Economic Research Institute Study Paper
ERI # 97-04

A FLEXIBLE PARAMETRIC GARCH MODEL WITH AN APPLICATION TO EXCHANGE RATES

By

Kai Li Wang
Chrisotpher Fawson
Christopher B. Barrett
James B. McDonald

DEPARTMENT OF ECONOMICS
UTAH STATE UNIVERSITY
LOGAN, UTAH

April 1998

A Flexible Parametric GARCH Model

With An Application To Exchange Rates

Kai Li Wang
Christopher Fawson
Christopher B. Barrett
James B. McDonald

March 1998

The authors are assistant professor, Department of International Trade, Tam-Kang University, associate professor and assistant professor, respectively, in the Department of Economics, Utah State University, and professor in the Department of Economics, Brigham Young University. Fawson is the corresponding author at Dept. of Economics, Utah State University, Logan, UT 84322-3530, email: cfawson@b202.usu.edu. We thank Basudeb Biswas and Terry Glover for comments, and Ron Schoenberg for programming advice. Remaining errors are ours. This work was supported by the Utah Agricultural Experiment Station and approved as journal paper 5068.

A Flexible Parametric GARCH Model

With An Application To Exchange Rates

Abstract: Many asset prices, including exchange rates, exhibit stable periods punctuated by substantial, often one-sided adjustments. Statistically, this generates empirical distributions of exchange rate changes that often have high peaks, long tails and are sometimes skewed. Existing time series estimation methods do not account for these characteristics satisfactorily. This paper introduces a more general GARCH model, with a flexible parametric error distribution based on the exponential generalized beta distribution. Applied to daily U.S. dollar exchange rate data for six major currencies, the GARCH-EGB2 model uniformly outperforms conventional GARCH models of exchange rate volatility in sample and generates superior near-term out-of-sample forecasts.

JEL Codes: C13, C22, F31, G15

I. Introduction

Contemporary modeling of exchange rate time series makes widespread use of generalized autoregressive conditional heteroskedastic (GARCH) models.[1] Not only can GARCH models capture the volatility clustering often found in exchange rate series, they also accommodate some of the leptokurtosis (i.e., thick tails) commonly found in exchange rate series. But GARCH models with conditionally normal errors generally fail to capture sufficiently the leptokurtosis evident in asset returns (Bollerslev 1987, Baillie and Bollerslev 1989, Hsieh 1989, Baillie and DeGennaro 1990, Wang, Barrett, and Fawson 1996). This has led to widespread adoption of nonnormal conditional error distributions, most commonly the student-t (Bollerslev 1987, Bollerslev et al. 1994).

What might cause the unconditional leptokurtosis commonly observed in exchange rate series and thus motivate use of GARCH modeling? Economic theories of exchange rate determination offer two likely explanations.[2] The first is the overshooting of floating nominal exchange rates associated with monetary or fiscal shocks in the presence of sticky prices (Dornbusch 1976). The other is speculative attacks against fixed exchange rates (Krugman 1979). Both models imply infrequent, extraordinarily sharp movements in exchange rates, i.e., the sorts of movements that appear as long (i.e., fat) tails in a distribution of differenced exchange rates. Sharp exchange rate movements do not necessarily imply leptokurtosis, however; they

---

[1] Bollerslev, Chou and Kroner (1992) offer a good survey.

[2] While we do not develop a formal derivation of the statistical implications of exchange rate determination models in this paper, it is nonetheless useful to consider this issue empirically. Taylor (1995) and Obstfeld and Rogoff (1996) offer excellent, formal treatments of exchange rate determination models.

could imply high variance in the time series. The key is that sticky prices in floating rate regimes, and especially fixed exchange rates, also generate modal daily exchange rate changes near zero (Obstfeld and Rogoff 1996). The implication is that exchange rate changes are concentrated near the mean but have long tails, and hence leptokurtosis.[3] The choice of a conditional distribution should accommodate both leptokurtosis and high peakedness in the exchange rate series. While the commonly used student-t is a leptokurtic distribution, it is also low-peaked, and perhaps inappropriate to time-series analysis of exchange rates.

Moreover, skewness might also be important in exchange rate series that exhibit episodes of sharp depreciation (appreciation) not offset by subsequent sharp appreciation (depreciation). There are two likely reasons for such phenomena. First, permanent shocks that lead to changes in the equilibrium exchange rate may be asymmetric; rapid improvements in Japanese productivity over the past thirty years seem an excellent example. Second, speculative attacks against a currency tend to be one-sided (causing depreciation/devaluation). The 1992-3 European, 1994 Mexican, and 1997-98 east Asian currency crises — including attacks against the British pound and the Italian lira of particular relevance to this study — are good recent examples of such episodes. Since significant skewness is often observed in exchange rate series (Boothe and Glassman 1987, Hsieh 1988, Peruga 1988), it seems advisable to employ estimation methods that

[3] An alternative way to view this, following Friedman (1953), is to recognize that profit-maximizing speculators' transactions stabilize transitory shocks to the exchange rate and accelerate movement in response to permanent shocks. If transitory shocks are far more common than permanent shocks, this will yield an empirical distribution of exchange rate changes that is high peaked and long tailed.

accommodate skewness.[4]

GARCH estimation based on conditional student-t distributions can capture leptokurtosis in exchange rate series, but fares less well in replicating their high peakedness and skewness. As a step toward resolving this deficiency, we introduce a GARCH model with a more flexible error distribution based on the exponential generalized beta (EGB) family of distributions (McDonald and Xu 1995).

The plan of the paper is as follows. Section II introduces the EGB family of distributions, including the specific variant used in this paper, the exponential generalized beta of the second kind (EGB2), and then develops a GARCH-EGB2 model. In section III we estimate time series models of the U.S. dollar exchange rates for six major industrial economy currencies at daily frequency. These exchange rate series exhibit high peakedness, leptokurtosis, and skeweness. The likelihood dominance criterion (for non-nested models), goodness of fit statistics and plots of the standardized residuals all indicate that the GARCH-EGB2 model systematically outperforms Gaussian GARCH and GARCH-t models in sample for each of the exchange rate series modeled. Section IV discusses the out-of-sample forecasting performance of the GARCH-EGB2 model. The concluding section summarizes our findings and highlights some implications for future research.

II. The GARCH-EGB2 Model

[4] Hansen's (1994) method does this, but by using a modified student-t distribution, so that he cannot simultaneously accommodate high peakedness. Moreover, Hansen's model depends on appropriate *ex ante* lag selection.

McDonald and Xu (1995) introduced the five-parameter generalized beta (GB) distribution and its logarithmic transform, the exponential generalized beta (EGB) distribution. The GB includes as special cases familiar distributions (e.g., lognormal, logistic, Pareto, generalized gamma, Burr, and Weibull). We use the EGB distribution because it can model both positive and negative random variables, while the GB models just positive random variables.

The EGB distribution is defined by the probability density function (pdf):

$$EGB(z;\delta,\sigma,c,p,q) = \frac{e^{\frac{p(z-\delta)}{\sigma}}\left(1-(1-c)e^{\frac{(z-\delta)}{\sigma}}\right)^{q-1}}{|\sigma|B(p,q)\left(1+ce^{\frac{z-\delta}{\sigma}}\right)^{p+q}}$$

$$for \quad -\infty < \frac{z-\delta}{\sigma} < \ln\left(\frac{1}{1-c}\right)$$

where δ is a location parameter that affects the mean of the distribution, σ reflects the peakedness of the density function, p and q are shape parameters that together determine the skewness and kurtosis of the distribution, and $0 \leq c \leq 1$. Of special interest are the exponential generalized beta distributions of the first and second kind, EGB1 and EGB2, respectively, which correspond to the limiting values of c=0 and c=1, respectively, and are alternative representations of the generalized exponential and generalized logistic distributions, respectively (Johnson and Kotz 1970, Patil et al. 1984). The associated probability density functions are

$$EGB1(z;\delta,\sigma,p,q) = EGB(z;\delta,\sigma,c=0,p,q)$$
$$= \frac{e^{\frac{p(z-\delta)}{\sigma}}\left(1-e^{\frac{p(z-\delta)}{\sigma}}\right)^{q-1}}{|\sigma|B(p,q)}$$

$$EGB2(z;\delta,\sigma,p,q) = EGB(z;\delta,\sigma,c=1,p,q)$$
$$= \frac{e^{\frac{p(z-\delta)}{\sigma}}}{|\sigma| B(p,q)\left(1+e^{\frac{z-\delta}{\sigma}}\right)^{p+q}}$$

Note that unlike the more general EGB distribution, EGB1 and EGB2 do not involve a nonlinear inequality constraint for the random variable. This feature makes numerical estimation of the latter distributions simpler than for the EGB. Furthermore, while the higher order moments of the EGB involve a relatively complex, hypergeometric series, the variance, skewness and kurtosis of EGB1 and EGB2 are relatively simple expressions. Table 1 presents equations for the variance, skewness and kurtosis of the EGB2 distribution employed in the empirical portion of this paper. Tractability therefore favors estimating four-parameter EGB1 or EGB2 distributions over the more general EGB form as long as the c parameter lies near one or zero. In this spirit McDonald and Xu (1995, p. 134) find that, "[t]he exponential generalized beta of the second kind (EGB2) provides the basis for partially adaptive estimation in regression and time series models to accommodate possibly thick-tailed and skewed error distributions." Since the prevailing concern about existing GARCH exchange rate models is unsatisfactory accommodation of leptokurtosis, skewness and high peakedness, the EGB or one of its two limiting distributions seems a promising conditional error distribution. In order to have a probability density function without restricted support and because preliminary results of estimating both the EGB and the EGB2 models suggest c is close to the unit boundary of the parameter space for each exchange rate, with the EGB2 specification uniformly favored over the EGB by likelihood ratio tests (Table 2), we use the EGB2 in this study.

An extensive literature finds time-varying conditional variance in asset prices, including exchange rates. GARCH models are commonly estimated under the assumption that the standardized residuals are normally distributed. Yet although the unconditional distribution of a GARCH process with normal errors is leptokurtic (Engle 1982, Bollerslev 1986) — i.e., its kurtosis is greater than 3.0, the benchmark value from the normal distribution — Gaussian GARCH models regularly fail to account adequately for the fat tails found in unconditional asset price distributions (Hsieh 1989; Wang, Barrett and Fawson 1996). As a consequence, many researchers now employ nonnormal conditional error distributions, particularly the student-t, in GARCH modeling.

Our concern about the GARCH modeling literature on exchange rates is that accommodation of leptokurtosis but not of the high peakedness or asymmetry commonly found in exchange rate series may lead to inappropriate choice of conditional error distributions. Given the problems associated with quasi-maximum likelihood GARCH estimation (Pagan and Sabau 1987, Lee and Hansen 1994, Deb 1996), incomplete accommodation of the statistical characteristics of exchange rates may yield inaccurate estimates of exchange rate dynamics. We therefore develop a GARCH model based on a more flexible EGB2 distribution.

We begin by adopting a general autoregressive moving average (ARMA) specification in the conditional mean equation with GARCH(1,1) errors.[5] With the right conditional distribution to describe the standardized errors, z_t, this specification can account for most of the characteristics observed in empirical financial distributions, including time-varying variance,

[5] The GARCH(1,1) specification we employ is generally excellent for a wide range of financial data (Bollerslev, Chou, and Kroner, 1992).

asymmetry, thick tails, and high peakedness. Denoting a time series dependent variable as y_t, the general form of this model is given by

ARMA(m,n) Conditional Mean Equation:

$$\phi_m(L)\, y_t = \mu + \phi_n(L)\, \epsilon_t$$

GARCH(1,1) Conditional Variance Equation:

$$\epsilon_t = h_t^{0.5} z_t$$
$$E(\epsilon^2_t \mid \psi_{t-1}) = h_t$$
$$= w + \alpha_1 \epsilon^2_{t-1} + \beta_1 h_{t-1}$$

Conditional Distribution :

$$\epsilon_t \mid \psi_{t-1} \sim D(0, h_t, \eta_t)$$

where the $\phi(L)$ are polynomials in the lag operator of order m and n, respectively, and w, α_1, $\beta_1 > 0$ to ensure strictly positive conditional variance. The errors follow the assumed zero mean conditional density function described by the variance h_t, and the parameter vector η_t. The latter are "shape" parameters, $\eta_t = \{p, q\}$ under EGB2, $\eta_t = \{\upsilon\}$ under the Student-t distribution, and η_t is the empty set under the normal distribution. To achieve efficiency, we jointly estimate the conditional mean and conditional variance equations with the conditional distribution by full information maximum likelihood using the GAUSS constrained maximum likelihood module.

For the standardized EGB2 distribution with the shape parameters p and q, the log-likelihood function of GARCH-EGB2 model is

$$LogL = T[\log(\sqrt{\Omega}) - \log(B(p,q)) + p\Delta] + \sum [(p(\frac{\sqrt{\Omega}\Delta\epsilon_t}{h_.}) - \log(h_t) - (p+q)\log(1+\exp(\frac{\sqrt{\Omega}\epsilon_t}{h_.} + \Delta\Omega)]$$

where

$$\Delta = \psi(p) - \psi(q)$$

$\Omega = \psi'(p) + \psi'(q)$

and $\psi(p)$ and $\psi'(p)$ represent digamma and trigamma functions, respectively. We show the detailed parameterization of the GARCH-EGB2 model in the technical appendix. For the Student-t distribution with ν degrees of freedom, the log-likelihood function of GARCH-t model is, as presented by Bollerslev (1987):

$$\log L = T[\log\Gamma(\frac{\upsilon+1}{2}) - \log\Gamma(\frac{\upsilon}{2}) - 0.5\log(\upsilon-2)] - 0.5\sum[\log h_t + (\upsilon+1)\log(1 + \frac{\epsilon_t^2}{h_t(\upsilon-2)}]$$

where Γ denotes the gamma function.

By adding just one extra parameter to be estimated, the GARCH-EGB2 model is able to account not only for the first, second and fourth moments of the conditional distribution of the variable of interest, as do popular Gaussian GARCH and GARCH-t models, it is also able to accommodate the third moment and high peakedness. Although economic theory suggests skewness and high peakedness should be common to exchange rates, they have been largely ignored in empirical work to date. EGB2 incorporates the normal distribution as a limiting case when p=q approaches infinity. It is symmetric for p=q and is positively (negatively) skewed for p > q (p < q) for σ > 0; the skewness results reverse for σ < 0. The EGB2 can accommodate coefficient of skewness values between -2 and 2 and coefficient of kurtosis values up to 9 (McDonald 1991), which will suffice for most data series, in particular the exchange rate data we study in this paper.

III. An Empirical Application to Six Daily Exchange Rates

The data are the daily noon spot U.S dollar exchange rate ($/local currency) for the

German deutsche mark (DM), British pound(£), Japanese yen(¥), French franc (FF), Belgian franc (BF) and Italian lira (IL) over the period January 1, 1985, to November 21, 1996 (3016 observations per series), as reported by the Exchange Rate Service of the Pacific Data Center at the University of British Columbia (http://pacific.commerce.ubc.ca/xr/). To achieve stationarity, we use first-differenced exchange rate series.[6] With R > 0 (R < 0) indicating currency appreciation (depreciation), the data are of the form[7]

$$R_{i\,t} = \ln[S_{i,\,t}/S_{i,\,t-1}]*\,100$$

where $R_{i\,t}$ = percentage change in the U.S.$/LC exchange rate of currency i at period t.
$S_{i\,t}$ = foreign exchange rate of currency i at period t, expressed as U.S.$/LC.

Table 3 presents descriptive statistics for each exchange rate series, including the coefficients of skewness[8] and kurtosis[9], inter-percentile ranges ($f_{0.75}$-$f_{0.25}$ and $f_{0.6}$-$f_{0.4}$), the Jarque-Bera asymptotic normality test statistics, and Ljung-Box-Pierce portmanteau test statistics. The yen, pound and lira all show significant skewness. The former is likely attributable to permanent structural shocks that led to the yen's dramatic appreciation over the sample period. The negative skewness in the pound and lira series probably reflect the autumn 1992 speculative attacks that knocked those currencies out of the European monetary system's exchange rate mechanism (ERM). As we will see in section IV, the GARCH-EGB2 model is especially appealing for

---

[6] Unit root test results demonstrating each series is I(1) are available from the authors.

[7] There is no adjustment is made for the weekend or holiday effects, so R indicates the exchange rate changes between two successive trading days.

[8] $E(R_t - \mu)^3/\sigma^3$, where μ is the mean and σ is the standard deviation.

[9] $E(R_t - \mu)^4/\sigma^4$, where μ is the mean and σ is the standard deviation.

currencies such as these, which exhibit significantly skewed percentage change distributions. The higher the coefficient of kurtosis (KUR), the less probability is concentrated around the mean, meaning that the distribution are more fat-tailed than normal distribution. For all currencies, the coefficients of kurtosis are greater than five and significantly different from the reference value of three drawn from the normal distribution. The high peakedness of each unconditional distribution is confirmed by inter-percentile ranges (e.g., $f_{\alpha 1} - f_{\alpha 2}$ indicates the range between the cumulative probabilities α_1 and α_2). Given α_1 and α_2, the lower the value of $f_{\alpha 1} - f_{\alpha 2}$, the higher the peakedness of the distribution. Across all six exchange rates, the value $f_{0.75}$-$f_{0.25}$ is uniformly less than 1.36, the reference range corresponding to the normal distribution. The unconditional distributions of these exchange rates have higher peaks than does a normal distribution around the central 50% of probability mass. The high peakedness is corroborated over the narrower interval $f_{0.6} - f_{0.4}$, for which all exchange rates' ranges are less than 0.5, the inter-percentile value of the standard normal over its central 20% of probability mass. Given skewness, leptokurtosis and high peakedness, it is not surprising that the null hypothesis of normality is strongly rejected by the Jarque-Bera (JB) asymptotic test for each exchange rate. Table 3 also presents the Ljung-Box test statistics for autocorrelation in R_{it} at a lag of 30 trading days (Q(30)), and in squared exchange rates changes ($Q^2(30)$), the latter serving as a test for GARCH effects. All the series exhibit GARCH effects. In summary, the descriptive statistics of Table 3 suggest the unconditional distributions of daily exchange rate changes are generally far from the traditional Gaussian assumption and also exhibit heteroskedasticity of the GARCH form.[10] These results are

[10] These characteristics are evident as well in a graphical appendix available by request from the authors.

consistent with previous empirical findings and economic theory (Boothe and Glassman 1987, Hsieh 1988, Wang, Barrett and Fawson 1996).

We began estimation by identifying and estimating a common ARMA process for the stationary R_{it}. First, Box-Jenkins techniques were used to reduce the set of prospective ARMA specifications. Next, we further narrowed the pool of possible models to those having a p-value for the Ljung-Box portmanteau Q(30) statistic of greater than 0.3, a significance level clearly supporting the assumption of white noise. Finally, we chose the ARMA specification having the lowest Schwarz Bayesian criterion (SBC) value from among the candidate models having passed the Box-Jenkins and Q(30) screens. In other words, the Ljung-Box Q statistic was used to identify a few possible models and then the information criterion (SBC) selected the final ARMA specification for the conditional mean equation.

Table 4 reports Ljung-Box portmanteau statistics for the squared standardized residual (z_t) for all currencies under homoskedastic (HOMO), Gaussian GARCH (GARCH), Student-t GARCH (GARCH-t) and GARCH-EGB2 specifications. The p-values of the test statistics (reported in brackets) clearly suggest that each of the GARCH specifications satisfactorily eliminates the serial correlation in conditional variance found in the homoskedastic model. Accommodating volatility clustering is not difficult in these exchange rate data.

While all the GARCH models appear to accommodate second-order serial correlation successfully, the issue of nonnormality remains. Skewness and excess kurtosis of the standardized residuals persist in all the Gaussian models (m_3^{GARCH} and m_4^{GARCH} in Table 5), although the leptokurtic characteristics (m_4^{GARCH}) have generally been muted somewhat relative to the homoskedastic model (m_4^{HOMO}).[11] As discussed in section II, Gaussian GARCH models inherently capture some unconditional leptokurtosis, but not always enough to represent exchange rate series accurately.

As a result, many applied econometricians have turned to using the student-t conditional error distribution to account for leptokurtosis. As measured by maximal log-likelihood values or likelihood ratio test statistics, the GARCH-t and GARCH-EGB2 models appear uniformly superior to the Gaussian GARCH model in fitting these exchange rate series (Table 6).[12] There are considerable gains to be had from capturing GARCH effects; our point is that there are considerable further gains to be had from accommodating nonnormal innovations. Note also the low estimated values in Table 5 for the degree of freedom parameter, υ, in each of the GARCH-t models ($\upsilon \geq 30$ indicates asymptotic normality). Moreover, Table 5 shows that the conditional kurtosis values predicted by the estimated values for the shape parameters (ν in the case of GARCH-t, p and q for GARCH-EGB2), ϕ_4^t and ϕ_4^{EGB2} under GARCH-t and GARCH-EGB2 assumptions, respectively, are reasonably close to the kurtosis of the standardized residuals (m_4^t and m_4^{EGB2}) for all series except the Japanese yen, where the GARCH-EGB2 model still performs very well (far better than the GARCH-t). This suggests that both the conditional Student-t and EGB2 distributions satisfactorily capture the leptokurtosis of exchange rate movements, permitting the applied econometrician significant gains in estimation accuracy.

The most important weaknesses remaining in the GARCH-t specification, however, are the apparent asymmetry of the standardized residuals ($m_3^t \neq 0$ in Table 5) and the high peakedness

---

[11] Milhϕj (1987), Hsieh (1989), and MaCurdy and Morgan (1987) found similar results.

[12] The Gaussian GARCH models are uniformly preferable to a homoskedastic null.

observed in the data (Table 3), which the Student-t distribution will systematically fail to capture. By contrast, the estimated coefficients of skewness in the standardized residuals of the GARCH-EGB2 model (m_3^{EGB2}) are reasonably close to the predicted coefficients implied by the estimated distribution parameters p and q (ϕ_3^{EGB2}). Unlike the conditional distributions commonly assumed in GARCH modeling — the normal and the Student-t — the more flexible EGB2 distribution appears to capture well all of the higher-order moments (skewness and kurtosis) observed in exchange rate series.

The superiority of the GARCH-EGB2 model in capturing the high peakedness inherent to most exchange rate series is most evident graphically. Figures 1-6 show paired plots of the densities of the observed standardized residuals and the corresponding predictions derived from the estimated shape parameters of GARCH-t (the top of each pair) and GARCH-EGB2 (the bottom of each pair) models. These figures clearly show that the observed standardized residuals generated by the GARCH-t model vary considerably from their assumed distribution, in particular exhibiting high peakedness, asymmetry, or both. The empirical density plots for the standardized residuals of the GARCH-EGB2 model, by contrast, are in remarkable agreement with the estimated EGB2 probability density functions for each exchange rate. While the GARCH-t model is only able to account for the fat tails dimension of nonnormality, the more flexible GARCH-EGB2 model can also accommodate skewness and high peakedness, which economic theory suggests are likely important features of exchange rate series, and perhaps other asset and commodity price series as well.

While both the Student-t and the EGB2 nest within them the normal distribution, enabling the likelihood ratio tests used in Table 6, GARCH-t and GARCH-EGB2 are not nested within each other, so some other criterion must be used to test formally the null hypothesis that the two models are equivalent in these data. An appropriate option is the likelihood dominance criterion (LDC) proposed by Pollak and Wales (1991), which offers an approach to nonnested model selection consistent with the conventional inferential approach to hypothesis testing. The idea of LDC is to nest two nonnested competing models — H_1 and H_2 — within a fictive composite and then consider a set of admissible composite parametric sizes. In most applications, the largest interesting sizes of the composite should range from one parameter more than the larger hypothesis (n_2+1) to one parameter more than the sum of the number of parameters in two hypothesis (n_1+n_2+1). In this sense, the LDC model selections rules are as follows:

(i) LDC prefers H_1 to H_2 if $L_2 - L_1 < [\ C(n_2 + 1) - C(n_1 + 1)\ ] /2$.
(ii) LDC is indecisive between H_1 and H_2 if $[\ C(n_2 - n_1 + 1) - C(1)\ ] /2 > L_2 - L_1 > [\ C(n_2 + 1) - C(n_1 + 1)\ ] /2$.
(iii) LDC prefers H_2 to H_1 if $L_2 - L_1 > [\ C(n_2 - n_1 + 1) - C(1)\ ] /2$.

where L_1, L_2 denote the maximum log likelihood values corresponding to the two models, and n_1 and n_2 are the numbers of parameters in H_1 and H_2, respectively. LDC also assumes that $n_1 < n_2$. $C(\gamma)$ is the critical values of the chi-square distribution with γ degree of freedom at the pre-specified significance level. In most practical situations the LDC proves decisive for model selection.

Because GARCH-EGB2 always involves one more parameter than the GARCH-t model, the value of the criterion ($[\ C(n_2 - n_1 + 1) - C(1)\ ] /2$) is fixed at 1.29 for the 1% significance level. For all six exchange rates, GARCH-EGB2 dominates GARCH-t in the LDC sense (Table 7). The superiority of the GARCH-EGB2 specification is especially evident in modeling the pound, yen and lira, each of which has a significantly skewed unconditional distribution. The returns to employing a more flexible conditional distribution appear greatest for asymmetric distributions.

The superiority of the GARCH-EGB2 model is further confirmed by goodness-of-fit test statistics (Table 8). These test statistics compare the frequency distribution of the residuals from the estimated models with the distribution predicted by the estimated distribution shape parameters, η_t. For each exchange rate, the test statistics for the GARCH-EGB2 model is far less than that from the GARCH-t model.[13]

Finally, Table 9 reports the parameter estimates and associated standard errors of the GARCH-EGB2 models fit to each exchange rate series. We report two standard errors for each estimated parameter: a conventional standard error and a White robust standard errors (White 1982). White showed that if the model is correctly specified, the different methods to compute the covariance matrix of the parameter will be (stochastically) the same. Our results routinely yield nearly identical standard error estimates by either method, providing informal evidence that the GARCH-EGB2 model is correctly specified.

IV. Out-of-Sample Forecast Performance

In addition to in-sample validation of the superior performance of the GARCH-EGB2 model, we compare its out-of-sample forecast performance against the GARCH-t model. One, seven, fourteen and thirty day ahead forecasts were constructed for both models. The choice of multiple forecast horizons facilitates a comparative assessment of forecast performance across

[13] The goodness-of-fit test has an asymptotic chi-squared distribution. For each estimated GARCH-EGB2 model, the test statistics support rejecting the hypothesis that the residuals are drawn from the assumed distribution at conventional levels of significance. However, this is common in large sample sizes, where this is a low power test.

models which potentially have different intertemporal characteristics. In particular, since the GARCH-EGB2 better captures the high peakedness of the series' distribution but longer forecast horizons increase the probability of drift, we expected the GARCH-EGB2 model's in-sample dominance to carry only to short horizons out-of-sample. At longer horizons, the GARCH-t model may forecast better precisely because of its inability to capture high peakedness leaves greater mass on its tails. Given that most empirical asset price forecasting applications are over short-run horizons, however, daily forecasts hold the greatest interest to practitioners.

For each currency, both the GARCH-EGB2 and GARCH-t models were estimated over the in-sample period (January 1, 1985, to November 21, 1996). Then a forecast was generated for each forecast horizon. The next day's observation was then added to the sample and the models were reestimated to generate the next set of forecasts. This sequential updating process was repeated for 130 new observations (November 22, 1996, to May 29, 1997), generating a sample of 130 forecasts for the one day horizon, 123 forecasts for the seven day horizon, 116 for the fourteen day, and 100 for the thirty day horizon..[14]

Our assessment of out-of-sample forecast performance is based on four statistics: the mean absolute error (ABSE), the root mean square error (RMSE), and the Diebold-Mariano (1995) test statistics (DM) for both the absolute error (DM^A) and squared error (DM^S) loss functions (Table 10).[15] A positive (negative) value of the DM statistic indicates better (worse) forecast performance by the GARCH-EGB2 model, as compared against the GARCH-t. While

[14] This forecast methodology is consistent with the idea that agents update forecasts as new observations become available.

[15] Stekler's (1987) rank-based test was also used to compare the forecast performance across all horizons. The results were qualitatively similar to the DM test results.

the evidence is clearly mixed and the statistical significance of the forecast differences is weak — as reflected in the bracketed p-values for the DM statistics — GARCH-EGB2 outperforms GARCH-t in four of the six currencies (German mark, British pound, French franc, and Italian lira) at the one-day forecast horizon by every measure. Comparative forecast ability is of greatest interest over short time horizons, particularly given that a forecaster would generally want to regularly update predictions based on model with a flexible error distribution. Longer run forecasts would be of lesser interest. At seven, fourteen, and thirty-day horizons, the GARCH-t forecast of the German mark and French franc outperforms the GARCH-EGB2 forecast. Like the German mark and the French franc, the Japanese yen forecast performance switches dominant models, but in the opposite direction, with the GARCH-t forecast dominating at the one and seven-day horizon, mixed results (based on the choice of loss function) at the fourteen day horizon, and GARCH-EGB2 dominating at the thirty-day horizon. The currencies which provide reasonably consistent results across forecast horizons are the British pound and Italian lira (GARCH-EGB2 dominant) and the Belgian franc (GARCH-t dominant). In summary, the GARCH-EGB2 is marginally better at out-of-sample forecasting at the one-day horizon, the GARCH-t is marginally better at the seven and fourteen-day horizon, and mixed results are evident at the thirty-day horizon.

V. Conclusions

Although GARCH modeling based on normal or student-t conditional distributions has proved useful in capturing the volatility clustering and leptokurtosis commonly present in asset price series, it cannot accommodate other commonly observed stylized effects in high frequency exchange rate data, notably high peakedness and skewness. Since economic theory suggests these are important statistical characteristics of the underlying series, we propose a GARCH model based on the more flexible EGB2 distribution. The GARCH-EGB2 specification can model either mesokurtic or leptokurtic data and can accommodate asymmetry, high peakedness, or both. An application to daily log changes in six major exchange rates over ten years reveals the GARCH-EGB2 model significantly outperforms commonly employed specifications in sample and is marginally better out-of-sample at the one-day forecast horizon. Since the improvements enjoyed due to employing a conditional EGB2 distribution are especially pronounced for unconditionally skewed data series, application to storable commodity price series appear especially promising.

References

Baillie R.T. and T.Bollerslev, "The message in daily exchange rate: a conditional-variance tale." *Journal of Business and Economic Statistics* 7, 1989, 297-305.

Baillie, R.T. and R.P. DeGennaro, "The Impact of Delivery Terms on Stock Return Volatility." *Journal of Financial Service Research* 3, 1990, 55-76.

Bollerslev, T., "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 3, 1986, 307-27.

Bollerslev, T., "A Conditional Heteroskedastic Time Series Model for Speculative Price and Rate of Return." *Review of Economics and Statistics* 9, 1987, 542-7.

Bollerslev, T., R.Y. Chou, and K.F. Kroner, "ARCH Modeling in Finance." *Journal of Econometrics* 52, 1992, 5-59.

Bollerslev, T., R.F. Engle, and D.B. Nelson, "ARCH Models," in R.F. Engle and D.L. McFadden, eds., *Handbook of Econometrics*, volume IV (Amsterdam: Elsevier Science, 1994).

Boothe, P. and D. Glassman, "The Statistical Distribution of Exchange Rates." *Journal of International Economics* 22, 1987, 297-319.

Davis, H. T., *Tables of the Higher Mathematical Functions*, Vol I and II, (Bloomington, IN: Principia Press, 1935).

Deb, P., "Finite Sample Properties of Maximum Likelihood and Quasi-Maximum Likelihood Estimators of EGARCH Models." *Econometric Reviews* 15, 1996, 51-68.

Diebold, F.X. and R. Mariano, "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13, 1995, 253-264.

Dornbusch, R., "Expectations and exchange rate dynamics." *Journal of Political Economy* 84, 1976, 1161-1176.

Engle, R.F., "Autoregressive Conditional Heteroskedasticity with the Estimates of the Variance of U.K. Inflation." *Econometrica* 50, 1982, 987-1008.

Friedman, M., "The Case for Flexible Exchange Rates." *Essays in Positive Economics* (Chicago: University of Chicago Press, 1953).

Hansen Bruce E., "Autoregressive Conditional Density Estimation." *International Economic Review* 35, 1994, 705-730.

Hsieh, D.A., “The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983.” *Journal of International Economics* 24, 1988, 129-145.

Hsieh, D.A., “Modeling Heteroskedasticity in Daily Foreign-Exchange Rates.” *Journal of Business and Economic Statistics* 7, 1989, 307-317.

Johnson, N.L. and S.Kotz, *Continuous univariate distributions,* Vol. 2 (New York: John D. Wiley, 1970).

Krugman, P.R., “A Model of Balance Of Payments Crises.” *Journal of Money, Credit and Banking* 11, 1979, 311-325.

Lee, S. and B.E. Hansen, “Asymptotic Theory for the GARCH(1,1) Quasi-maximum Likelihood Estimator.” *Econometric Theory* 10, 1994, 29-52.

MaCurdy, T.H. and L. Morgan, “Tests of the Martingale Hypothesis for Foreign Currency Futures with Time Varying Volatility.” *International Journal of Forecasting* 3, 1987, 131-48.

McDonald James B., "Parametric Models for Partially Adaptive Estimation with Skewed and Leptokurtic Residuals.” *Economics Letters* 37, 1991, 273-278.

McDonald J.B. and Y.J. Xu, "A Generalization of the Beta Distribution with Applications.” *Journal of Econometrics* 66, 1995, 133-152.

Milh∅j, A., “A Conditional Variance Model for Daily Deviations of an Exchange Rate.” *Journal of Business and Economic Statistics* 5, 1987, 99-103.

Obstfeld, M. and K. Rogoff, *Foundations of International Macroeconomics* (Cambridge: MIT Press, 1996).

Pagan, A.R. and H. Sabau, “On the Inconsistence of the MLE in Certain Heteroskedasticity Regression Model.” mimeo, 1987, University of Rochester.

Patil, G.P., M.T. Boswell and M.V. Ratnaparkhi, *Dictionary and classified bibliography of statistical distributions in scientific work: continuous univariate models* (Burtonsville, MD: International Cooperative Publishing House, Burtonsville, MD, 1984).

Peruga, R., "The Distributional Properties of Exchange Rate Changes Under a Peso Problem.” Ph.D. dissertation, University of California, San Diego, Dept. of Economics, 1988.

Pollak, R.A. and T. J.Wales, "The Likelihood Dominance Criterion.” *Journal of Econometrics* 47, 1991, 227-242.

Stekler, H.O., “Who forecasts better?” *Journal of Business and Economic Statistics* 5, 1987, 155-158.

Taylor, M.P., “The Economics of Exchange Rates.” *Journal of Economic Literature* 33, 1995, 13-47.

Wang, K.L., C.B. Barrett and C. Fawson, “Estimating East Asian Exchange Rates at Different Frequencies.” Utah State University Economic Research Institute Study Paper, 1996, #96-43.

White, H., "Maximum Likelihood estimation of misspecified models.” *Econometrica* 50, 1982, 1-25.

Table 1. The moments of the EGB2 distribution

Mean	$\delta + \sigma[\psi(p) - \psi(q)]$
Variance	$\sigma^2[\psi'(p) + \psi'(p)]$
Skewness	$(\psi''(p) - \psi''(q))/(\psi'(p) + \psi'(q))^{1.5}$
Kurtosis	$[(\psi'''(p) + \psi'''(q)) + 3(\psi'(p) + \psi'(q))^2]/(\psi'(p) + \psi'(q))^2$

Note: where $\psi(\ )$, $\psi'(\ )$, $\psi''(\ )$, and $\psi'''(\ )$ are digamma, trigamma, tetragamma, pentagamma functions, respectively (Davis, 1935).

Table 2. Comparison of EGB and EGB2 Estimation Results

	DM	£	¥	FF	BF	IL
c	0.9999	1.0000	1.0000	1.0000	1.0000	1.0000
LLH^{EGB}	3244.41	3122.94	3087.57	3114.19	3225.66	3109.09
LLH^{EGB2}	3235.40	3078.89	3047.50	3097.70	3210.02	3090.04
LR	18.02*	88.10*	80.14*	32.98*	31.28*	38.10*

c is the parameter estimated from EGB distribution.

LLH^{EGB} and LLH^{EGB2} represent the negative of the maximal log-likelihood value of the models under the EGB and EGB2 distributions, respectively.

LR represents the likelihood ratio test statistic of GARCH-EGB2 model against the corresponding GARCH-EGB model.

* denotes statistical significance at the 1% level.

Table 3. Descriptive Statistics

	SK	Kur	$f_{0.75}$-$f_{0.25}$	$f_{0.6}$-$f_{0.4}$	JB	Q(30)	Q^2(30)
DM	-0.037	5.1	1.13	0.41	566.73*	30.03	391.36
	(0.045)	(0.09)				[0.46]	[0.00]
£	-0.12	5.2	1.08	0.40	604.92*	37.98	451.30
	(0.045)	(0.09)				[0.15]	[0.00]
¥	0.286	6.1	1.02	0.38	1282.9*	37.80	233.98
	(0.045)	(0.09)				[0.16]	[0.00]
FF	0.02	6.0	1.14	0.39	511.89*	41.16	391.67
	(0.045)	(0.09)				[0.08]	0.00]
BF	0.024	5.0	1.12	0.41	521.53*	41.69	370.85
	(0.045)	(0.09)				[0.08]	[0.00]
IL	-0.616	8.8	1.14	0.41	4377.65*	33.39	641.75
	(0.045)	(0.09)				[0.31]	[0.00]

SK = coefficient of skewness.
KUR = coefficient of kurtosis (the value for the normal distribution is 3.0)
The asymptotic standard errors of SK and KUR are reported in parentheses and computed as $(6/T)^{0.5}$ and $(24/T)^{0.5}$, respectively.
JB = Jarque-Bera normality test statistic.
Q and Q^2 represent the Ljung-Box test statistics for up to 30th order serial correlation for each exchange rate series. Similar results obtain at different orders. P-values against the null hypothesis of white noise are reported in brackets.
* denotes statistical significance at the 1% level.

Table 4. Tests for Serially Correlated Conditional Variance

	DM	£	¥	FF	BF	IL
HOMO	403.20 [0.00]	447.50 [0.00]	239.53 [0.00]	404.10 [0.00]	370.75 [0.00]	620.03 [0.00]
GARCH	29.28 [0.50]	24.36 [0.76]	28.72 [0.53]	25.50 [0.70]	35.82 [0.21]	23.67 [0.79]
GARCH-t	28.71 [0.53]	24.46 [0.75]	29.06 [0.52]	25.41 [0.71]	34.73 [0.25]	26.22 [0.66]
GARCH-EGB2	28.96 [0.52]	24.64 [0.74]	29.52 [0.49]	25.50 [0.70]	35.01 [0.24]	25.03 [0.72]

The figure in brackets is the p-value of the Ljung-Box Q(30) test against the null hypothesis of no serial correlation.

Table 5. Skewness and kurtosis of sample standardized residuals and predicted values

	DM	¥	¥	FF	BF	IL
HOMO						
m_3^{HOMO}	-0.034	-0.100	0.281	0.026	0.036	-0.586
m_4^{HOMO}	5.054	5.053	6.099	4.943	4.966	8.550
GARCH						
m_3^{GARCH}	0.075	-0.110	0.464	0.094	0.114	-0.118
m_4^{GARCH}	4.419	4.365	6.154	4.350	4.402	4.745
GARCH-t						
m_3^{t}	0.083	-0.108	0.502	0.103	0.127	-0.160
m_4^{t}	4.451	4.345	6.373	4.398	4.462	5.172
ϕ_4^{t}	4.433	4.652	15.49	4.361	4.445	4.680
ν	6.093	5.816	4.240	6.210	6.075	5.785
GARCH-EGB2						
m_3^{EGB2}	0.083	-0.115	0.502	0.102	0.125	-0.148
ϕ_3^{EGB2}	0.088	-0.071	0.326	0.087	0.076	-0.025
m_4^{EGB2}	4.451	4.395	6.355	4.394	4.457	5.065
ϕ_4^{EGB2}	4.584	4.778	5.356	4.541	4.601	4.909

m_3 is the coefficient of skewness of the standardized residuals from the estimated model.

m_4 is the coefficient of kurtosis of the standardized residuals from the estimated model.

For each model, the asymptotic standard error of the coefficients of skewness and kurtosis are 0.045 and 0.089, respectively.

ν is the degree of freedom estimate from GARCH-t model.

ϕ_4^{t} is the predicted kurtosis coefficient of student-t distribution = $3 (\upsilon -2) / (\upsilon - 4)$, $\upsilon > 4$.

ϕ_3^{EGB2} is predicted skewness coefficient of EGB2 distribution = $[\psi''(p) - \psi''(q)] / [\psi'(p) + \psi'(q)]^{1.5}$

ϕ_4^{EGB2} is predicted skewness coefficient of EGB2 distribution = $\{[\psi'''(p) + \psi'''(q)] + 3[\psi'(p) + \psi'(q)]^2\} / (\psi'(p) + \psi'(q))^2$

Table 6. Comparisons of Alternative Specifications

	HOMO	GARCH		GARCH-t		GARCH-EGB2	
	LLH^{HOMO}	LLH^{GARCH}	LR^{GARCH}	LLH^{t}	LR^{t}	LLH^{EGB2}	LR^{EGB2}
DM	-3342.19	-3228.52	227.34*	-3163.46	130.12*	-3161.31	134.42*
£	-3200.57	-3059.32	282.50*	-2991.03	136.58*	-2987.83	142.98*
¥	-3236.48	-3121.97	229.02*	-2970.24	303.46*	-2963.68	316.58*
FF	-3197.86	-3085.85	224.02*	-3024.20	123.30*	-3022.09	127.52*
BF	-3309.58	-3201.02	217.12*	-3137.07	127.90*	-3134.82	132.40*
IL	-3251.40	-3078.24	346.32*	-3005.76	144.96*	-3002.04	152.40*

LLH^{HOMO}, LLH^{GARCH}, LLH^{t} and LLH^{EGB2} represent the maximal log-likelihood value of HOMO, Gaussian GARCH, GARCH-t and GARCH-EGB2 models, respectively.

LR^{GARCH} indicates the likelihood ratio test statistic for the Gaussian GARCH models against the HOMO model.

The LR^{t} and LR^{EGB2} statistics are for the GARCH-t and GARCH-EGB2 against the corresponding Gaussian GARCH models, respectively.

* denotes statistical significance at the 1% level (using the $\chi^2(3)$, $\chi^2(1)$ and $\chi^2(2)$ distributions for the LR^{GARCH}, LR^{t}, and LR^{EGB2}, respectively).

Table 7. Comparisons of GARCH-t and GARCH-EGB2 Models by LDC

	n^t	n^{EGB2}	$LLH^{EGB2} - LLH^t$	$[C(n_2 + 1) - C(n_1 + 1)]/2$	$[C(n_2 - n_1 + 1) - C(1)]/2$
DM	6	7	2.15	0.81	1.29
£	6	7	3.20	0.81	1.29
¥	6	7	6.56	0.81	1.29
FF	6	7	2.11	0.81	1.29
BF	8	9	2.26	0.77	1.29
IL	8	9	3.72	0.77	1.29

n^t and n^{EGB2} are the number of parameters in the GARCH-t and GARCH-EGB2 models, respectively.

$LLH^{EGB2} - LLH^t$ is the difference of log-likelihood value between the GARCH-EGB2 and GARCH-t models.

$[C(n_2 + 1) - C(n_1 + 1)]/2$ is the critical value to determine if GARCH-t is preferred to GARCH-EGB2 when LLH^t is greater than LLH^{EGB2}; whereas $[C(n_2 - n_1 + 1) - C(1)]/2$ is used to determine if GARCH-EGB2 is preferred to GARCH-t when LLH^{EGB2} is greater than LLH^t. The critical values are derived from the χ^2 distribution evaluated at the 1% significance level.

Table 8. Chi-Square Goodness of Fit Test Statistics

Currency	*GARCH-t test statistic*	*GARCH-EGB2 test statistic*
DM	194.41	102.79
£	197.00	83.45
¥	364.63	87.04
FF	179.64	81.32
BF	194.76	99.88
IL	239.28	99.82

The test statistics are obtained by evaluating $\sum(f_i\text{-}F_i)^2 / F_i$, where f_i is the observed count frequency of the standardized residuals, F_i is the predicted count frequency, and i=1,...,40. The χ^2 critical value at the 1% level is 63.69.

Table 9. Parameter Estimates from GARCH-EGB2 Models

	DM	£	¥	FF	BF	IL
Conditional mean equation parameters						
C	0.021 (0.012)^{SE} (0.012)^{RSE}	0.020 (0.012)^{SE} (0.011)^{RSE}	0.023 (0.012)^{SE} (0.012)^{RSE}	0.020 (0.012)^{SE} (0.012)^{RSE}	0.024 (0.013)^{SE} (0.013)^{RSE}	0.014 (0.012)^{SE} (0.012)^{RSE}
AR(1)	0.378 (0.183)^{SE} (0.186)^{RSE}	0.345 (0.183)^{SE} (0.191)^{RSE}	0.040 (0.178)^{SE} (0.189)^{RSE}	0.392 (0.184)^{SE} (0.186)^{RSE}		
AR(3)						0.218 (0.179)^{SE} (0.180)^{RSE}
AR(6)					-0.248 (0.179)^{SE} (0.183)^{RSE}	-0.189 (0.174)^{SE} (0.171)^{RSE}
AR(8)					-0.319 (0.177)^{SE} (0.179)^{RSE}	
MA(1)					0.033 (0.019)^{SE} (0.019)^{RSE}	0.006 (0.019)^{SE} (0.019)^{RSE}
Conditional variance equation parameters						
w	0.015 (0.005)^{SE} (0.006)^{RSE}	0.01 (.)^{SE} (.)^{RSE}	0.011 (0.004)^{SE} (0.006)^{RSE}	0.016 (0.006)^{SE} (0.007)^{RSE}	0.015 (0.006)^{SE} (0.007)^{RSE}	0.022 (0.007)^{SE} (0.008)^{RSE}
α1	0.908 (0.018)^{SE} (0.020)^{RSE}	0.920 (0.008)^{SE} (0.009)^{RSE}	0.922 (0.018)^{SE} (0.026)^{RSE}	0.897 (0.021)^{SE} (0.024)^{RSE}	0.908 (0.020)^{SE} (0.024)^{RSE}	0.876 (0.025)^{SE} (0.030)^{RSE}
β1	0.065 (0.012)^{SE} (0.013)^{RSE}	0.061 (0.010)^{SE} (0.010)^{RSE}	0.058 (0.013)^{SE} (0.017)^{RSE}	0.071 (0.014)^{SE} (0.015)^{RSE}	0.065 (0.013)^{SE} (0.014)^{RSE}	0.081 (0.016)^{SE} (0.019)^{RSE}
Distribution parameters						
p	0.746 (0.122)^{SE} (0.118)^{RSE}	0.596 (0.099)^{SE} (0.097)^{RSE}	0.425 (0.077)^{SE} (0.077)^{RSE}	0.775 (0.128)^{SE} (0.123)^{RSE}	0.730 (0.122)^{SE} (0.118)^{RSE}	0.538 (0.108)^{SE} (0.131)^{RSE}
q	0.698 (0.112)^{SE} (0.109)^{RSE}	0.625 (0.106)^{SE} (0.103)^{RSE}	0.351 (0.061)^{SE} (0.061)^{RSE}	0.724 (0.117)^{SE} (0.112)^{RSE}	0.690 (0.113)^{SE} (0.110)^{RSE}	0.548 (0.112)^{SE} (0.135)^{RSE}

Standard errors reported in parentheses. ( )^{SE} indicate the conventional standard error, while ( )^{RSE} is the White robust standard error.

(.) indicates the standard error cannot be estimated because the parameter estimate lies on the boundary of the feasible parameter space.

Table 10.Comparisons of GARCH-t and GARCH-EGB2 Models for Multiperiod Ahead Forecasts

		1-day horizon		7-day horizon		14-day horizon		30-day horizon	
		Student-t	EGB2	Student-t	EGB2	Student-t	EGB2	Student-t	EGB2
Mark	ABSE	0.554	0.552	0.529	0.530	0.517	0.518	0.517	0.519
	DM^A	1.28 [0.22]		**-1.29 [0.20]**		**-1.30 [0.20]**		**-1.31 [0.19]**	
	RMSE	0.714	0.711	0.676	0.677	0.663	0.663	0.660	0.661
	DM^S	**1.33 [0.19]**		-1.12 [0.26]		-1.13 [0.26]		-1.15 [0.25]	
£	ABSE	0.525	0.515	0.482	0.481	0.469	0.469	0.454	0.453
	DM^A	**1.35 [0.18]**		0.74 [0.46]		0.88 [0.38]		1.03 [0.30]	
	RMSE	0.674	0.660	0.624	0.623	0.576	0.575	0.550	0.549
	DM^S	**1.36 [0.18]**		0.75 [0.45]		0.68 [0.50]		0.92 [0.36]	
¥	ABSE	0.346	0.348	0.496	0.498	0.507	0.509	0.527	0.529
	DM^A	-0.88 [0.38]		-1.23 [0.22]		-1.19 [0.23]		0.18 [0.86]	
	RMSE	0.475	0.477	0.700	0.700	0.715	0.715	0.745	0.744
	DM^S	-1.04 [0.30]		-0.09 [0.92]		0.16 [0.87]		0.49 [0.63]	
FF	ABSE	0.534	0.532	0.505	0.507	0.499	0.500	0.494	0.495
	DM^A	**1.31 [0.19]**		-1.24 [0.21]		-1.27 [0.21]		**-1.29 [0.20]**	
	RMSE	0.682	0.679	0.646	0.647	0.640	0.641	0.628	0.629
	DM^S	**1.33 [0.19]**		-1.12 [0.26]		-1.15 [0.25]		-1.18 [0.24]	
BF	ABSE	0.574	0.578	0.566	0.567	0.535	0.535	0.542	0.543
	DM^A	**-1.34 [0.18]**		-0.27 [0.79]		-0.52 [0.6]		-1.05 [0.30]	
	RMSE	0.766	0.770	0.754	0.754	0.706	0.707	0.550	0.549
	DM^S	**-1.33 [0.19]**		-0.39 [0.70]		-0.90 [0.37]		-0.88 [0.38]	
IL	ABSE	0.510	0.509	0.470	0.471	0.444	0.442	0.459	0.457
	DM^A	1.00 [0.32]		-1.22 [0.22]		**1.36 [0.18]**		**1.40 [0.16]**	
	RMSE	0.701	0.700	0.682	0.682	0.670	0.669	0.675	0.674
	DM^S	0.12 [0.90]		0.811 [0.42]		0.988 [0.33]		1.09 [0.28]	

Student-t and EGB2 represent the GARCH-t and GARCH-EGB2 models, respectively.

DM^A and DM^S represent the Diebold and Mariano statistics for the absolute error and squared error loss functions, respectively A positive value indicatesGARCH-EGB2 outperforms GARCH-t. The bracketed statistics are p-values (the null hypothesis is equal forecast accuracy for both models).

Statistics in **bold** have p-values$\leq$ 0.2

Figures 1-6: Empirical and predicted distributions of standardized residuals from alternative GARCH models

Figure 1a: GARCH-t model for German DM

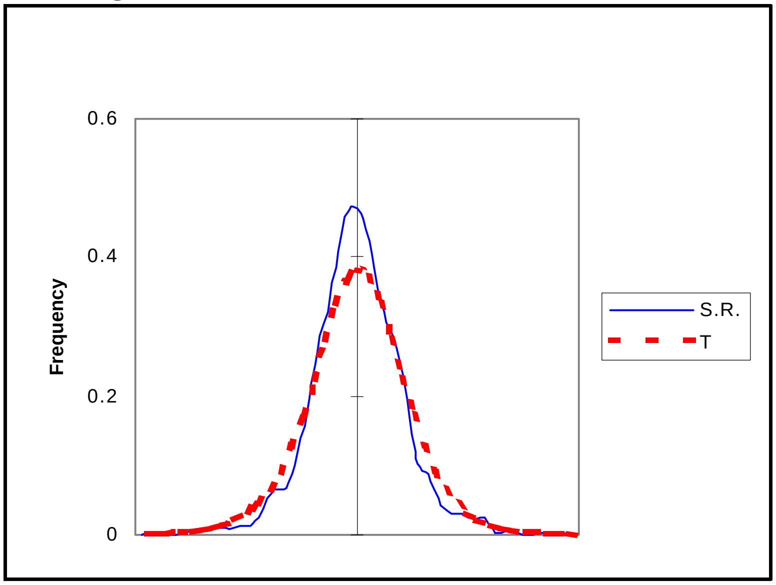

Figure 1b: GARCH-EGB2 model for German DM

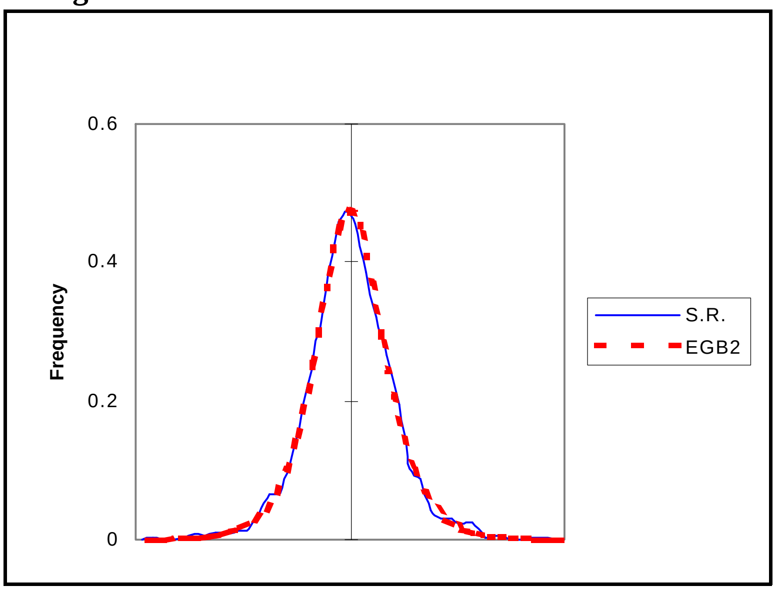

Figure 2a: GARCH-t model for British Pound

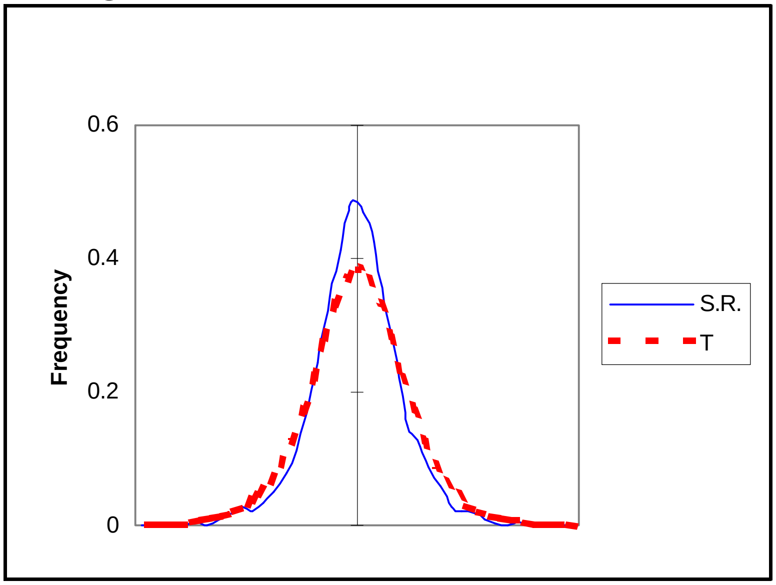

Figure 2b: GARCH-EGB2 model for British Pound

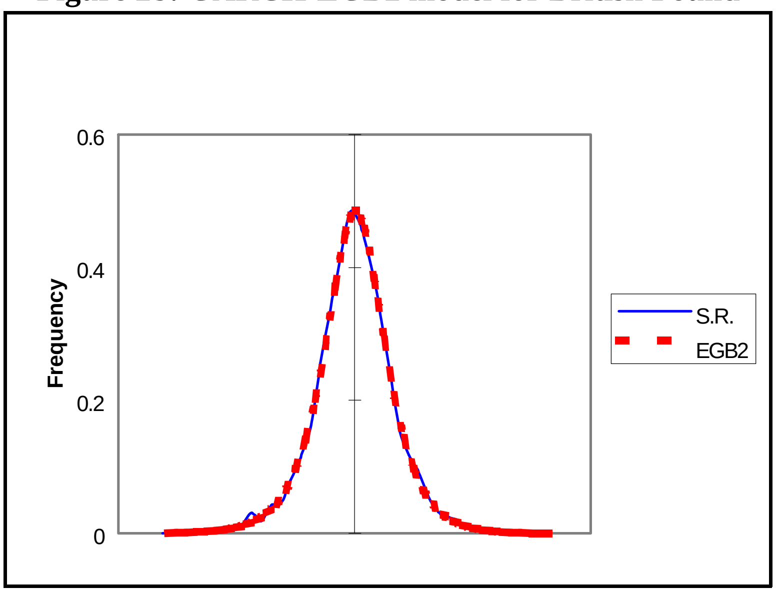

Figure 3a: GARCH-t model for Japanese Yen

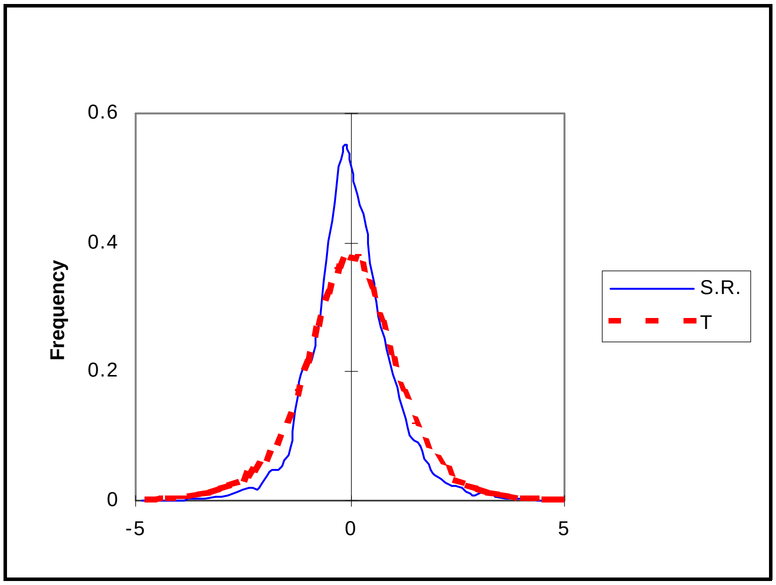

Figure 3b: GARCH-EGB2 model for Japanese Yen

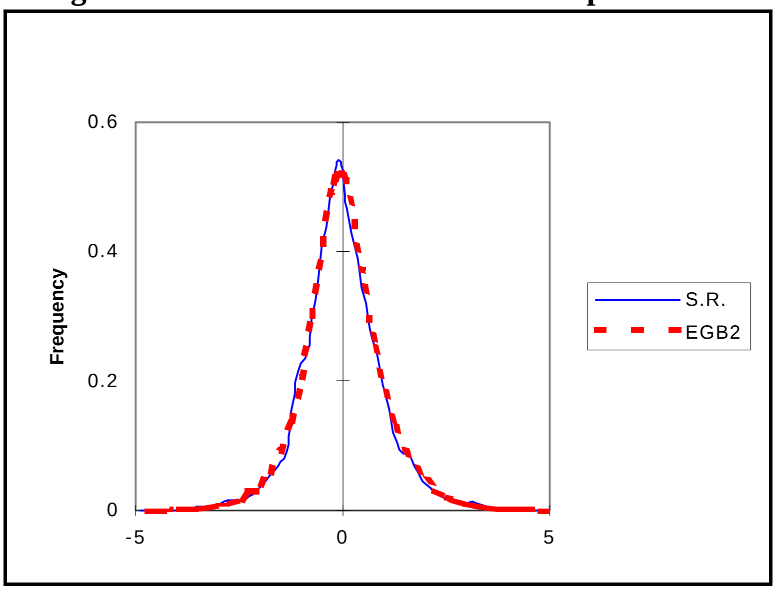

Figure 4a: GARCH-t model for French Franc

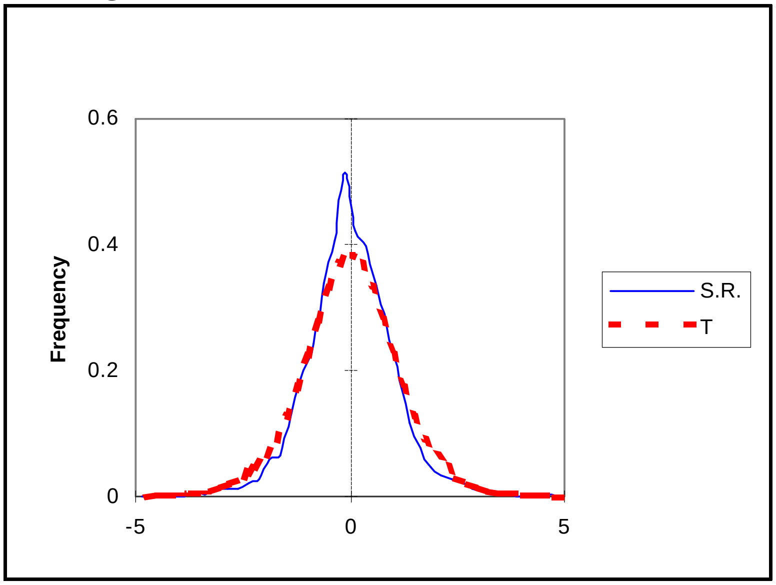

Figure 4b: GARCH-EGB2 model for French Franc

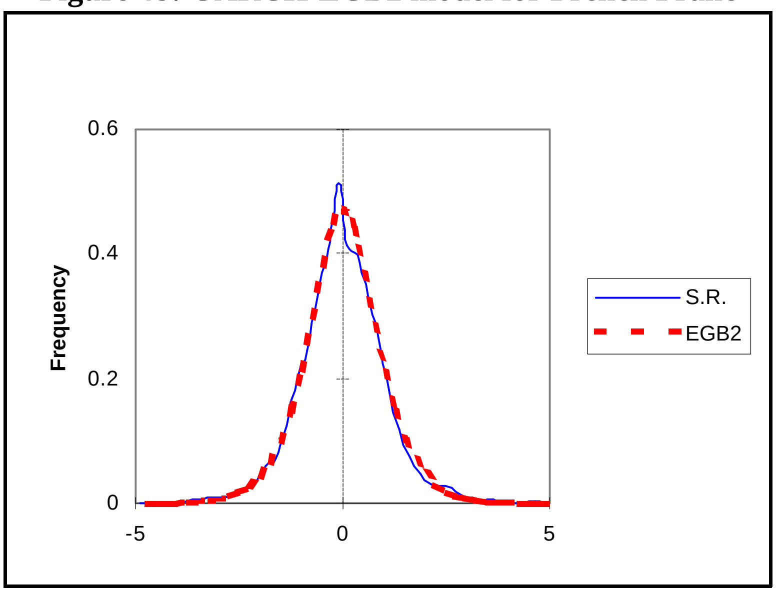

Figure 5a: GARCH-t model for Belgian Franc

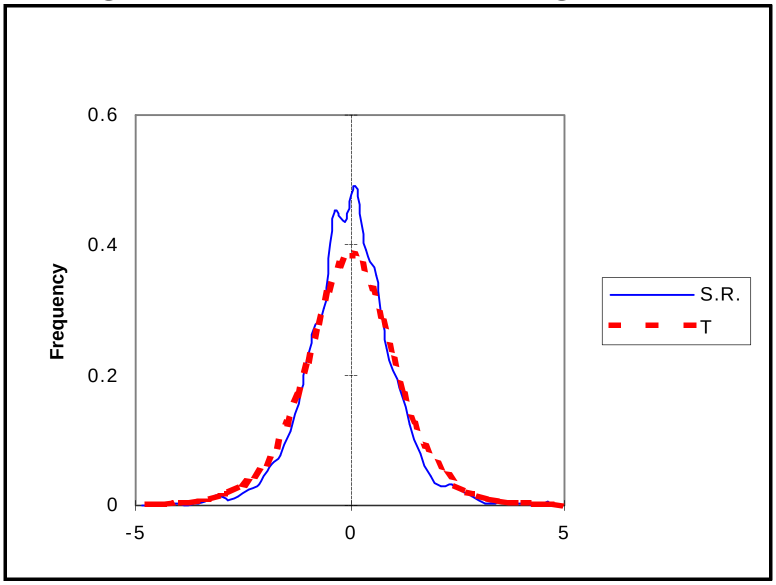

Figure 5b: GARCH-EGB2 model for Belgian Franc

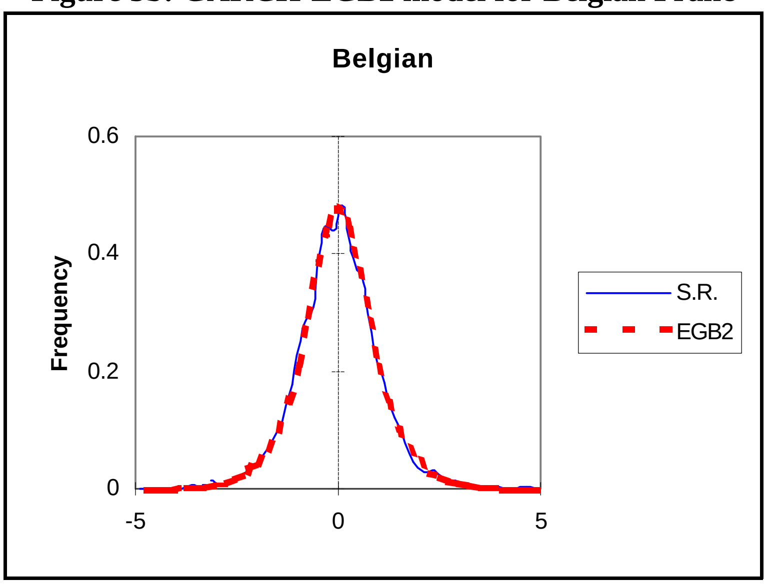

Figure 6a: GARCH-t model for Italian Lira

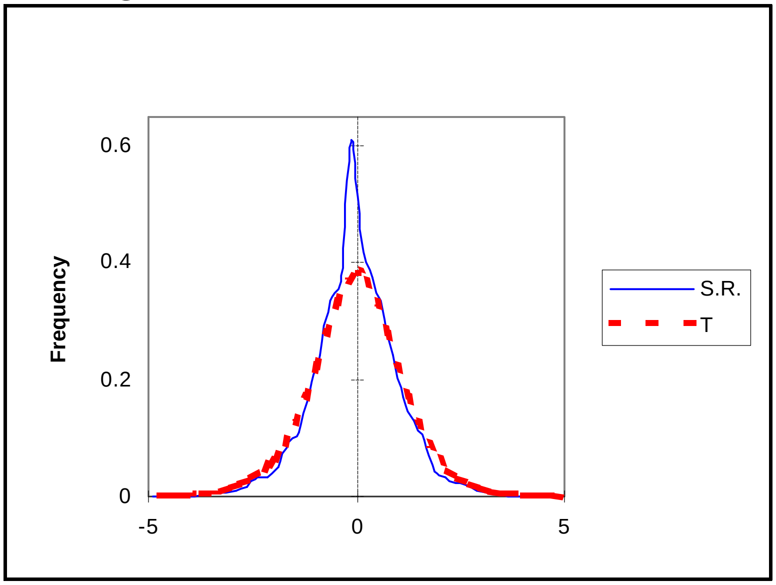

Figure 6b: GARCH-EGB2 model for Italian Lira

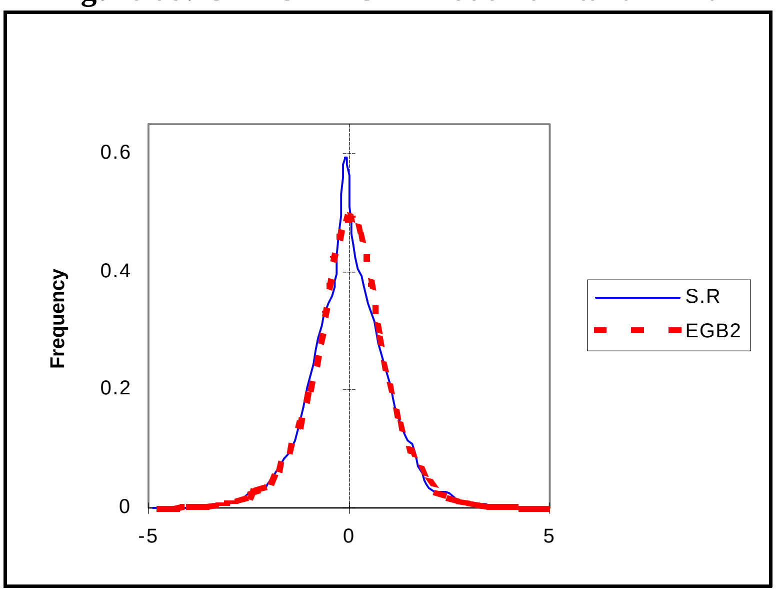

Technical Appendix

Following the traditional definition of a GARCH process, suppose that

$$\epsilon_t = h_t^{0.5} z_t \tag{A1}$$

where $\{\epsilon_t\}$ is the error term sequence from the conditional mean equation and $\{z_t\}$ is an i.i.d. sequence with zero mean and unit variance. Let h_t evolve according to a GARCH(1,1) process

$$h_t = \alpha_0 + \alpha_1 h_{t-1} + \beta_1 \epsilon_t^2 \tag{A2}$$

If z_t is drawn from a EGB2 distribution, then the density is given by

$$EGB2(\epsilon;\delta,\sigma,p,q) = \frac{\exp\frac{p(\epsilon-\delta)}{\sigma}}{|\sigma| B(p,q)\,(1+e^{\frac{\epsilon-\delta}{\sigma}})^{p+q}} \tag{A3}$$

The mean and variance of z are then as follows:

$$\mathrm{Var}(z) = \sigma^2 \left(\psi'(p) + \psi'(q)\right) = 1 \tag{A4}$$

$$E(z) = \delta + \sigma[\psi(p) - \psi(q)] = 0 \tag{A5}$$

Hence, solving for σ and δ in terms of Δand Ω

$$\sigma = \sqrt{\frac{1}{\psi'(p) + \psi'(q)}} = \sqrt{\frac{1}{\Omega}} \tag{A6}$$

$$\delta = -\sigma[\psi(p) - \psi(q)] = -\Delta\sqrt{\frac{1}{\Omega}} \tag{A7}$$

where:

$$\Delta = \psi(p) - \psi(q) \tag{A8}$$

$$\Omega = \psi'(p) + \psi'(q) \tag{A9}$$

Substituting those expressions for δ and σ back into the EGB2 distribution yields an EGB2 density function with zero mean and unit variance as

$$EGB2(z;p,q) = \frac{\sqrt{\Omega}\exp\left(p\left(z+\frac{\Delta}{\sqrt{\Omega}}\right)\sqrt{\Omega}\right)}{B(p,q)\left(1+\exp\left(\left(z+\frac{\Delta}{\sqrt{\Omega}}\right)\sqrt{\Omega}\right)\right)^{p+q}} \tag{A10}$$

According to the assumption (A1),

$$\epsilon_t = \frac{z_t}{\sqrt{h_t}} \tag{A11}$$

Changing the variable from z to ϵ as follows: (dz = dϵ/h)

$$EGB2(\epsilon;h,p,q) = \frac{\sqrt{\Omega}\exp(p(\frac{\epsilon}{\sqrt{h}} + \frac{\Delta}{\sqrt{\Omega}})\sqrt{\Gamma\Omega})}{\sqrt{h}B(p,q)(1+\exp((\frac{\epsilon}{\sqrt{h}} + \frac{\Delta}{\sqrt{\Omega}})\sqrt{\Gamma\Omega}))^{p+q}} \tag{A12}$$

Algebraic manipulation then yields

$$EGB2(\epsilon;h,p,q) = \frac{\sqrt{\Omega}\exp(p(\frac{\sqrt{\Omega}}{\sqrt{h}}\epsilon + \Delta))}{\sqrt{h}B(p,q)(1+\exp(\frac{\sqrt{\Omega}}{\sqrt{h}}\epsilon + \Delta)^{p+q}} \tag{A13}$$