JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ. **17**: 535–548 (2002)
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/jae.685

BRIDGING THE GAP BETWEEN THE DISTRIBUTION OF REALIZED (ECU) VOLATILITY AND ARCH MODELLING (OF THE EURO): THE GARCH-NIG MODEL

LARS FORSBERG[a]* AND TIM BOLLERSLEV[b]

[a] *Department of Information Science, Division of Statistics, Uppsala University, Sweden*
[b] *Department of Economics and Finance, Duke University, NC, USA and NBER*

SUMMARY

This paper bridges the gap between traditional ARCH modelling and recent advances on realized volatilities. Based on a ten-year sample of five-minute returns for the ECU basket currencies versus the US dollar, we find that the realized volatilities constructed from the summation of the high-frequency intraday squared returns conditional on the lagged squared daily returns are approximately Inverse Gaussian (IG) distributed, while the distribution of the daily returns standardized by their realized volatilities is approximately normal. Moreover, the implied daily GARCH model with Normal Inverse Gaussian (NIG) errors estimated for the ECU returns results in very accurate out-of-sample predictions for the three years of actual daily Euro/US dollar exchange rates. Copyright © 2002 John Wiley & Sons, Ltd.

1. INTRODUCTION

The seminal ARCH paper by Engle (1982) triggered one of the most active and fruitful areas of research in econometrics over the past two decades. The success of the ARCH/GARCH class of models at capturing volatility clustering in financial markets is well documented (see, for example, Bollerslev, Chou and Kroner, 1992). Meanwhile the inability of the ARCH/GARCH models coupled with the auxiliary assumption of conditionally normally distributed errors to fully account for all of the mass in the tails of the distributions of, say, daily returns is also well recognized. Indeed, several alternative error distributions were proposed in the early ARCH literature to better account for the deviations from normality in the conditional distributions of the returns, including the *t*-distribution of Bollerslev (1987), the General Error Distribution (GED) of Nelson (1991), and more recently, the Normal Inverse Gaussian (NIG) distribution of Barndorff-Nielsen (1997), Andersson (2001) and Jensen and Lunde (2001). The motivation behind these alternative error distributions has been almost exclusively empirical and pragmatic in nature. In the present paper, building on the Mixture-of-Distributions-Hypothesis (MDH) (Clark, 1973) along with the recent idea of so-called Realized Volatilities (RV) (Andersen *et al.*, 2001b, 'Modeling and forecasting realized volatility', unpublished manuscript, 2002; Barndorff-Nielsen and Shephard,

* Correspondence to: Lars Forsberg, Department of Information Science, Division of Statistics, Uppsala University, Sweden. E-mail: Lars.Forsberg@dis.uu.se
Contract/grant sponsor: The Swedish Foundation for International Cooperation in Research and Higher Education (STINT); Contract/grant number: IG 2002-02 045.
Contract/grant sponsor: Jan Wallanders and Tom Hedelius Foundation; Contract/grant number: J 01-28.
Contract/grant sponsor: NSF.

Received 31 October 2001
Revised 21 May 2002

2001a,b, 2002a; Comte and Renault, 1998), we provide a sound empirical foundation for the distributional assumptions underlying the GARCH-NIG model.

Consistent with the absence of arbitrage and a time-changed Brownian motion (see, for example, Ane and Geman, 2000; Andersen, Bollerslev, and Diebold, 2002), the MDH postulates that the distribution of returns is normal, but with a stochastic (latent) variance. In the original formulation in Clark (1973), the variance is assumed to be *i.i.d.* lognormally distributed, resulting in a lognormal-normal mixture distribution for the returns. Numerous theoretical extensions and empirical investigations of these ideas, involving various proxies for the mixing variable have been advanced in the literature (important early contributions include Epps and Epps, 1976; Taylor, 1982; Tauchen and Pitts, 1983). Importantly, to explicitly account for the volatility clustering effect Taylor (1982, 1986) proposed an extension of the MDH setup by making the (latent) logarithmic variances follow a Gaussian autoregression, resulting in the lognormal Stochastic Volatility (SV) model; see also Andersen (1996). Since the joint distribution of the returns in the SV model is not known in closed form, both estimation and inference for these types of models are considerably more complicated than for the ARCH/GARCH class of models (see, e.g., Shephard, 1996).

In contrast to the existing SV literature, which treats the mixing variable as latent, we show that by measuring the daily variance by the corresponding realized volatility constructed from the sum of intraday high-frequency returns, the daily return standardized by the realized volatility is approximately normally distributed. Therefore, even though the realized volatilities are subject to measurement error *vis-à-vis* the true daily latent volatilities (see, for instance, Andreou and Ghysels, 2002; Barndorff-Nielsen and Shephard, 2001b, 2002b; Meddahi, 'A theoretical comparison between integrated and realized volatilities', unpublished manuscript, 2002), the (approximate) normality of the standardized returns is consistent with the basic tenets of the MDH and the use of the realized volatility as the underlying mixing variable. Moreover, we find that the distribution of the realized volatility conditional on the past squared daily returns (as well as the unconditional distribution of the realized volatility) is closely approximated by an Inverse Gaussian (IG) distribution. Taken together, these results imply that in practical modelling situations where the high-frequency data are not actually available, the daily returns should be well described by a GARCH-NIG model.

The empirical analysis outlined above is based on a ten-year sample of high-frequency five-minute returns for the ECU basket of currencies versus the US dollar, spanning the period from 3 January 1989 to 30 December 1998. In addition, we show that the implied daily GARCH-NIG model, estimated for the ten years of daily (non-traded) ECU returns up to 30 December 1998, provides very accurate out-of-sample predictive distributions for the actual (traded) Euro/US dollar exchange rate from 2 January 1999 to 31 December 2001.

Our results build directly on recent empirical findings and related theoretical developments in the literature. First, the idea of explicitly modelling realized volatility proxies has a long history in empirical finance (see, for example, Schwert, 1989; Hsieh, 1991, and, more recently Andersen *et al*., unpublished manuscript, 2002; Maheu and McCurdy, 2002). Second, empirical results reported in Andersen *et al*. (2000, 2001a) have previously demonstrated the approximate normality of the returns when standardized by the realized volatility for other asset classes and time periods. Third, Barndorff-Nielsen and Shephard (2002a) report that the unconditional realized volatilities for other exchange rates are well approximated by the Inverse Gaussian distribution, and that, empirically, the IG distribution is almost indistinguishable from the lognormal distribution typically employed in SV models. Hence, by combining and extending these concurrent ideas within a unified empirical framework, we offer a new basis for gauging the distributional assumptions underlying the 'classical' daily GARCH-NIG modelling approach recently advocated in the literature.

The rest of the paper is organized as follows. The following section outlines the basic building blocks behind the GARCH-NIG model within the MDH framework. Section 3 describes the data and briefly reviews the notion of realized volatility. This section also provides our characterization of the raw and standardized return distributions, as well as the conditional and unconditional distributions of the realized volatilities. The in-sample daily GARCH-NIG model estimates for the ECU and the out-of-sample Euro predictions are briefly discussed in Section 4. Section 5 concludes.

2. THE MDH AND THE GARCH-NIG MODEL

The Mixture-of-Distributions Hypothesis (MDH) starts from the premise that the distribution of discretely sampled returns, conditional on some latent information arrival process, is Gaussian. This assumption is justified theoretically if the underlying price process follows a continuous sample path diffusion (see, for instance, the discussion in Andersen *et al.*, 2002, and Barndorff-Nielsen and Shephard, 2001b). However, the integrated volatility process that serves as the mixture variable in this situation is not directly observable. As noted above, this has spurred numerous empirical investigations into the use of alternative volatility proxies and/or mixture variables.

Meanwhile, in the diffusion setting the integrated volatility may, in theory, be estimated arbitrarily well by the summation of finely sampled squared high-frequency returns, or so-called realized volatilities. This suggests the following empirically testable starting point for the MDH,

$$f\,(r_t|RV_t) \sim N(0, RV_t) \tag{1}$$

where r_t refers to the discretely sampled one-period returns from time $t-1$ to t, and RV_t denotes the corresponding realized volatility proxy measured over the same time interval. Consistent with earlier related empirical results in Andersen *et al.*(2000), the results for the high-frequency foreign exchange rates discussed in the next section are broadly consistent with this hypothesis.

The distribution in (1), and the unconditional and conditional distributions of RV_t, jointly determine the distribution of the returns. In the original MDH formulation advocated by Clark (1973), the latent mixing variable is simply assumed to be *i.i.d.* lognormally distributed, resulting in an unconditional lognormal normal mixture distribution. This is also consistent with the empirical results reported in Andersen *et al.* (2001a,b), which suggest that the unconditional distributions of the logarithmic realized exchange rate and equity volatilities are approximately normal.

Unfortunately, the lognormal distribution is not closed under temporal aggregation, and the density function for the resulting lognormal normal mixture distribution is only available in integral form. Partly in response to these concerns, Barndorff-Nielsen and Shephard (2002a) have recently demonstrated that the unconditional distribution of the realized volatility may be equally well approximated by the Inverse Gaussian (IG) distribution. That is,

$$f\,(RV_t) \sim IG_{\sigma^2}\left(\sigma^2, \alpha\right) \tag{2}$$

where the density function for the IG distribution may be expressed in standardized form as,[1]

$$IG_{\sigma^2}\left(z;\sigma^2,\alpha\right)=\frac{\left(\dfrac{1}{\alpha\sigma^2}\right)^{-1/2}z^{-3/2}}{(2\pi)^{1/2}}\exp\left\{\alpha-\frac{1}{2}\left(\frac{\alpha\sigma^2}{z}+\frac{\alpha z}{\sigma^2}\right)\right\}$$

In contrast to the lognormal distribution, the IG distribution is formally closed under temporal aggregation. Hence, if, say the daily realized volatility is IG distributed, the weekly realized volatility, defined by the summation of the daily realized volatilities within the week, will also be IG distributed. The empirical results for the realized exchange rate volatility series in the next section, strongly support the IG distribution in equation (2).

Now, combining the distributional assumptions in (1) and (2) the implied unconditional distribution for the returns should be Normal Inverse Gaussian (NIG),

$$f\left(r_t\right)=\int_0^{\infty}f\left(r_t|RV_t\right)f\left(RV_t\right)dRV_t\sim NIG_{\sigma^2}\left(\sigma^2,0,0,\alpha\right) \tag{3}$$

with the following closed-form density function

$$NIG_{\sigma^2}\left(z;\sigma^2,\alpha\right)=\frac{\alpha^{1/2}}{\pi\sigma}\exp\left(\alpha\right)q\left(\frac{z}{\sigma\alpha^{1/2}}\right)^{-1}K_1\left(\alpha q\left(\frac{z}{\sigma\alpha^{1/2}}\right)\right)$$

where $q(x)=\sqrt{1+x^2}$ and $K_1(z)$ denotes the modified Bessel function of third order and index one.[2] The NIG distribution was first used for modelling speculative returns in Barndorff-Nielsen (1997). It may be viewed as a special case of the Generalized Hyperbolic Distribution in Barndorff-Nielsen (1978).[3]

Although the NIG distribution in equation (3) may adequately capture the fat-tailed unconditional return distributions, it does not account for the well-documented volatility clustering, or ARCH, effects. In order to incorporate conditional heteroscedasticity in the return process within the MDH framework, define the I_{t-1} information set generated by the past daily returns,

$$I_{t-1}=\sigma\left\{r_{t-1},r_{t-2},\ldots\right\}$$

The standard ARCH class of models then essentially entails the joint specification of a conditional density for r_t given I_{t-1} along with a parametric expression for the conditional variance. In the next section, we show that the corresponding *conditional* distribution for RV_t is similarly well approximated empirically by the IG distribution,

$$f\left(RV_t|I_{t-1}\right)\sim IG_{\sigma^2}\left(\sigma_t^2,\alpha\right) \tag{4}$$

[1] We use the notation $z\sim IG_{\sigma^2}\left(\alpha,\sigma^2\right)$ to indicate the standardized parameterization. Note that $cz\sim IG_{\sigma^2}\left(\alpha,c^2\sigma^2\right)$, so that the parameter α does not change under scaling. The mean, variance, skewness and kurtosis of this IG distribution are, respectively, $E(z)=\sigma^2$, $V(z)=\sigma^4/\alpha$, $S(z)=3/\sqrt{\alpha}$ and $K(z)=3+15/\alpha$.

[2] The first four central moments of the NIG distribution as parameterized here, are $\mu_1=0$, $\mu_2=\sigma^2$, $\mu_3=0$, and $\mu_4=3\sigma^4\left[1/\alpha+1\right]$, respectively. Non-zero skewness in the distribution of the returns could easily be accommodated by the estimation of a non-zero asymmetry parameter in the NIG distribution; see Jensen and Lunde (2001).

[3] The same arguments for the IG distribution outlined here easily extend to the corresponding Generalized IG (GIG) mixture.

where σ_t^2 depends exclusively on the coarser I_{t-1} information set. Specifically, for the results reported below, the conditional mean in the IG distribution is assumed to follow the recursive GARCH(1,1) like structure,

$$\sigma_t^2 = \rho_0 + \rho_1 r_{t-1}^2 + \pi_1 \sigma_{t-1}^2 \tag{5}$$

The GARCH(1,1)-NIG model now arises naturally by combining equation (1) augmented with the I_{t-1} information set, $f(r_t|RV_t, I_{t-1}) \sim N(0, RV_t)$, with the conditional distribution for RV_t in equations (4) and (5),

$$f(r_t|I_{t-1}) = \int f(r_t|RV_t, I_{t-1}) f(RV_t|I_{t-1})\, dRV_t \sim NIG_{\sigma^2}\left(\sigma_t^2, 0, 0, \alpha\right) \tag{6}$$

The ARCH-NIG model was first proposed by Barndorff-Nielsen (1997) and later extended to allow for more flexible GARCH type parameterizations in Andersson (2001) and Jensen and Lunde (2001). By explicitly equating the mixing variable with the observable realized volatility, equations (1) and (4) allow for direct empirical investigation and verification of the salient distributional features underlying the GARCH-NIG model, a task to which we now turn.

3. DISTRIBUTIONAL PROPERTIES OF REALIZED VOLATILITY AND RETURNS

3.1. Data Sources and the Construction of Realized Volatility

Our primary data set consists of daily returns and realized volatilities for the ECU/US Dollar exchange rate from 3 January 1989 to 30 December 1998. All of the raw data were obtained from Olsen and Associates in Zürich, Switzerland. Following standard practice in the literature, the daily realized volatilities were constructed by summing squared five-minute returns. Formally,

$$RV_t = \Sigma_{i=1}^{288} r_{(288)}^2 (t + (i/288)) \tag{7}$$

where $r_{(288)}^2 (t + (i/288))$ denotes the continuously compounded return for day t over the ith five-minute interval calculated on the basis of the linearly interpolated logarithmic midpoint of the bid–ask prices. We explicitly omit non-trading days and weekend periods as described in Andersen *et al.* (2001b). All in all, this leaves us with a total of 2428 days. [4] Time series plots of the relevant returns and realized volatilities are given in the two panels in Figure 1.

In addition to the daily ECU returns and realized volatilities, our out-of-sample predictive analysis is based on the 780 daily returns for the Euro from 5 January 1999 to 31 December 2001. We do not have access to the high-frequency data underlying the construction of the realized volatilities over this more recent time period. Hence, our empirical verification of the various distributional hypotheses involving the RV_t variable outlined in the previous section will be based exclusively on the ECU data.

[4] We also excluded nine days in January and February 1989 on which the realized volatility was less than 0.005. These days are directly associated with problems in the data-feed early on in the sample. None of the results are sensitive to these additional exclusions.

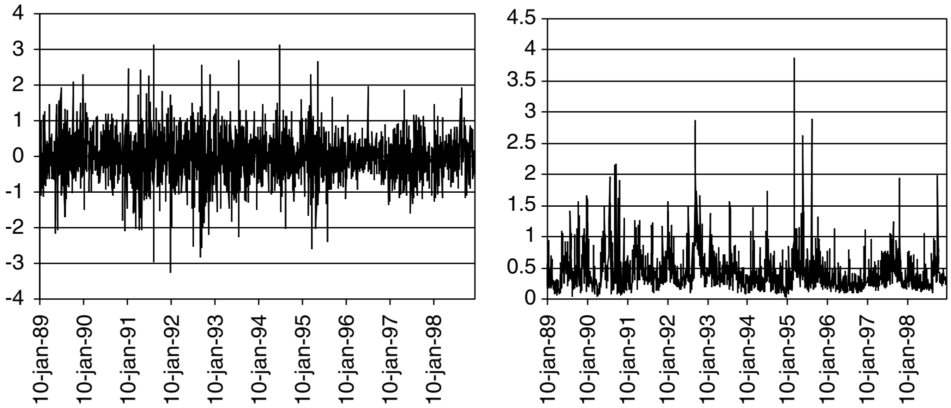

Figure 1. Daily returns and realized volatilities. (Note: The two panels in the figure show the daily returns and RV for the ECU/USD exchange rate from 1989-98)

3.2. Unconditional Distributions

Figure 2 gives the QQ-plot for the probability integral transform, or PIT, for the raw daily ECU returns under the assumption of normality.[5] The figure shows the typical S-shape associated with a fat-tailed unconditional return distribution. Further evidence is provided by the unconditional sample kurtosis which equals 5.426. The standard Jarque–Bera test statistic of 598.1 also firmly rejects the null hypothesis of normally distributed returns when evaluated in the asymptotic chi-square distribution with two degrees of freedom.

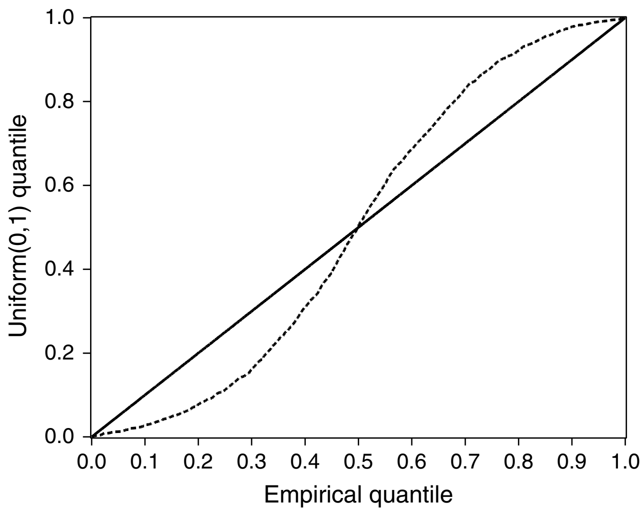

Figure 2. QQ-plot of daily raw returns. (Note: The plot shows the PIT of the raw daily returns for the ECU/USD 1989-98 data, assuming them to be $N(0,1)$ against the quantiles of the $U(0,1)$ distribution)

[5] The PIT is defined as $z_t^{PIT} = \int_{-\infty}^{x_t} f(u)\,du$. If $f(u)$ is the correct distribution, then $z_t^{PIT} \sim U(0,1)$.

Meanwhile, the QQ-plot for the cumulative distribution function of the daily returns, standardized by the realized volatility, $r_t/\sqrt{RV_t}$, reported in Figure 3, indicates a very good fit by the standard normal distribution. The corresponding Jarque–Bera test statistic for normality also drops to only 4.318, which has an asymptotic p-value of 0.115. This approximate normality of the standardized daily returns is, of course, directly in line with the distribution postulated in equation (1).

Turning to the distribution of the realized volatility, the QQ-plot for the IG distribution in Figure 4 shows an equally good empirical fit of the unconditional volatility distribution in equation (2). The maximum likelihood estimates for σ^2 and α in the IG distribution, with asymptotic standard errors in parentheses, are 0.427 (0.006) and 2.028 (0.051), respectively. The estimate for σ^2 is identically equal to the sample mean of the realized volatilities. The implied standard deviation, equal to 0.300, also closely matches the sample standard deviation of 0.311. It is worth noting that the conditional IG distributions for the temporally aggregated weekly and monthly realized volatilities provide similar close approximations to the empirical distributions (for further discussion of these results see Forsberg, 2002).

The distribution for the returns in equation (1) and the realized volatilities in equation (2) imply that the returns should be unconditionally NIG distributed as in Equation (3). The QQ-plot for the daily raw returns in Figure 5, does indeed indicate a very close fit of the NIG distribution. The corresponding maximum likelihood estimates for the NIG parameters are reported in the second column in Table I. The unconditional kurtosis implied by the estimates is equal to 5.636, which is close to the aforementioned sample value of 5.426.

These results further underscore the applicability of the NIG in characterizing the unconditional distribution of speculative returns. However, the unconditional NIG distribution obviously does not account for the well-documented volatility clustering phenomenon. In order to do so, we now turn to a discussion of the conditional return and volatility distributions.

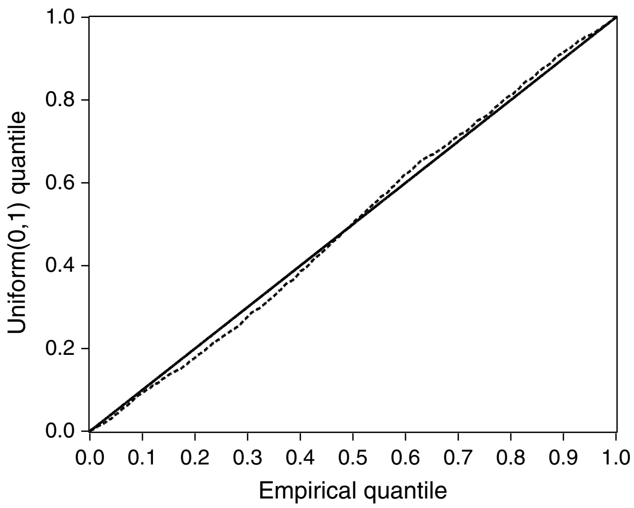

Figure 3. QQ-plot of PIT of daily returns standardized by the realized volatility. (Note: The plot shows the PIT of the daily returns for the ECU/USD 1989-98 data standardized by the realized volatility, assuming them to be $N(0,1)$ against the quantiles of the $U(0,1)$ distribution)

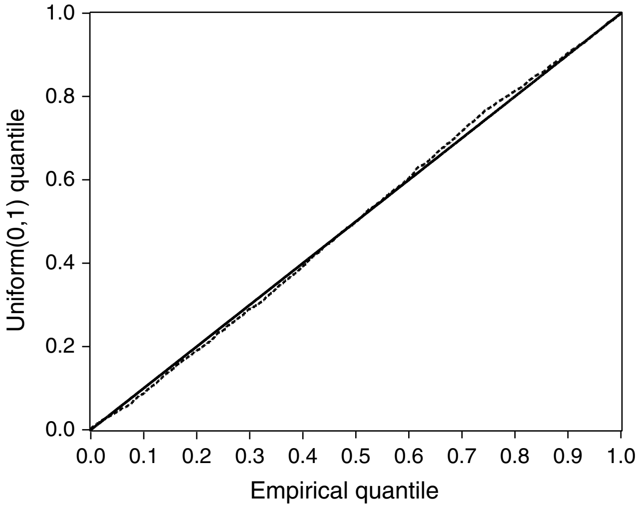

Figure 4. QQ-plot of daily unconditional realized volatility. (Note: The figure shows a QQ-plot of the PIT for the unconditional daily realized volatility for the ECU/USD 1989-98 assumed to be $IG(0.427, 2.028)$, against the quantiles of the $U(0,1)$ distribution)

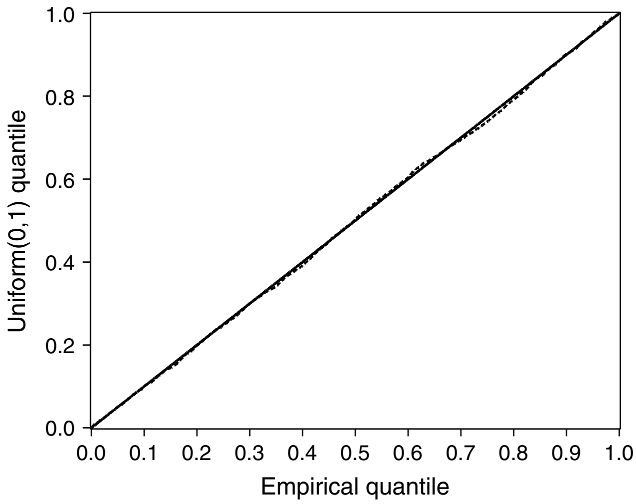

Figure 5. QQ-plot of daily raw returns. (Note: The figure shows a QQ-plot of the PIT of the daily demeaned returns for the ECU/USD 1989-98 data assumed to be $NIG(1.138, 0, 0, 0.407)$, against the quantiles of the $U(0,1)$ distribution)

3.3. Conditional Distributions

Not surprisingly, the Q(1) and Q(10) Ljung–Box portmanteau statistics reported in Table I are generally consistent with the absence of important own serial correlation in the daily returns. Meanwhile, the corresponding portmanteau statistics for the squared daily returns indicate highly

Table I. Distributions and models fitted to the daily ECU/USD 1989-1998 return

	Distributions		Models	
	Normal	NIG	GARCH-N	GARCH-NIG
μ	0.002 (0.013)	0.005 (0.011)	0.007 (0.012)	0.005 (0.011)
α	—	1.138 (0.167)	—	1.705 (0.287)
σ^2	0.407 (0.008)	0.407 (0.017)	—	
ρ_0	—	—	0.006 (0.001)	0.004 (0.002)
ρ_1	—	—	0.049 (0.006)	0.041 (0.008)
π_1	—	—	0.937 (0.007)	0.951 (0.010)
Log Lik.	−2352.9	−2251.2	−2237.9	−2177.9
AIC	4711.9	4506.3	4483.9	4365.9
BIC	4723.5	4517.9	4507.1	4394.8
$Q(1)$	0.040 (0.840)	0.040 (0.840)	0.044 (0.832)	0.070 (0.790)
$Q(10)$	20.270 (0.026)	20.270 (0.026)	15.780 (0.106)	16.355 (0.089)
$Q^2(1)$	34.4 (0.000)	34.4 (0.000)	0.088 (0.766)	0.324 (0.569)
$Q^2(10)$	208.5 (0.000)	208.5 (0.000)	4.128 (0.941)	4.963 (0.893)

Notes: AIC and BIC refer to the standard model selection criteria.
$Q(n)$ denotes the Ljung–Box statistic for serial correlation in the standardized returns for up to lag n. $Q^2(n)$ denotes the Ljung–Box statistic for the squared standardized residuals.

significant volatility clustering effects. At the same time, the $Q^2(1)$ statistic for r_t^2 drops from 34.4 to only 1.208 for the squared standardized returns, r_t^2/RV_t. Hence, this suggests that the close fit afforded by the unconditional normal distribution for $r_t/\sqrt{RV_t}$ in Figure 3 carries over to the conditional distribution in equation (1) when augmented by the I_{t-1} information set.[6]

The significant ARCH effects in the squared returns and the lack of any serial correlation in the squared standardized returns is further manifested in the strong temporal dependencies in the realized volatility series. For instance, the Ljung–Box Q(1) statistic for testing first order serial correlation in RV_t, is 667.9. This highlights the importance of incorporating conditional dependencies when characterizing the distribution of RV_t in equation (4).

Following the discussion in the previous section, we proceed to do so by estimating an IG distribution for the standardized realized volatility, RV_t/σ_t^2, where σ_t^2 is determined by the GARCH(1,1) model structure in equation (5). Of course, richer parameterizations involving longer lags and past values of the realized volatility may result in a better fit. However, since we are primarily concerned with the justification of the distributional assumptions underlying standard

[6] This is further corroborated by simple regression based tests of $r_t/\sqrt{RV_t}$ and r_t^2/RV_t on additional lagged daily and daily squared returns. For instance, including a single lag, the t-statistics for zero slope coefficients equal −0.842 and 0.012 in each of the two regressions, respectively.

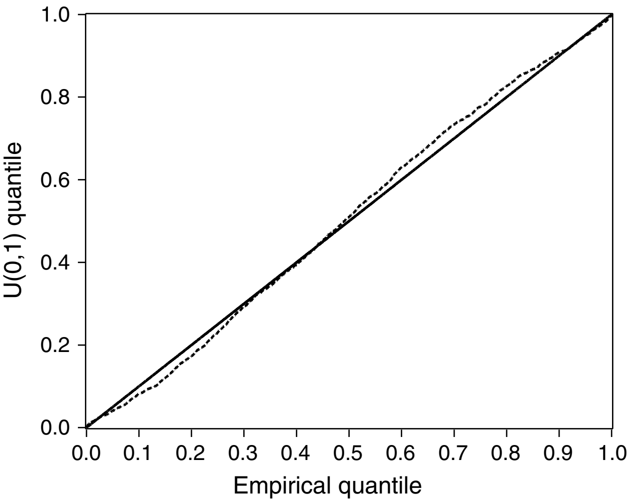

Figure 6. QQ-plot of GARCH(1,1)-IG standardized realized volatility. (Note: The figure shows a plot of the PIT of the daily realized volatility for the ECU/USD 1989-98 data standardized by the GARCH(1,1)-IG model, against the quantiles of the $U(0,1)$ distribution)

daily ARCH models, we explicitly restrict the conditioning information set to the past daily returns only, or I_{t-1}. The resulting maximum likelihood estimates for the GARCH(1,1)-IG model are: $\alpha = 2.732(0.062)$, $\rho_0 = 0.037(0.003)$, $\rho_1 = 0.077(0.006)$ and $\pi_1 = 0.841(0.012)$. These parameter estimates for the realized volatility, are directly in line with the estimates for traditional GARCH models based on daily data reported in the extant literature. More importantly from the present perspective, however, is the quality of the fit afforded by the error distribution. Again, the QQ-plot for the PIT in Figure 6 indicates that the implied IG distribution provides a very accurate approximation. Guided by these distributional features, we turn next to a discussion of the results for the estimation and empirical fit of the GARCH(1,1)-NIG model implied by the mixture in equation (6).

4. THE GARCH-NIG MODEL

The empirical analysis in the previous section explicitly relies on the availability of high-frequency intraday data. In contrast, all of the results discussed in this section are based on daily data only.

4.1. In-sample ECU Estimates

The parameter estimates and summary statistics for the daily models in Table I are directly in line with the results in the extant literature. In particular, the estimates for α clearly indicate significant excess kurtosis, both in the unconditional NIG distribution and in the conditional GARCH-NIG model.[7] The superior fit of the GARCH-NIG error distribution is further illustrated by the QQ-plot for the PITs in Figure 7. Of course, for the model to be correctly conditionally calibrated, the

[7] Necessary and sufficient conditions for the existence of all moments in the GARCH(p, q) model have recently been established by Ling and McAleer (2002). In particular, the fourth unconditional moment for the GARCH(1,1)-NIG model

corresponding sequence of PITs, say $\{z_t\}$, should also be *i.i.d.* through time (this idea of checking model adequacy has previously been used in a similar context by Diebold, Gunther and Tay, 1998, and Kim, Shephard and Chib, 1998, among others). Consistent with this hypothesis, the Ljung–Box portmanteau statistics for first-order serial correlation in z_t and z_t^2 equal to 2.076 and 1.358, respectively, do not suggest any significant temporal dependencies.

It is also noteworthy that the AIC and BIC model selection criteria for the maximum likelihood estimates for the commonly employed GARCH-*t* model equal 4367.4 and 4396.4, respectively, while the same two criteria for the popular GARCH-GED model equal 4374.7 and 4403.7, respectively. Thus, the GARCH-NIG model is the preferred specification according to both criteria. Of course, the MDH outlined in Section 2, along with the distributional results for the realized volatility discussed in Section 3, provide further in-sample support for the model.

4.2. Out-of-sample Euro Predictions

Additional out-of-sample support for the NIG error distribution is provided by the results in Table II. The table reports the unconditional coverage probabilities for the one-day-ahead VaRs, or quantile predictions, obtained by applying the GARCH-NIG model for the 1989-1998 ECU returns in Table I to the post-sample 1999-2001 Euro data. The results are striking. All of the empirical VaRs are easily within the standard confidence bands around their respective nominal sizes.[8] Moreover, the corresponding sequence of PITs indicates no apparent temporal dependency in the accuracy of the out-of-sample predictions from the model.

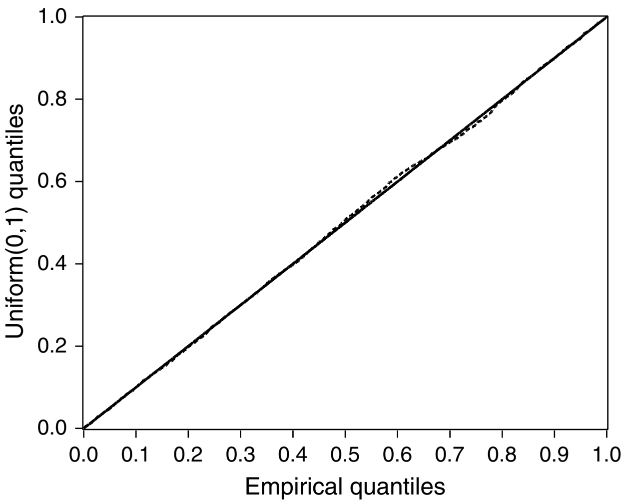

Figure 7. QQ-plot of GARCH standardized daily returns. (Note: The figure shows the PIT of the daily ECU/USD 1989-98 returns standardized by the GARCH-NIG mode)

exists if and only if $\pi_1^2 + 2\pi_1\rho_1 + 3\left(1+\alpha^{-1}\right)\rho_1^2 < 1$, which is trivially satisfied for the parameter estimates reported in Table I. The corresponding implied kurtosis from the model equals 7.887, which actually exceeds the sample kurtosis of 5.426.

[8] Using the standard binomial approximation, a 95% confidence interval is readily calculated as $\pm 1.96\sqrt{\widehat{\kappa}\left(1-\widehat{\kappa}\right)/780}$, where $\widehat{\kappa}$ refers to the empirical size.

Table II. Value-at-Risk for GARCH(1,1)-NIG model

Nominal	Sizes (%)					
	1	5	10	90	95	99
Empirical	1.538	5.513	10.00	88.97	94.49	99.10

Notes: The table displays empirical VaRs for the daily Euro/USD over the 1999-2001 time period based on the GRACH(1,1)-NIG model estimates for the 1989-1998 daily ECU/USD returns reported in Table I.

For instance, the tests for first-order serial correlation in the PITs and the PITs squared equal 0.201 and 0.041, respectively, with corresponding p-values of 0.653 and 0.838, respectively. These results are particularly noteworthy insofar as they represent true out-of-sample verifications of the distributional assumptions underlying the in-sample GARCH-NIG model estimates.[9]

5. CONCLUDING REMARKS

This paper illustrates how realized volatility measures, constructed from high-frequency intraday returns, may be used in guiding the specification of traditional ARCH models at lower interdaily frequencies. Our empirical results corroborate the arguments in favour of the GARCH-NIG model recently proposed in the literature, as providing accurate and parsimonious representations of daily exchange rate dynamics. It would be interesting to see how these results carry over to other asset classes and different time periods. In the case of equities, skewness in the returns may necessitate the use of a more general NIG mixture distribution. A more thorough investigation of issues related to temporal aggregation would also be interesting. Of course, in situations when the high-frequency data underlying the construction of the realized volatility measures are actually available, even better empirical representations of the low-frequency return dynamics may, in theory, be obtained by directly modeling the high-frequency data. However, such models are invariably complicated by a host of market microstructure features. Directly modelling the realized volatility, as recently advocated by Andersen *et al.*(unpublished manuscript. 2002), may provide a simple, yet effective, way of utilizing the inherent information in the high-frequency data. Meanwhile, the availability of reliable high-frequency data remains the exception rather than the norm for most speculative assets. As such, the traditional ARCH class of models pioneered by Engle (1982), and further analysed here, is likely to remain the workhorse in volatility modelling for years to come.

ACKNOWLEDGEMENTS

The authors are grateful to Olsen and Associates for making the intraday exchange rate quotations for the ECU/USD available. We also would like to thank participants at the Financial Econometrics Lunch Group at Duke University, the conference on 'Modeling and Forecasting Financial Volatility' at the University of Western Australia, 7-9 September 2001, Anders Eriksson, Eric Ghysels, two anonymous referees, and the editors Philip Hans Franses and Michael McAleer

[9] It is also noteworthy that when evaluating the log likelihood functions for the GARCH-N and GARCH-NIG models for the Euro sample at the parameter estimates in Table I, the log likelihood values for the two models are −679.21 and −640.45, respectively.

for many helpful comments and suggestions. The main part of the paper was written while Forsberg visited the Department of Economics, Duke University, whose hospitality is gratefully acknowledged. The work was supported by a grant from The Swedish Foundation for International Cooperation in Research and Higher Education (STINT) grant number: IG 2002-02 045 and Jan Wallanders and Tom Hedelius Foundation grant number: J 01-28 (Forsberg) and a grant from the NSF to the NBER (Bollerslev).

REFERENCES

Andersen TG, Bollerslev T, Diebold FX, Ebens H. 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics* **61**: 43–76.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2000. Exchange rate returns standardized by realized volatility are nearly Gaussian. *Multinational Finance Journal* **4**: 159–179.

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2001b. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* **96**: 42–55.

Andersen TG, Bollerslev T, Diebold FX. 2002. Parametric and nonparametric volatility measurement. In *Handbook of Financial Econometrics*, Hansen LP, Ait-Sahalia Y (eds). North-Holland: Amsterdam; (forthcoming).

Andersen TG. 1996. Return volatility and trading volume: an information flow interpretation of stochastic volatility. *Journal of Finance* **56**: 169–204.

Andersson J. 2001. On the normal inverse Gaussian stochastic volatility model. *Journal of Business and Economic Statistics* **19**: 44–54.

Andreou E, Ghysels E. 2002. Rolling-sample volatility estimators: some new theoretical, simulation and empirical results. *Journal of Business and Economic Statistics* (forthcoming).

Ane T, Geman H. 2000. Order flow, transaction clock, and normality of asset returns. *Journal of Finance* **60**: 2259–2284.

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

Barndorff-Nielsen OE, Shephard N. 2002a. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B* **64**: (forthcoming).

Barndorff-Nielsen OE, Shephard N. 2002b. How accurate is the asymptotic approximation to the distribution of realized volatility? In *Identification and Inference for Econometric Models*, Andrews DWF, Powell JL, Ruud PA, Stock JH (eds). forthcoming.

Barndorff-Nielsen OE. 1978. Hyperbolic distributions and distributions on hyperbolae. *Scandinavian Journal of Statistics* **5**: 151–157.

Barndorff-Nielsen OE. 1997. Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics* **24**: 1–13.

Bollerslev T. 1987. A conditionally heteroscedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* **69**: 542–547.

Bollerslev T, Chou RY, Kroner KF. 1992. ARCH modeling in finance: a review of the theory and empirical evidence. *Journal of Econometrics* **39**: 5–59.

Clark PK. 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* **41**: 135–155.

Comte F, Renault E. 1998. Long memory in continuous time stochastic volatility models. *Mathematical Finance* **8**: 291–323.

Diebold FX, Gunther T, Tay AS. 1998. Evaluating density forecasts, with applications to financial risk management. *International Economic Review* **39**: 863–883.

Engle RF. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation. *Econometrica* **50**: 987–1008.

Epps TW, Epps ML. 1976. The stochastic dependence of security prices changes and transaction volumes: implications for the mixture-of-distributions hypothesis. *Econometrica* **44**: 305–321.

Forsberg L. 2002. *On the Normal Inverse Gaussian Distribution in Modeling Volatility in the Financial Markets*. Doctoral thesis, Studia Statistica Upsaliensia **5**, Christofferson A, Jöreskog KG (eds). Uppsala University Library.

Hsieh DA. 1991. Chaos and nonlinear dynamics: application to financial markets. *Journal of Finance* **46**: 1839–1877.

Jensen MB, Lunde A. 2001. The NIG-S&ARCH model: a fat-tailed, stochastic, and autoregressive conditional heteroscedastic volatility model. *Econometrics Journal* **4**: 319–342.

Kim S, Shephard N, Chib S. 1998. Stochastic volatility: likelihood inference and comparison with ARCH models. *Review of Economic Studies* **65**: 361–393.

Ling S, McAleer M. 2002. Necessary and sufficient moment conditions for the GARCH(p,q) and asymmetric power GARCH(r,s) models. *Econometric Theory* **18**: 722–729.

Maheu JM, Mc Curdy TH. 2002. Nonlinear features of realized FX volatility. *Review of Economics and Statistics* (forthcoming).

Nelson D. 1991. Conditional heteroscedasticity in asset returns: A new approach. *Econometrica* **59**: 347–370.

Schwert GW. 1989. Why does stock market volatility change over time? *Journal of Finance* **44**: 1115–1153.

Shephard N. 1996. Statistical aspects of ARCH and stochastic volatility. In *Time Series Models in Econometrics, Finance and other Fields*, Cox DR, Barndorff-Nielsen OE, Hinkley DV (ed.). Chapman and Hall: London; 1–67.

Tauchen GE, Pitts M. 1983. The price variability–volume relationship on speculative markets. *Econometrica* **51**: 485–505.

Taylor SJ. 1982. Financial returns modeled by the product of two stochastic processes—a study of the daily sugar prices 1961–75. In *Time Series Analysis: Theory and Practice*, Anderson OD (ed.). North-Holland: Amsterdam; 203–226.

Taylor SJ. 1986. *Modelling Financial Time Series*. John Wiley: Chichester.

