0.007288	λ_y	-1.634E-12	5.321
w	-9.932E-07	0.002992	w	-4.302E-17	18.08
b	0.8933*	0.001643	b	0.979*	0.003658
a	2.583E-06	0.001098	a	3.064E-07	0.002515
c	179*	0.00002709	c	135.3*	0.0001447
η_0	0.0247*	0.004493	η_0	0.0266*	0.002828
μ	-0.01241*	0.004113	ρ	0.9713*	0.002297
τ	0.01879*	0.001799	γ	0.8608*	0.009808
			μ	-0.004641*	0.00171
			τ	0.006705*	0.002265
Ln likelihood	19 675		19 783		

Table 3: This table shows ML estimates of the component GARCH model parameters for daily returns on three indices from 1988-2011. Estimates of the parameters are obtained using daily total returns from January 4, 1988 to December 31, 2010 on the Dow Jones Industrial Average (DJIA), the Deutscher Aktien Index (DAX 30) and the FTSE 100 indices. Results that are significant at the 5% significance level are reported with a *.

	DJIA		DAX30		FTSE100	
Parameter	Estimate	Standard error	Estimate	Standard error	Estimate	Standard error
λ	2.041*	0.01736	0.8354*	0.008361	1.758*	0.01505
$\tilde{\beta}$	0.6774*	0.002489	0.6238*	0.004278	0.941*	0.006124
α	2.438E-07	0.01794	3.754E-06	0.009018	2.367E-06	0.005406
γ_1	3383*	0.001181	209.5*	0.00003462	124.9*	0.00001346
ω	1.616E-06	0.001543	3.802E-06	0.002335	9.902E-07	0.003756
ρ	0.9844*	0.001581	0.979*	0.002421	0.9907*	0.003765
φ	2.932E-06	0.0009801	4.517E-06	0.002363	1.852E-06	0.005322
γ_2	101.4*	0.0003313	102.6*	0.0003509	94.78*	0.0007927
Ln likelihood	19667		17952		19582	

Table 4: This table shows ML estimates of the AR-Jump model parameters for the three indices' daily returns, 1988-2011. Estimates of the parameters are obtained using daily total total returns from January 4, 1988 to December 31, 2010 on the Dow Jones Industrial Average (DJIA), the Deutscher Aktien Index (DAX 30) and the FTSE 100 indices. Results that are significant at the 5% significance level are reported with a *.

	DJIA		DAX30		FTSE100	
Parameter	Estimate	Standard error	Estimate	Standard error	Estimate	Standard error
λ_z	3.629*	0.00575	2.454*	0.004563	3.052*	0.00581
λ_y	-1.635E-12	27.26	-1.635E-12	6.521	-1.635E-12	3.448
w	-4.302E-17	41.39	-4.302E-17	10.55	-4.302E-17	2.148
b	0.9777*	0.00346	0.9793*	0.002791	0.9894*	0.002248
a	3.398E-07	0.002595	4.075E-07	0.002559	2.788E-07	0.002835
c	143.4*	0.0001273	121.8*	0.0001147	9.943*	0.0001262
η_0	0.02964*	0.002312	0.02707*	0.002901	0.02485*	0.002264
ρ	0.9654*	0.002342	0.9698*	0.002332	0.9769*	0.002561
γ	0.9999*	0.4766	0.8409*	0.0105	0.9999*	0.3896
μ	-0.003805*	0.001508	-0.004573*	0.001667	-0.003562*	0.001395
τ	0.006958*	0.001029	0.009137*	0.0009207	0.006212*	0.001042
Ln likelihood	19827		18233		19719	

Table 5: This table shows call option data and the root mean squared error (RMSE) plus the ratio RMSE by moneyness and maturity, 2008. The ML estimates from Tables 1 and 2 are used to compute the root mean squared error (RMSE) by moneyness and maturity during 2008. Option prices for the component GARCH model are calculated from the closed-form solution defined in the text, the Monte Carlo simulation with 100 000 iterations is used to calculate prices for the AR-Jump model. The options data in Panel A represents weekly prices of European call options on the Standard and Poors 500 index. Panel B shows the RMSE for the component GARCH and Panel C shows the ratio of RMSE for the AR-Jump model over RMSE for the component GARCH. DTM refers to the number of days to maturity and S/K refers to the ratio of underlying index level to strike price.

Moneyness	$S/K < 0.975$	$0.975 < S/K < 1$	$1 < S/K < 1.025$	$1.025 < S/K < 1.05$	$1.05 < S/K < 1.075$	$1.075 < S/K$	Total
Panel A. Number of call option contracts							
$DTM < 25$	207	40	38	36	39	335	695
$25 < DTM < 45$	173	39	35	38	28	338	651
$DTM > 45$	152	50	43	49	40	454	788
Total	532	129	116	123	107	1127	2134
Panel B. Component GARCH RMSE							
$DTM < 25$	2.72	8.79	7.45	11.53	11.43	7.01	8.69
$25 < DTM < 45$	4.72	8.82	9.43	12.02	5.83	11.96	9.22
$DTM > 45$	4.93	6.89	9.64	10.64	7.75	8.68	8.3
Total	4.24	8.22	8.89	11.41	8.65	9.44	8.75
Panel C. Ratio of AR-Jump to component GARCH RMSE							
$DTM < 25$	0.67	0.59	0.7	0.72	0.89	0.9	0.77
$25 < DTM < 45$	1.04	0.79	0.79	0.99	1.18	0.74	0.88
$DTM > 45$	1.57	1.25	1.1	1.16	1.66	1.29	1.29
Total	1.27	0.86	0.91	0.96	1.19	0.96	0.99

Table 6: Data for the European call options on the Standard and Poor's 500 index prices are taken at 6 p.m. (CST) on each Thursday during 2008.

Moneyness	$S/K < 0.975$	$0.975 < S/K < 1$	$1 < S/K < 1.025$	$1.025 < S/K < 1.05$	$1.05 < S/K < 1.075$	$1.075 < S/K$	Total
Maturity				Number of open call option contracts			
$DTM < 25$	207	40	38	36	39	335	695
$25 < DTM < 45$	173	39	35	38	28	338	651
$DTM > 45$	152	50	43	49	40	454	788
Total	532	129	116	123	107	1,127	2,134

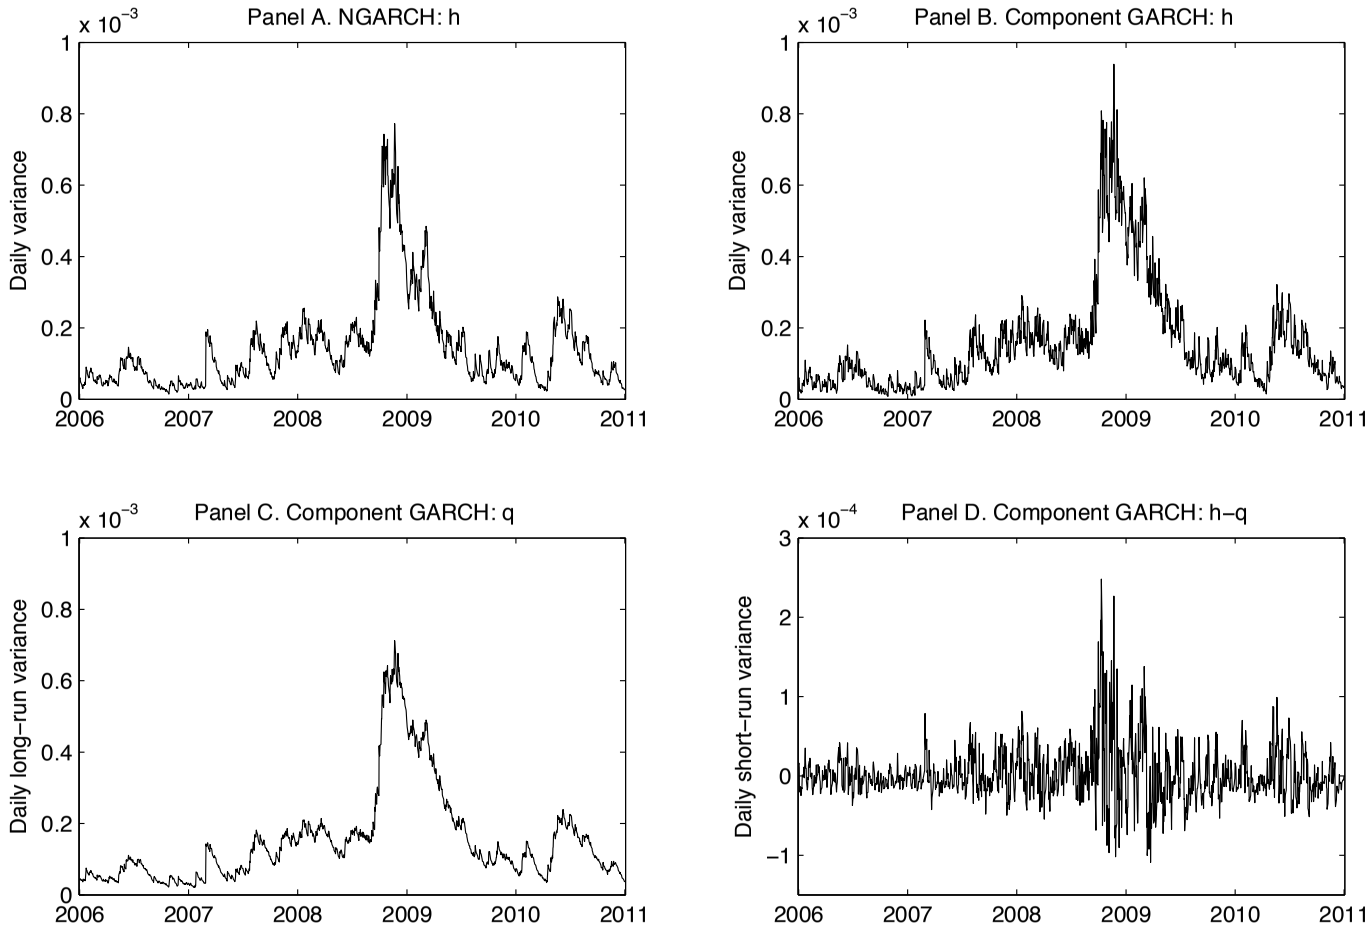

Figure 1: These panels demonstrate the spot variance of the component GARCH and NGARCH models. The top panels depict the dynamics for the conditional variance while the bottom panels show the long-run (left-hand panel) and the short-run (right-hand panel) components. The scales are identical across the rows of Panel A, Panel B and Panel C. The parameter values are taken from the maximum likelihood estimates reported in Table 1.

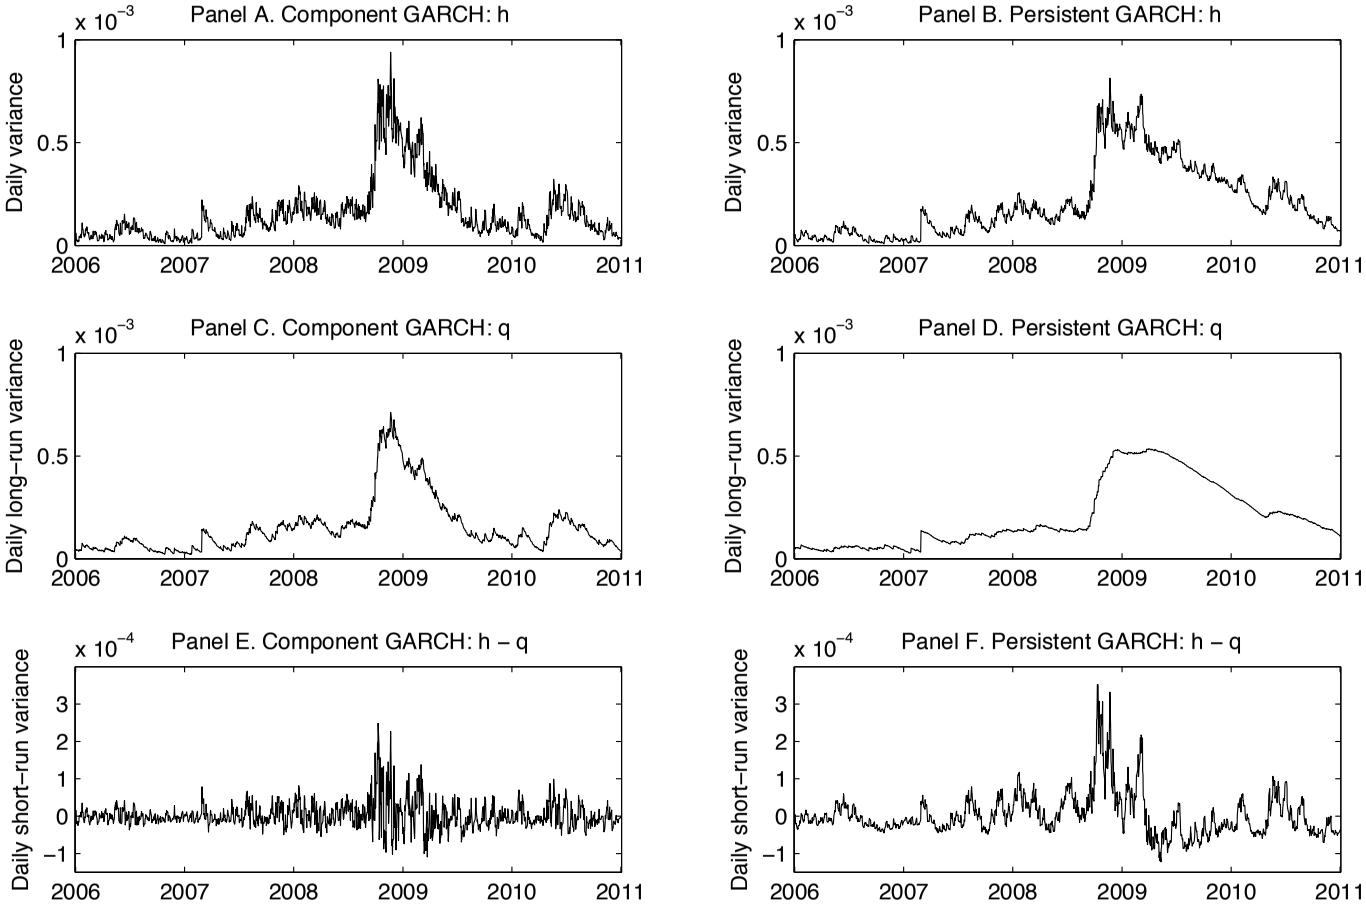

Figure 2: These panels show the spot variance of the component and persistent models. The left-hand panels depict the conditional variance (top panel), the long-run component (centre panel) and the short-run component (bottom panel) paths from the component model. The right-hand panels depict respective paths from the persistent model. The parameter values are taken from the maximum likelihood estimates reported in Table 1.

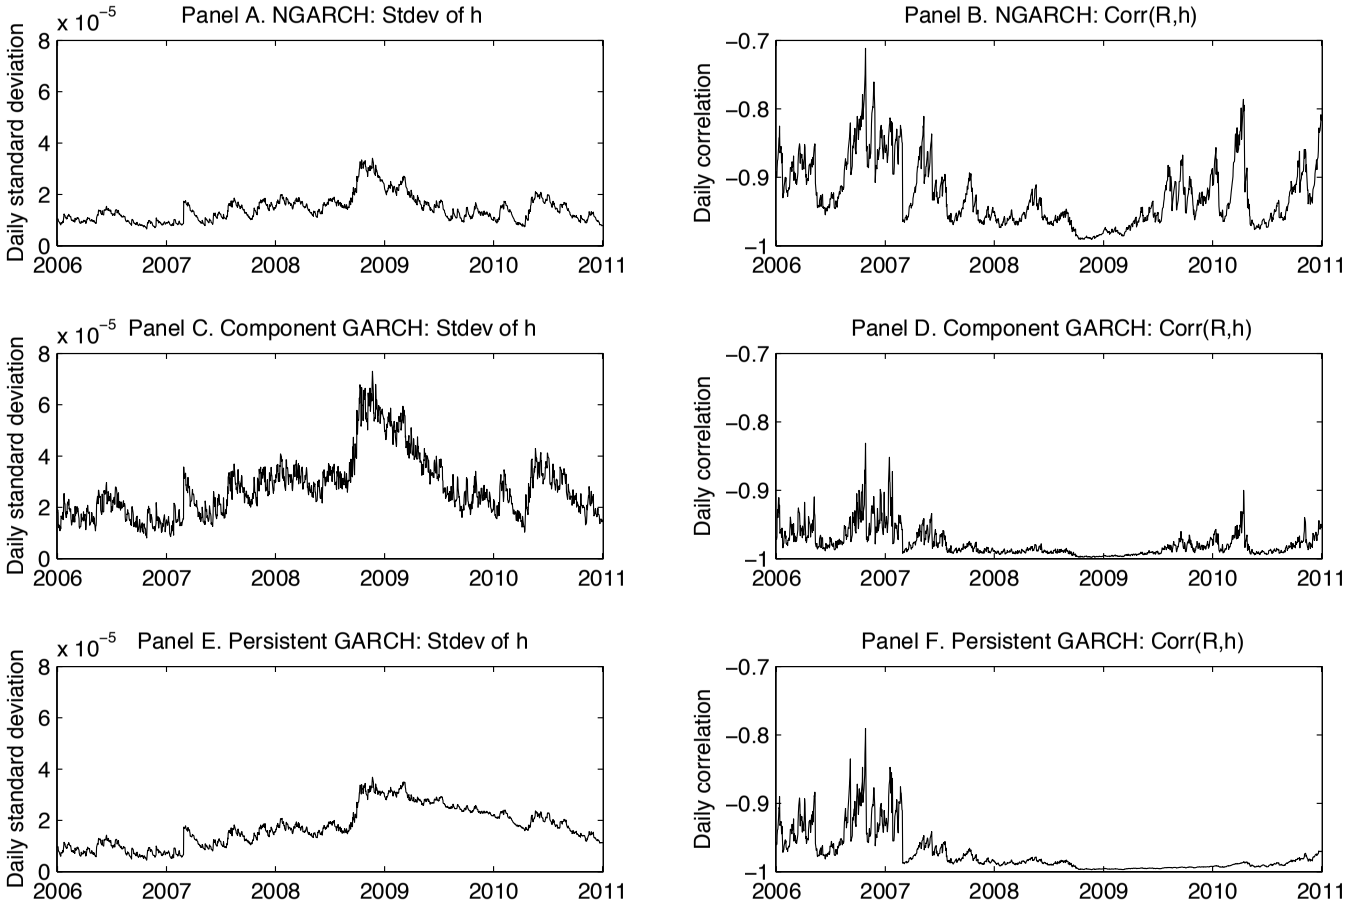

Figure 3: These panels demonstrate the conditional standard deviation of variance and conditional correlation. The left-hand panels depict the conditional standard deviation of conditional variance from the NGARCH (top panel), the component GARCH (centre panel) and the persistent GARCH (bottom panel) models. The right-hand panels depict the conditional correlation between the return and variance for the three models, respectively. The parameter values are taken from the maximum likelihood estimates reported in Table 1.

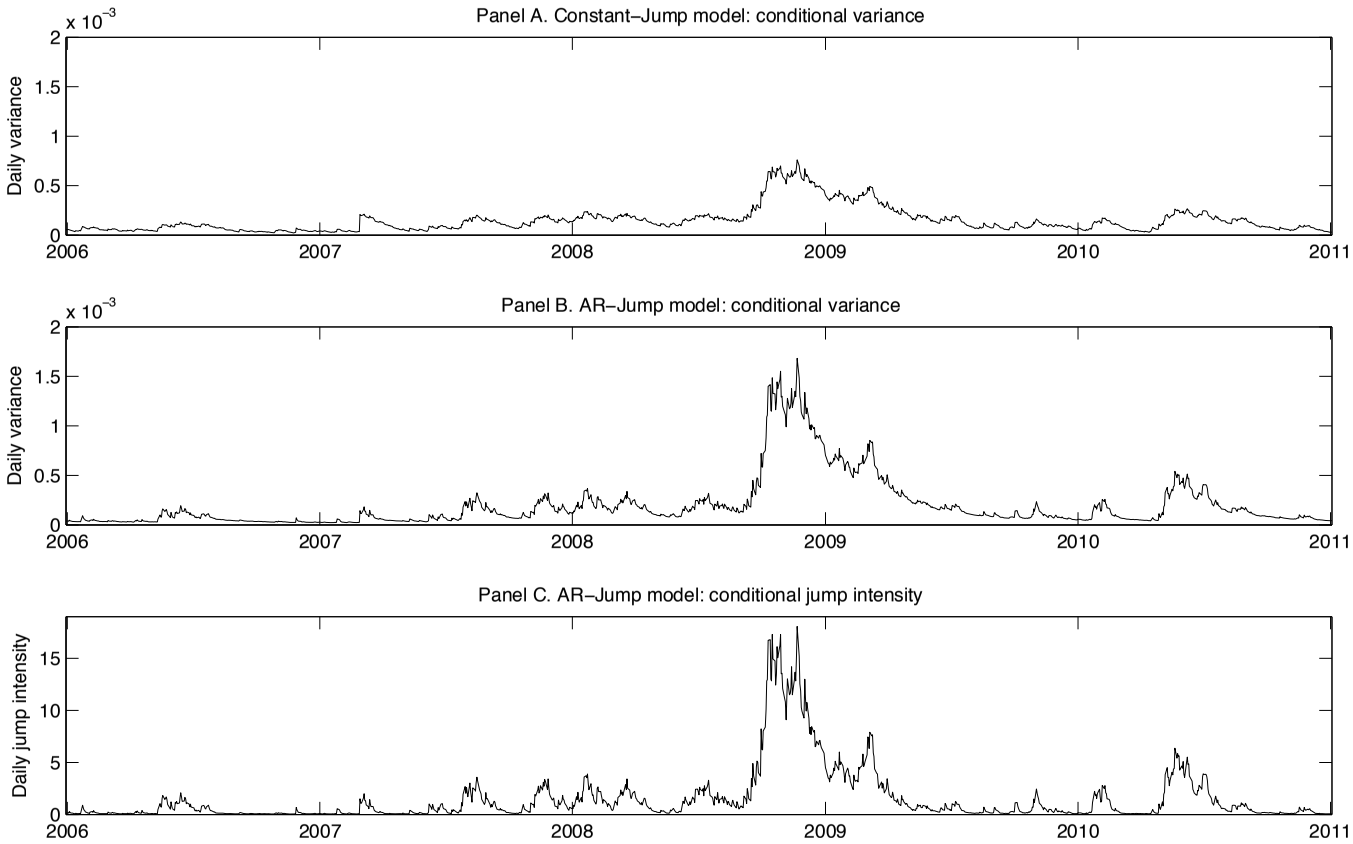

Figure 4: These panels show the spot variance of the Constant-Jump and AR-Jump models. Panels A and B depict the conditional variance from the Constant-Jump and AR-Jump models, respectively. Panel C depicts the conditional jump arrival intensity. The parameter values are taken from the maximum likelihood estimates reported in Table 2.

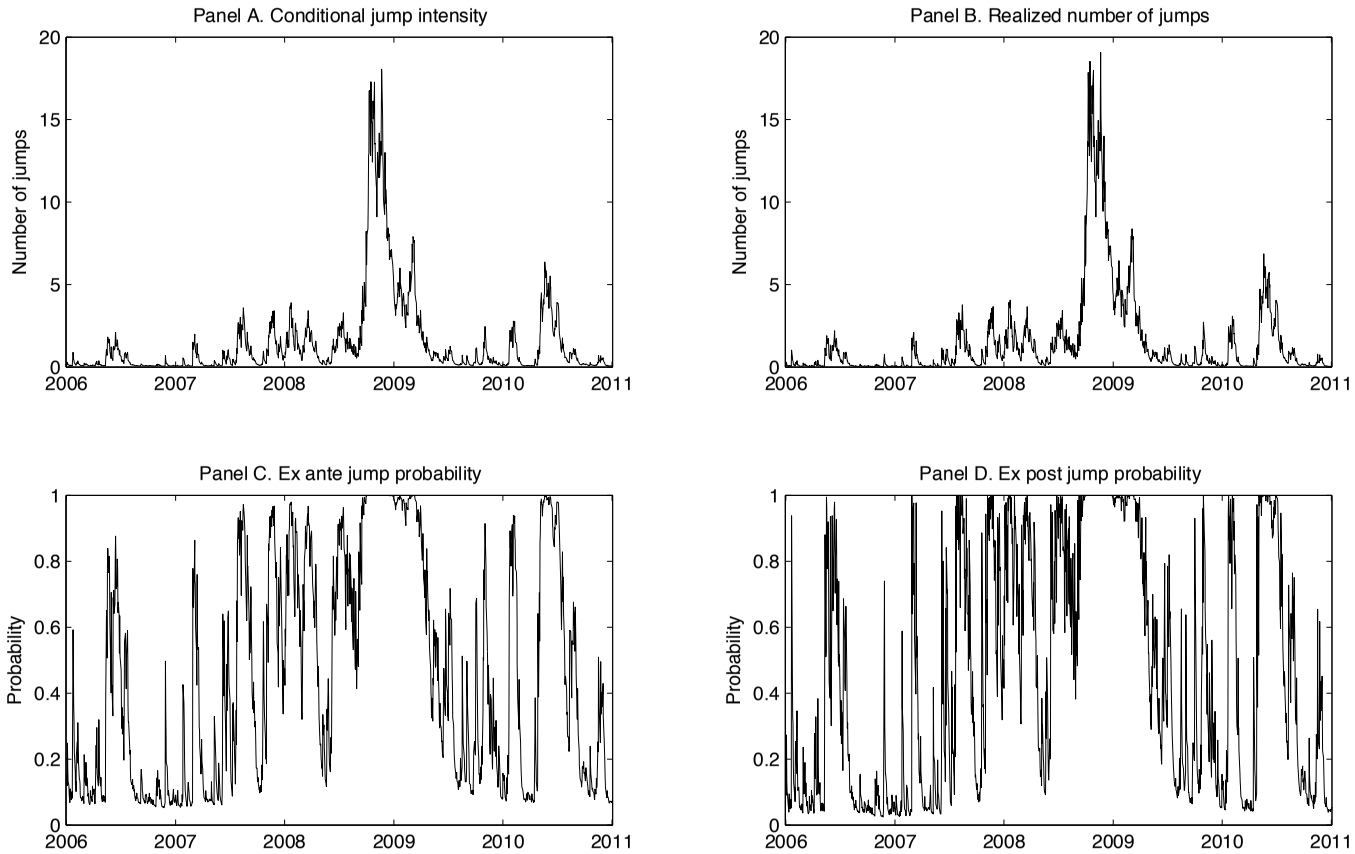

Figure 5: These panels demonstrate the conditional jump intensity and the realized number of jumps for ex ante and ex post probabilities of jump from the AR-Jump model. The top panels show the conditional jump intensity (left-hand panel) and the realized number of jumps (right-hand panel). The bottom panels depict the ex ante (left-hand panel) and ex post (right-hand panel) jump probabilities. The parameter values are taken from the maximum likelihood estimates reported in Table 2.

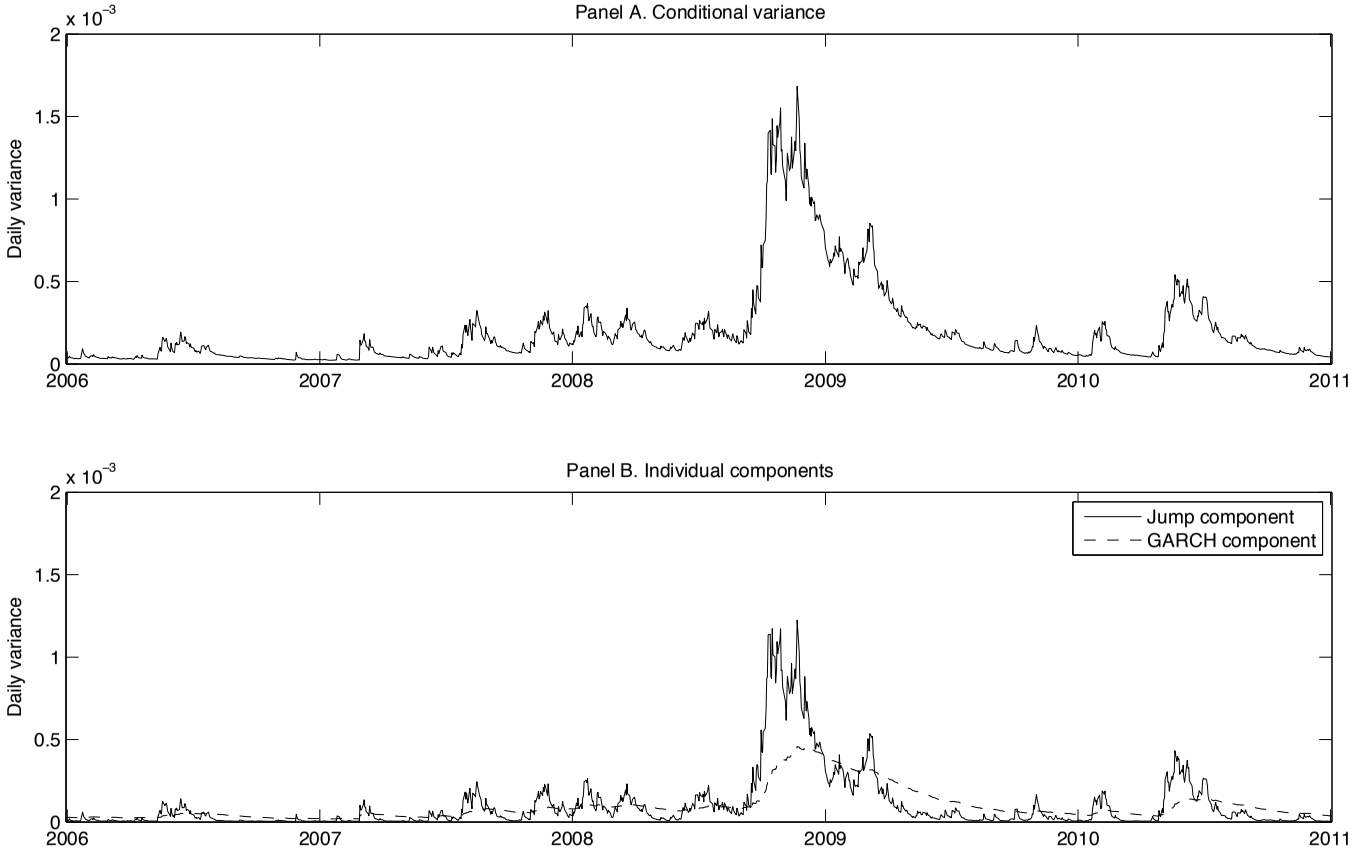

Figure 6: These panesl show the spot variance decomposition in AR-Jump model. The top panel depicts the conditional variance path from the AR-Jump model and the bottom panel depicts the individual component paths. The parameter values are taken from the maximum likelihood estimates reported in Table 2.

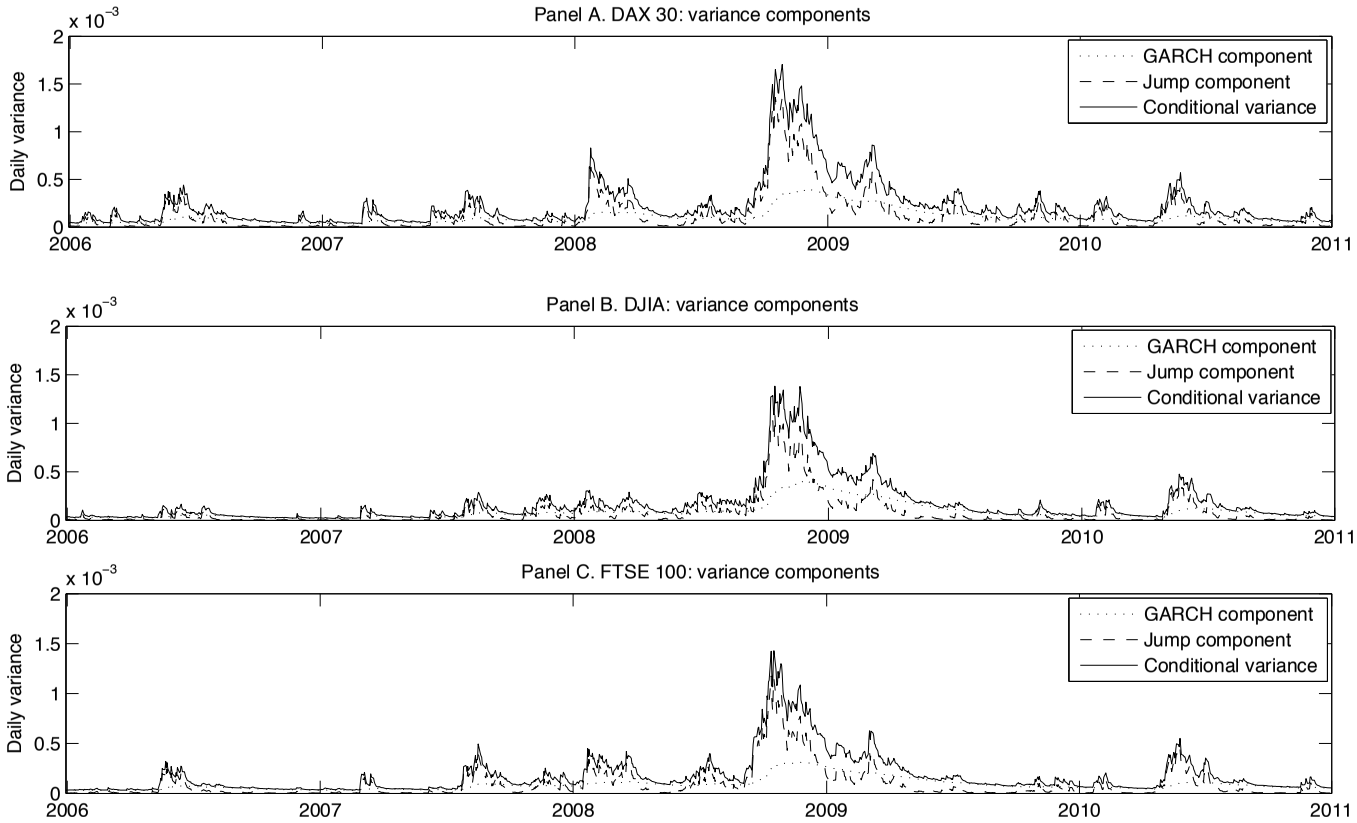

Figure 7: These panels show spot variance paths from the AR-Jump model. The panels depict the spot variance and the individual variance components of the total returns on three indices: DAX (top panel), DJIA (centre panel) and FTSE 100 (bottom panel). The parameter values are taken from the maximum likelihood estimates reported in Table 4.

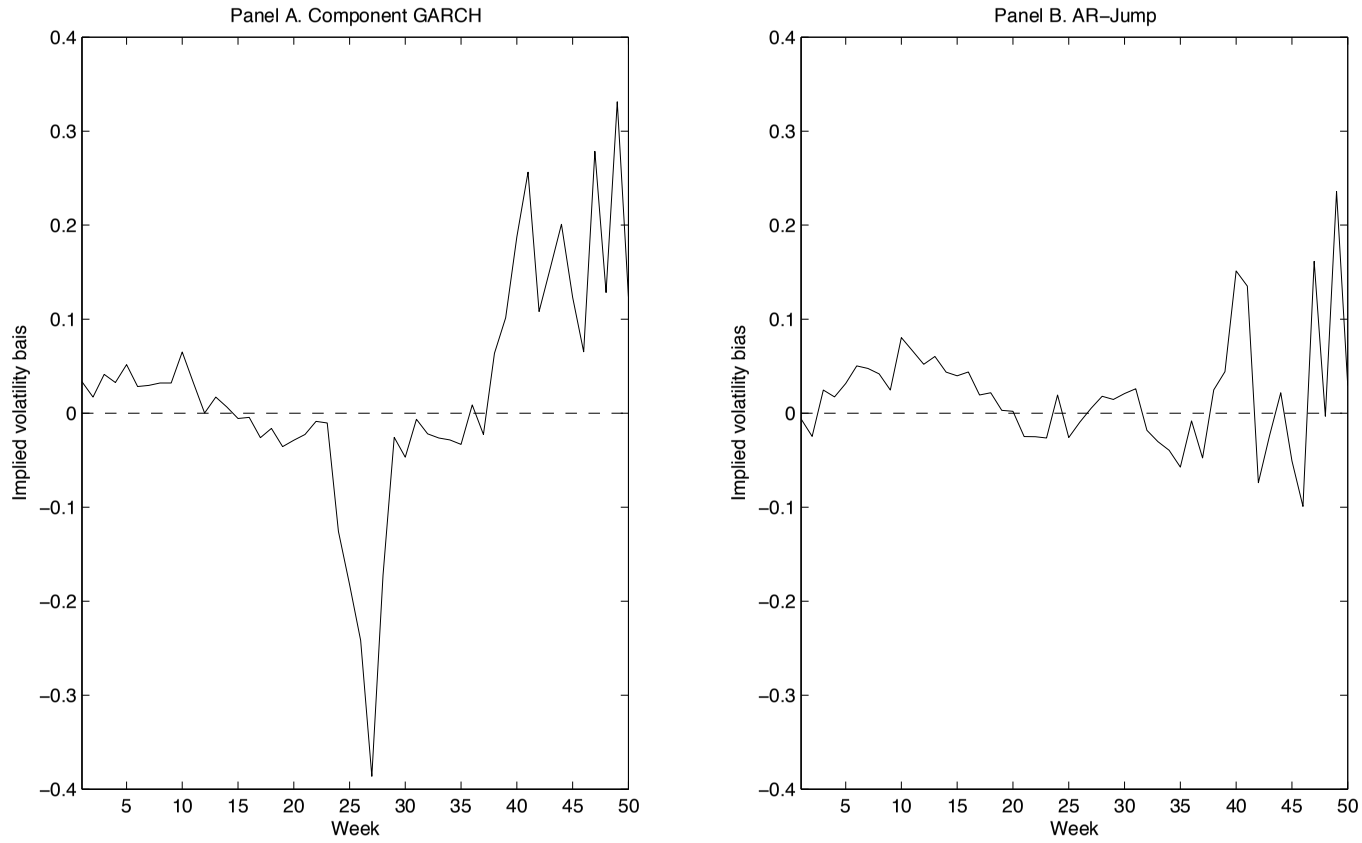

Figure 8: Panel A shows the average weekly volatility bias for at-the-money options (the difference between the market and implied volatility values)implied by the component GARCH model. Panel B depicts the average weekly volatility bias implied by the AR-Jump model. Black and Scholes implied volatility for at-the-money options are calculated as contracts with a strike price within 2.5% of underlying index. The parameter values are taken from the maximum likelihood estimates reported in Tables 1 and 2.