

Forecasting Agricultural Commodity Prices with Asymmetric-Error GARCH Models

Octavio A. Ramírez and Mohamadou Fadiga

The performance of a proposed asymmetric-error GARCH model is evaluated in comparison to the normal-error- and Student-*t*-GARCH models through three applications involving forecasts of U.S. soybean, sorghum, and wheat prices. The applications illustrate the relative advantages of the proposed model specification when the error term is asymmetrically distributed, and provide improved probabilistic forecasts for the prices of these commodities.

Key words: GARCH, nonnormality, skewness, time-series forecasting, U.S. commodity prices

Introduction

Producing reliable forecasts is often a key objective in agricultural economics research. A reliable forecast should be unbiased or at least consistent, should provide a narrow confidence interval for the expected value of the economic variable of interest, and should incorporate confidence bands that adequately portray the likelihood of the variable's occurrences. Time-series models have been widely used for these purposes. Among them, the generalized autoregressive conditional heteroskedastic process (GARCH) (Bollerslev 1986) and its predecessor, the autoregressive conditional heteroskedastic process (ARCH) (Engle), have proven useful for modeling a variety of time-series phenomena because many time-series variables exhibit autocorrelation as well as dynamic heteroskedasticity. Some of these variables, however, are also nonnormally distributed.

Agricultural economics applications of standard GARCH models include analyses by Moss, and by Moss, Shonkwiler, and Ford. Bollerslev (1987) proposed a nonnormal-error GARCH model of speculative prices and rates of return based on the Student-*t* distribution (*t*-GARCH), which is leptokurtic but symmetric.

Yang and Brorsen, concerned with the nonnormality of daily cash prices, explored the use of a mixed diffusion-jump process, a deterministic chaos model, and the *t*-GARCH model to explain the stochastic behavior of these prices. They concluded that, while the *t*-GARCH model best explains daily cash price behavior, "it is not well calibrated" because it cannot explain all of the observed nonnormality (p. 714)—referring to the *t*-GARCH model's inability to account for the skewness in the distribution of cash prices.

As shown by Pagan and Sabau; Lee and Hansen; and Deb, in finite sample sizes, misspecification of the error-term distribution results in poor statistical properties, and an

Octavio A. Ramírez is associate professor and Mohamadou Fadiga is graduate research assistant, Department of Agricultural and Applied Economics, Texas Tech University, Lubbock, Texas. The authors thank two anonymous reviewers for their valuable comments and suggestions to improve this manuscript. Any remaining errors are the responsibility of the authors.

Review coordinated by Gary D. Thompson.

unreliable quasi-maximum likelihood estimator. Ramírez and Shonkwiler report that symmetric-error GARCH models tend to underestimate the true standard errors of the intercept and slope parameter estimators when the true underlying error-term distribution is asymmetric. Because most GARCH applications occur with small- or moderate-sized samples, a flexible specification which can accommodate both error-term skewness and kurtosis is important to improve the reliability of quasi-maximum likelihood estimation of GARCH models.

To address the problem of unreliable quasi-maximum likelihood estimation of GARCH models, Wang et al. recently proposed an asymmetric-error GARCH model based on the Exponential Generalized Beta 2 (EGB2) family of distributions and applied it in the modeling of exchange rates. According to McDonald and White, however, the EGB2 can only accommodate positive standardized (normal kurtosis = 0) kurtosis coefficients from 0 to 6, and skewness coefficients between -2 and 2. These ranges of kurtosis and skewness coefficients might be a limitation in certain applications.

In this study, an arguably more flexible asymmetric-error GARCH (A-GARCH) model is developed based on an expansion of the S_U family of distributions (Johnson, Kotz, and Balakrishnan) which can accommodate any level of leptokurtosis and right or left skewness—specifically, kurtosis coefficients from 0 to ∞ , and skewness coefficients from $-\infty$ to ∞ . The flexibility of the expanded S_U family is exploited to estimate an A-GARCH in which error-term kurtosis and skewness are allowed to systematically change through time. A-GARCH models of U.S. soybean, sorghum, and wheat prices are estimated and their forecasting performance is evaluated in comparison to Bollerslev's normal-error (N-GARCH) and t-GARCH models.¹

The A-GARCH(p, q) Model

An A-GARCH(p, q) model analogous to Bollerslev's (1986) N-GARCH(p, q) is written as:

$$(1) \quad y_t = \mathbf{x}_t' \mathbf{b} + \varepsilon_t, \quad \varepsilon_t \sim NN(0, h_t), \quad h_t = \alpha_0 + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^p \beta_i h_{t-i},$$

where $NN(0, h_t)$ represents a family of nonnormal distributions with mean zero and variance h_t . In the case of the t-GARCH models, $\varepsilon_t \sim NN(0, h_t)$ is assumed to follow a Student- t distribution with possibly variable degrees of freedom. In the EGB2, ε_t is assumed to follow an exponential generalized Beta of the second kind. In the A-GARCH(p, q) model, it is assumed ε_t can be adequately represented by the following expansion of the S_U family of distributions:

$$(2) \quad \varepsilon_t = \left[\{h_t / G(\Theta, \mu)\}^{1/2} \{ \sinh(\Theta v_t) - F(\Theta, \mu) \} \right] / \Theta, \quad v_t \sim N(\mu, 1),$$

$$F(\Theta, \mu) = E[\sinh(\Theta v_t)] = \exp(\Theta^2/2) \sinh(\Theta \mu), \quad \text{and}$$

$$G(\Theta, \mu) = \{ \exp(\Theta^2) - 1 \} \{ \exp(\Theta^2) \cosh(-2\Theta \mu) + 1 \} / 2\Theta^2,$$

where $-\infty < \Theta < \infty$, $-\infty < \mu < \infty$, and $0 < \sigma < \infty$ are distributional parameters.

¹ The A-GARCH is not directly compared to the EGB2-GARCH because the main means of comparison utilized in this analysis, i.e., the computation of confidence bands for the dependent variable occurrences, require simulation of random draws from the assumed error-term distribution. To our knowledge, there are no procedures to simulate draws from an EGB2.

From the results of Johnson, Kotz, and Balakrishnan, it follows:

$$(3) \quad E[\varepsilon_t] = 0, \quad \text{Var}[\varepsilon_t] = h_t,$$

$$(4) \quad \text{Skew}[\varepsilon_t] = S(\Theta, \mu) = -\frac{1}{4}w^{3/2}(w-1)^2[w(w+2)\sinh(3\Omega) + 3\sinh(\Omega)]/G(\Theta, \mu)^{3/2},$$

$$(5) \quad \text{Kurt}[\varepsilon_t] = K(\Theta, \mu) = \left\{ \frac{1}{8}(w-1)^2[w^2(w^4 + 2w^3 + 3w^2 - 3)\cosh(4\Omega) + 4w^2(w+2)\cosh(2\Omega) + 3(2w+1)]/G(\Theta, \mu)^2 \right\},$$

where $w = \exp(\Theta^2)$, $\Omega = -\Theta\mu$, and $\text{Skew}[\varepsilon_t]$ and $\text{Kurt}[\varepsilon_t]$ refer to the standardized measures of skewness and excess kurtosis. As in the previously discussed t - and EGB2-GARCH models, equation (3) implies $E[y_t]$ and $\text{Var}[\varepsilon_t]$ are the same as in Bollerslev's N-GARCH model. The error-term skewness and kurtosis are determined by the parameters Θ and μ . According to equations (4) and (5), if $\Theta = 0$ and $\mu = 0$, then $\text{Kurt}[\varepsilon_t] = 0$ and $\text{Skew}[\varepsilon_t] = 0$, ε_t follows a normal distribution, and the proposed A-GARCH reduces to Bollerslev's N-GARCH.

Higher absolute values of Θ cause increased positive kurtosis, up to infinity. If $\Theta \neq 0$ and $\mu = 0$, $\text{Kurt}[\varepsilon_t] \neq 0$ but $\text{Skew}[\varepsilon_t] = 0$, which means that ε_t follows a leptokurtic but symmetric distribution, such as the Student- t . If $\Theta \neq 0$ and $\mu > 0$, $\text{Kurt}[\varepsilon_t] \neq 0$ and $\text{Skew}[\varepsilon_t] > 0$, implying leptokurtosis and right-skewness, while $\mu < 0$ results in $\text{Skew}[\varepsilon_t] < 0$ (left-skewness). Further, note that as μ goes to ∞ ($-\infty$), $\text{Skew}[\varepsilon_t]$ also approaches ∞ ($-\infty$). Although higher absolute values of μ increase both skewness and kurtosis, kurtosis can be lowered by reducing Θ .

In practice, under error-term normality, Θ and μ will approach zero and the A-GARCH will approach Bollerslev's N-GARCH. This is another advantage of the proposed specification: the null hypothesis of an N-GARCH versus the alternative of a leptokurtic but symmetric-error A-GARCH can be directly tested by $H_0: \Theta = \mu = 0$ versus $H_A: \Theta > 0$. The null hypothesis of a leptokurtic but symmetric-error GARCH versus the alternative of a full (leptokurtic and skewed error) A-GARCH is given by $H_0: \Theta > 0, \mu = 0$ versus $H_A: \Theta > 0, \mu \neq 0$. The null of an N-GARCH versus the alternative of a full A-GARCH can be directly tested as well.

Another advantage of the A-GARCH is that the degree of skewness and kurtosis of the error-term distribution can be assumed variable across observations without interfering with the estimation of the linear regression and GARCH-process parameters. This is achieved by making either Θ or μ , or both, a function of time or any other potentially relevant factor. Given equations (1) and (2), the concentrated log-likelihood function for estimating the A-GARCH model is obtained using the transformation technique (Mood, Graybill, and Boes):

$$(6) \quad LL = \sum_{t=1}^T \ln(G_t) - 0.5 \times \sum_{t=1}^T H_t^2,$$

where $G_t = \{h_t/G(\Theta, \mu)(1 + R_t^2)\}^{-1/2}$; $H_t = \{\sinh^{-1}(R_t)/\Theta\} - \mu$; $R_t = \{[\Theta(y_t - \mathbf{x}_t' \mathbf{b})/h_t/G(\Theta, \mu)]^{-1/2} + F(\Theta, \mu)\}$; $t = 1, \dots, T$ refers to the observations; $\sinh^{-1}(R_t) = \ln\{R_t + (1 + R_t^2)^{1/2}\}$ is the inverse hyperbolic sine function; and h_t , $F(\Theta, \mu)$, and $G(\Theta, \mu)$ are as given in equations (1) and (2).

Maximization of (6) with respect to the model's parameters is achieved through numerical optimization algorithms, which are available in most econometric software packages, including the GAUSS-386i program (Aptech Systems, Inc.) used in this study.

These pre-programmed procedures only require a few standard command lines and the log-likelihood function. These procedures achieve convergence when the gradients of all parameter estimates are less than some arbitrarily small amount. In addition to parameter estimates, they provide standard errors based on a numerical estimate of the Hessian matrix of this function. In most applications, the optimization algorithm converges quickly and properly, as long as the starting values are reasonable. The intercept, slope, and GARCH process parameter estimates from an N-GARCH model are excellent starting values, while it is best to start with $0.1 < \Theta < 0.5$ and $\mu = 0$.

Although there were no convergence problems in any of the applications discussed below, these could arise when working with small samples ($T < 40$) or when there is a highly insignificant parameter. Then, the log-likelihood function will converge to a maximum but not be able to compute standard error estimates. This problem is attributed to the extreme flatness of the log-likelihood function at the maximum or to excessive differences in the magnitudes of the elements of the numerical estimate of the Hessian matrix that has to be inverted.

As in any other quasi-maximum likelihood estimator, such as the N-, t- or EGB2-GARCH, if the regressors are fixed in relation to the error term, the A-GARCH estimators for the slope parameters will be unbiased. Also, as McDonald and Newey point out, so long as the error term is independent of the regressors, any quasi-maximum likelihood estimator of the mean of the distribution of y_t conditional on \mathbf{x}_t would be consistent. Thus, there is no need to assume that ε_t is a member of the expanded S_U family to guarantee a consistent forecast.

With regard to the flexibility of the proposed A-GARCH model, Johnson, Kotz, and Balakrishnan note that both the Gaussian and the lognormal family of densities are limiting cases of the S_U family, which also provides for a close approximation of the Pearson family of distributions. They demonstrate that the S_U family can accommodate any kurtosis-skewness combination below the lognormal line. Because these results apply to the expanded form of the S_U family underlying the proposed asymmetric-error GARCH, it is clear the A-GARCH allows for any mean and variance, as well as for any combination of right or left skewness and leptokurtosis values which could be exhibited by a continuous unimodal variable. Under zero skewness, the A-GARCH allows for any possible mean-variance-leptokurtosis combination: it can precisely fit the first four central moments of any symmetric "thick"-tailed distribution.

Normal-Error versus t- and A-GARCH Models of U.S. Commodity Prices

The modeling strategy of Engle and Kraft, also used by Bollerslev (1986) to illustrate his GARCH expansion of Engle's ARCH process, is adopted here. In particular, it is assumed that real U.S. commodity prices may be forecasted on the basis of their past behavior, denoted by the following model:

$$(7) \quad \begin{aligned} P_t &= b_0 + b_1 P_{t-1} + b_2 P_{t-2} + b_3 P_{t-3} + b_4 P_{t-4} + \varepsilon_t, \\ \varepsilon_t &\sim NN(0, h_t), \quad h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1}. \end{aligned}$$

All models are estimated using the Newton-Raphson algorithm (cubic step-length calculation method) preprogrammed in GAUSS-386i constrained maximum-likelihood

(CML) module. The model in equation (7) is first estimated for real, quarterly U.S. soybean (1924–2000), sorghum (1933–2000), and wheat (1913–2000) prices received by farmers, under the assumption of error-term normality. These nominal price data were compiled from the U.S. Department of Agriculture/National Agricultural Statistics Service (USDA/NASS) database.² The price data were adjusted for inflation, with year 2000 as a base, using the producer price index for nonprocessed agricultural products from the U.S. Department of Labor, Bureau of Labor Statistics.³

All price series are found to be stationary according to augmented Dickey-Fuller tests. Also following Bollerslev (1986), the standardized GARCH residuals ($\varepsilon_t/h_t^{1/2}$) are tested for independence. In the case of soybeans and wheat, several autocorrelations and partial autocorrelations are statistically significant at the 5% level, and the Box-Pierce Q -statistic rejects the null hypothesis that the first 12 autocorrelation coefficients are all zero, also at the 5% level. In the case of sorghum, only the seventh partial autocorrelation is significant at the 5% level, but the Q -statistic does not reject the joint null hypothesis of error-term independence at the 10% level. Closer examination of the correlograms suggests the autocorrelation patterns in the residuals of the soybean and wheat models are likely due to seasonality.

Therefore, the data are seasonally adjusted following a standard centered multiplicative moving-average procedure. The deseasonalized data are also stationary, according to augmented Dickey-Fuller tests. N-, t - and A-GARCH models are then estimated with the deseasonalized price data. In the case of soybean and wheat prices, the standardized residuals do not show any statistically significant autocorrelations, partial autocorrelations, or Q -statistics, at the 10% level. Deseasonalizing real prices solves the autocorrelation problem. Because deseasonalizing actually creates a residual autocorrelation problem in the case of sorghum, the models are estimated with seasonally adjusted soybean and wheat data, and with the unadjusted sorghum data (table 1).

All GARCH parameters (α_0 , α_1 , and β_1) are statistically different from zero at the 5% level, and all models satisfy the variance stationarity condition ($0 < \alpha_1 + \beta_1 < 1$). Also, following Bollerslev (1986), the D'Agostino skewness test statistic (S), Anscombe and Glynn kurtosis test statistic (K), and the D'Agostino-Pearson omnibus normality test statistic (K^2) (D'Agostino, Belanger, and D'Agostino) are applied to the standardized residuals ($\varepsilon_t/h_t^{1/2}$) in the case of the N-GARCH models because, if the true errors are nonnormal, the residuals from an estimated N-GARCH are expected to be nonnormally distributed.

Alternatively, because, in theory, the error term of the A-GARCH model is related to a normally distributed error through equation (2), if an estimated A-GARCH has appropriately accounted for residual nonnormality, the residuals transformed by the inverse of (2) should be normally distributed. Therefore, to evaluate the former, the S , K , and K^2 tests are applied to the standardized and normalized residuals ($\{\sinh^{-1}([\varepsilon_t/h_t/G(\Theta, \mu)]^{-1/2}] + F(\Theta, \mu)/\Theta\} - \mu$) from the A-GARCH models. Because a Student- t random variable cannot be expressed as an algebraic transformation of a normal, the residuals from the t -GARCH models can be standardized but not normalized, making it difficult to ascertain how well the t -distribution is able to model error-term nonnormality in a particular application.

² The USDA/NASS database can be accessed online at <http://www.usda.gov/nass/pubs/histdata.htm>.

³ The U.S. Department of Labor/Bureau of Labor Statistics price index can be accessed online at <http://www.bls.gov/>.

Table 1. Parameter Estimates for Final N-, t-, and A-GARCH Models: U.S. Soybean, Sorghum, and Wheat Prices

A. SOYBEAN PRICES (1924–2000)						
Parameter	N-GARCH		t-GARCH		A-GARCH	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
α_0	0.012**	0.007	0.022*	0.016	0.017**	0.010
α_1	0.142**	0.039	0.140**	0.074	0.159**	0.060
β_1	0.827**	0.043	0.814**	0.089	0.796**	0.068
b_1	1.148**	0.228	1.102**	0.203	1.209**	0.209
b_2	-0.043	0.033	-0.076**	0.028	-0.049*	0.030
b_3	1.074**	0.061	1.050**	0.058	0.969**	0.057
b_4	-0.421**	0.092	-0.420**	0.083	-0.349**	0.083
b_5	0.268**	0.090	0.282**	0.078	0.241**	0.076
b_6	-0.103*	0.061	-0.082	0.058	-0.050	0.050
θ_0	—	—	7.800**	3.475	—	—
θ_1	—	—	-1.873	1.184	0.323*	0.059
μ_0	—	—	—	—	3.308**	1.822
μ_1	—	—	—	—	-0.934**	0.604
2*MVLLF	65.668		93.618		113.058	
Skewness statistic (S)	4.488		—		-0.792	
Kurtosis statistic (K)	3.290		—		0.392	
Normality statistic (K^2)	30.969		—		0.780	
B. SORGHUM PRICES (1933–2000)						
Parameter	N-GARCH		t-GARCH		A-GARCH	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
α_0	0.013**	0.005	0.006*	0.004	0.006**	0.003
α_1	0.041**	0.019	0.034**	0.018	0.027**	0.015
β_1	0.917**	0.027	0.942**	0.020	0.945**	0.019
b_1	1.594**	0.367	1.512**	0.300	1.589**	0.301
b_2	-0.229**	0.072	-0.212**	0.058	-0.228**	0.058
b_3	0.723**	0.063	0.697**	0.048	0.703**	0.048
b_4	-0.263**	0.077	-0.146**	0.060	-0.200**	0.074
b_5	-0.007	0.075	-0.025	0.055	0.040	0.067
b_6	0.258**	0.061	0.183**	0.048	0.167**	0.045
θ_0	—	—	3.520**	0.809	0.657**	0.120
θ_1	—	—	—	—	—	—
μ_0	—	—	—	—	0.827**	0.367
μ_1	—	—	—	—	—	—
2*MVLLF	1.972		56.653		69.404	
Skewness statistic (S)	8.131		—		0.147	
Kurtosis statistic (K)	6.819		—		0.500	
Normality statistic (K^2)	112.61		—		0.272	

(continued...)

Table 1. Continued

C. SOYBEAN PRICES (1913–2000)						
Parameter	N-GARCH		t-GARCH		A-GARCH	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
α_0	0.007**	0.003	0.006**	0.004	0.006**	0.003
α_1	0.243**	0.071	0.227**	0.083	0.214**	0.079
β_1	0.724**	0.061	0.750**	0.073	0.760**	0.071
b_1	0.649**	0.165	0.561**	0.152	0.557**	0.151
b_2	-0.093**	0.027	-0.082**	0.025	-0.081**	0.025
b_3	1.096**	0.059	1.072**	0.056	1.072**	0.056
b_4	-0.374**	0.087	-0.328**	0.080	-0.330**	0.079
b_5	0.308**	0.088	0.243**	0.080	0.244**	0.079
b_6	-0.147**	0.055	-0.090*	0.055	-0.089*	0.054
Θ_0	—	—	4.924**	1.336	0.642**	0.101
Θ_1	—	—	—	—	—	—
μ_0	—	—	—	—	—	—
μ_1	—	—	—	—	—	—
2 * MVLLF	449.86		471.328		472.591	
Skewness statistic (S)	0.483		—		0.558	
Kurtosis statistic (K)	3.823		—		-0.040	
Normality statistic (K^2)	14.846		—		0.313	

Notes: Single and double asterisks (*) denote statistical significance at the 10% and 5% levels, respectively, based on one-tailed t -tests for the GARCH process parameters, which are supposed to be positive, and two-tailed t -tests for all other parameters. The b_2 , Θ_1 , and μ_1 parameter and standard error estimates have been multiplied by 100. MVLLF indicates the maximum value reached by the concentrated log-likelihood function.

Under the null hypothesis of normality, S and K are distributed approximately normal, and K^2 follows a $\chi_{(2)}^2$ distribution. Note, in the N-GARCH models, the null hypothesis of residual normality is rejected at the 1% level of statistical significance in all three cases. The S and K statistics indicate statistically significant error-term skewness and kurtosis in the case of soybean and sorghum prices, but only kurtosis in the case of wheat prices (table 1).

Fully parameterized A-GARCH models (available from the authors on request) were initially estimated specifying the kurtosis and the skewness parameters (Θ and μ) as linear functions of time ($\Theta_t = \Theta_0 + \Theta_1 t$ and $\mu_t = \mu_0 + \mu_1 t$). Single-parameter likelihood-ratio (LR) ($\chi_{(1)}^2$) tests and the usual t -tests are conducted separately to determine the statistical significance of Θ_0 , Θ_1 , μ_0 , and μ_1 . As expected, the LR and the t -tests yield similar results. The final A-GARCH models presented in table 1 were obtained by setting any of these four parameters (Θ_0 , Θ_1 , μ_0 , or μ_1) equal to zero if it did not result in statistical significance at the 10% level. Consequently, the final models for soybean, sorghum, and wheat prices only include Θ_1 , μ_0 , and μ_1 ; Θ_0 and μ_0 ; and Θ_0 , respectively.

In the case of soybean prices (panel A of table 1), $\Theta_0 = 0$ implies the error-term distribution was nearly normal during the 1920s. The positive estimate for Θ_1 combined with a negative estimate for μ_1 in this case results in the kurtosis coefficient increasing through time at a faster rate than the skewness coefficient. Specifically, the skewness and kurtosis coefficients, calculated using the formulas in equations (4) and (5), are both

below 0.5 until 1943. By 1960, skewness reaches 1.2 and kurtosis 6.26. The estimated year 2000 error-term distribution is considerably skewed (3.34) and highly kurtotic (53.27).

In the final A-GARCH sorghum price model (panel B, table 1), $\Theta_1 = 0$ and $\mu_1 = 0$, but Θ_0 and μ_0 are statistically different from zero. The positive estimate for μ_0 implies a right-skewed and leptokurtic error distribution with constant skewness and kurtosis coefficients of 1.46 and 9.84, respectively. Note that some of these skewness-kurtosis combinations are not allowed by the EGB2 family. The EGB2-GARCH would not be as theoretically suitable to model soybean and sorghum prices. In the final A-GARCH wheat model (panel C, table 1), only Θ_0 is statistically different from zero, implying an error-term distribution with a constant kurtosis coefficient of 3.78.

In short, two of the three estimated A-GARCH models exhibit error-term distributions that are substantially right-skewed, suggesting relatively high positive forecasting errors are more likely than negative errors of the same high magnitude. Right-skewed forecasting errors are expected because agricultural price series, even when expressed in real terms, usually exhibit more pronounced peaks on the up side than on the down side. In all three cases, likelihood-ratio tests ($\chi^2_{[3]} = 113.06 - 65.67 = 47.39$; $\chi^2_{[2]} = 69.40 - 1.97 = 67.43$; and $\chi^2_{[1]} = 472.59 - 449.86 = 22.73$) rejected $H_0: \Theta_1 = \mu_0 = \mu_1 = 0$, $H_0: \Theta_0 = \mu_0 = 0$, and $H_0: \Theta_0 = 0$ at the 1% level (table 1), implying the final A-GARCH models are significantly more likely to have generated the observed soybean, sorghum, and wheat price data than final N-GARCH models.

Also note that the values taken by the S , K , and K^2 statistics for the standardized and normalized residuals from the final A-GARCH models are all less than one. These findings confirm the "normalizing" transformation implied by the expanded S_U family was successful in removing the nonnormality from the GARCH residuals in all cases. In other words, the expanded S_U family provides a good approximation to the nonnormal distribution of those residuals.

Initial t-GARCH models were also estimated assuming that the degrees-of-freedom parameter (Θ), which determines the degree of leptokurtosis of the error-term distribution, was a function of time ($\Theta_t = \Theta_0 + \Theta_1 t$). As before, the final t-GARCH models were obtained by setting to zero either of the two parameters (Θ_0 or Θ_1) not statistically significant at the 10% level (table 1). The parameter Θ_1 is only significant in the case of the soybeans model, resulting in the degrees-of-freedom parameter declining from 7.8 in 1924 to 2.1 in the 2000. As in the A-GARCH model of soybean prices, the kurtosis coefficient is markedly increasing through time. Also as in the A-GARCH models, the error-term kurtosis in the t-GARCH models of sorghum and wheat prices has remained constant through time.

The t-GARCH approaches an N-GARCH as the degrees-of-freedom parameter (Θ) approaches infinity, i.e., at the boundary of the admissible parameter space. This relationship implies the distribution of a likelihood-ratio test statistic in which the N-GARCH is viewed as a restricted t-GARCH would be more concentrated toward zero than the probability distribution of a $\chi^2_{[1]}$ variable (Bollerslev 1987). By comparing the likelihood-ratio test statistic with a $\chi^2_{[1]}$ (or a $\chi^2_{[2]}$, when $\Theta = \Theta_t = \Theta_0 + \Theta_1 t$), it follows that rejection of H_0 : N-GARCH in favor of H_A : t-GARCH would be a safe conclusion. Such likelihood-ratio tests allow for rejection of H_0 : N-GARCH in favor of H_A : t-GARCH in all three commodity price models ($\chi^2_{[2]} = 93.62 - 65.67 = 27.95$; $\chi^2_{[1]} = 56.65 - 1.97 = 54.68$; and $\chi^2_{[1]} = 471.33 - 449.86 = 21.47$).

With regard to forecasting performance, because of the partially adaptive nature of the A-GARCH models (McDonald and White), the standard error estimates for the intercept and “slope” parameters ($b_1, b_2, b_3, b_4, b_5,$ and b_6) in these models are always (and sometimes substantially) lower than in the N-GARCH models. Also note the standard error estimates must be evaluated in light of the previously discussed evidence about the unreliability of quasi-maximum-likelihood estimation results when working with finite samples and a misspecified error-term distribution. In particular, symmetric-error GARCH models tend to underestimate the true standard errors of the intercept and slope parameter estimators when the true underlying error-term distribution is asymmetric.

Given the A-GARCH is based on the most flexible error-term distribution, and that statistically significant skewness is detected in the cases of soybean and sorghum prices, the A-GARCH standard error estimates should be considered more reliable than those of the N- and t-GARCH models. The latter models likely underestimate the true standard errors. In the absence of significant skewness, as in the case of wheat price, the t- and A-GARCH models yield nearly identical results.

Because of the differences in the intercept and slope parameter estimates, the point forecasts from the N-, t-, and A-GARCH models differ noticeably from one another. Specifically, the averages of the absolute differences between the within-sample forecasts from the N- versus the A-GARCH are 5.4¢, 6.8¢, and 1.8¢/bushel in the case of the soybean, sorghum, and wheat models, respectively, while the maximum differences are 35¢, 46¢, and 10¢/bushel, respectively. Interestingly, larger average absolute differences are found between the N- and A-GARCH models in which the error terms exhibit more pronounced nonnormality. In spite of the large average absolute differences, the within-sample root mean squared prediction errors (RMSE) and the squares of the correlation coefficients between the observed and the predicted dependent variable values (R^2 s) are almost identical under these three types of models (table 2). Nearly equal RMSEs, however, do not imply that the N-, t-, and A-GARCH model predictions and their associated confidence intervals are equally reliable.

A reliable prediction is characterized by narrow and accurate confidence intervals for the dependent variable occurrences. In nonlinear models, confidence intervals are usually obtained by applying the numerical technique of Krinsky and Robb. Specifically, let $\hat{\mathbf{b}}_i$ ($i = N, t,$ or A) be the $k \times 1$ vector of maximum-likelihood estimators for \mathbf{b}_i , the vector of true population parameters underlying the N-, t-, or A-GARCH models, and $\text{Var}[\hat{\mathbf{b}}_i]$ be the $(k \times k)$ estimated covariance matrix for $\hat{\mathbf{b}}_i$. Then, a draw from the joint probability distribution of $\hat{\mathbf{b}}_i$ is simulated by:

$$(8) \quad \mathbf{z} \times \text{Chol}(\text{Var}[\hat{\mathbf{b}}_i]) + \hat{\mathbf{b}}_i, \quad i = N, t, \text{ or } A,$$

where \mathbf{z} is a $1 \times k$ vector of independent standard normal draws, $\text{Chol}(\cdot)$ denotes the Cholesky decomposition, and $\hat{\mathbf{b}}_i$ the $1 \times k$ vector of ML parameter estimates obtained from $\hat{\mathbf{b}}_i$. Repeated application of equation (8) yields an $m \times k$ matrix (\mathbf{S}_i) of random variables with mean $\hat{\mathbf{b}}_i$ and covariance matrix $\text{Var}[\hat{\mathbf{b}}_i]$. Under the correct model specification, $\hat{\mathbf{b}}_i$ is a consistent estimator for \mathbf{b}_i , and $\text{Var}[\hat{\mathbf{b}}_i]$ is a consistent estimate for the theoretical covariance matrix of $\hat{\mathbf{b}}_i$, implying \mathbf{S}_i is a theoretically correct probabilistic statement about \mathbf{b}_i .

Thus, the boundaries of a $(1 - \alpha)\%$ confidence interval for the expected soybean, sorghum, and wheat prices under the N-, t-, or A-GARCH models at time period t can

Table 2. Confidence Band Statistics for the N-, t, and A-GARCH Models: U.S. Soybean, Sorghum, and Wheat Prices

Description	SOYBEAN PRICE			SORGHUM PRICE			WHEAT PRICE		
	N-GARCH	t-GARCH	A-GARCH	N-GARCH	t-GARCH	A-GARCH	N-GARCH	t-GARCH	A-GARCH
Average width of confidence band ^a	2.026	2.247	1.958	2.234	2.007	2.136	1.175	1.137	1.121
% of Observations:^b									
▸ below 80% LB	6.3	6.6	8.2	4.9	7.9	8.6	7.2	9.8	8.6
▸ above 80% UB	9.9	13.5	11.5	7.9	13.9	9.7	8.1	11.2	11.0
▸ below 90% LB	2.0	2.6	4.9	1.1	2.6	4.5	4.6	4.9	4.9
▸ above 90% UB	6.6	6.9	6.9	4.9	8.2	4.9	5.5	5.8	5.5
▸ below 95% LB	1.6	1.3	2.6	0.7	0.7	1.5	2.6	2.9	2.9
▸ above 95% UB	4.9	3.9	2.6	3.4	4.1	3.4	3.2	3.2	2.3
▸ below 99% LB	0.7	0.3	0.3	0.0	0.0	0.4	1.7	0.6	0.3
▸ above 99% UB	1.6	0.0	0.0	2.2	0.7	0.4	0.9	0.9	0.9
▸ below LB of all confidence bands	51.4	52.9	89.6	35.7	67.4	90.2	86.2	100.5	91.4
▸ above UB of all confidence bands	129.7	132.2	118.7	97.4	156.6	101.3	95.5	116.1	109.0
RMSE of Price Predictions ^c	0.382	0.385	0.383	0.459	0.466	0.458	0.124	0.124	0.124
R^2 ^d	0.734	0.733	0.733	0.646	0.646	0.647	0.861	0.861	0.861

^aAverage width of confidence band is the average width of the 20 (80% to 99%) confidence bands for the observed prices. The average lower bound, upper bound, and width of each of the 20 bands are calculated as the simple means of the bounds and widths of the $n - 4$ confidence intervals obtained for the $n - 4$ predictions.

^bThe % of observations below LB (above UB) refers to the percentage of observations that were vertically below (above) the lower (upper) bound of the corresponding confidence band. The % of observations below LB (above UB) of all confidence bands refers to the percentage of observations that were vertically below (above) the lower (upper) bound of all 20 (80% to 99%) confidence bands. The percentages are in reference to the theoretical number of observations that should be below or above the specified bound.

^cRMSE denotes root mean squared error.

^dThe R^2 is calculated as the square of the correlation coefficient between the observed and the predicted prices.

be obtained by extracting the m sets of simulated intercept and slope parameter values from \mathbf{S}_N , \mathbf{S}_t , or \mathbf{S}_A to obtain m “predicted” price values for time t , and finding the $(\alpha/2) \times m$ th and the $[(1 - \alpha) + \alpha/2] \times m$ th largest of these m price values.

The next step in computing the confidence intervals for the actual price occurrences requires simulation of m draws of the error term as well. In the case of the N-GARCH, these are obtained by extracting the m simulated values for the GARCH process parameters (α_0 , α_1 , and β_1) from \mathbf{S}_N and using them to obtain m simulated values of h_t according to the GARCH(1,1) variance function [equation (7)], and multiplying their square roots by m independent draws from a standard normal random variable. In the case of the A- and t-GARCH models, the process of simulating the h_t values is similar, except that the m sets of simulated GARCH process parameters are extracted from \mathbf{S}_A and \mathbf{S}_t .

For the A-GARCH, these m simulated h_t values are coupled with the m simulated Θ_0 , μ_0 , Θ_1 , and μ_1 values, also obtained from \mathbf{S}_A , and with m independent standard normal draws. Then, m nonnormal error-term values are simulated by applying equation (2). For the t-GARCH, the square roots of the m simulated h_t values are coupled with the m simulated Θ_0 and Θ_1 values obtained from \mathbf{S}_t , and multiplied by m independent draws from a Student- t distribution with $\Theta_0 + \Theta_1 t$ degrees of freedom. The final step in constructing the boundaries of a $(1 - \alpha)\%$ confidence interval for the actual price observations is to add the m simulated error-term values to the corresponding m simulated price “predictions” and find the $(\alpha/2) \times m$ th and the $[(1 - \alpha) + \alpha/2] \times m$ th largest of the resulting m simulated price realization values.

The process described above is programmed in GAUSS-386i with $m = 10,000$, starting at $t = 1$ assuming $\varepsilon_0 = 0$ and $h_0 = \alpha_0 / (1 - \alpha_1 - \beta_1)$ (the unconditional GARCH process variance), and recursively repeated up to $t = T$, under the N-, t-, and A-GARCH soybean, sorghum, and wheat price models. The resulting boundaries are joined to obtain $(1 - \alpha)\%$ confidence bands for the T within-sample price predictions from each of the estimated models. The process is repeated from $\alpha = 0.20$ to $\alpha = 0.01$ at 0.01 decrements to obtain 80% to 99% confidence bands.

Figure 1 shows the 80% confidence bands for the soybean price occurrences versus the sample data under the N- and A-GARCH models. The difference between these two confidence bands starts to become visually obvious at about 1950, when the estimated error term becomes substantially nonnormal. Under the N-GARCH, the observed prices tend to be closer to the middle of the interval. Only 19 observations trespass the lower bound, while 30 observations exceed the upper bound of the 80% confidence band, versus the theoretically required $T \times 0.20/2 = 304 \times 0.20/2 = 30.4$. Given the pattern of the soybean price observations, the symmetry of the assumed error-term distribution results in a lower bound that is unnecessarily low in order to ensure more of the observed price spikes do not surpass the upper bound.

Under the A-GARCH, the right-skewness in the estimated error-term distribution allows for a substantially higher lower bound, which is closer to the mass of the observations, coupled with an upper bound that is still high enough to avoid a theoretically excessive number of observations surpassing it. Specifically, 25 observations are found under the lower bound and 35 exceed the upper bound of the 80% confidence band, respectively, versus the theoretically required 30.4. The t-GARCH bands are symmetric, like the N-GARCH bands, and follow a similar visual pattern. The 80% band leaves 20 observations under its lower bound and 41 observations above its upper bound. At the same time, the average width of the 80% A-GARCH band is \$1.35/bushel, versus \$1.51/bushel for the N-GARCH and \$1.41/bushel in the t-GARCH.




Figure 1. 80% confidence bands for the soybean price occurrences versus 1950–2000 data under the N- and A-GARCH models




Figure 2. Theoretically required versus actual number of observations below the 80% to 99% confidence bands under the N-, t-, and A-GARCH models of soybean prices

Figure 2 shows the number of observations left below the lower bounds of the 80% to 99% confidence bands from the N-, t-, and A-GARCH models of soybean prices, versus the theoretically required numbers. Both the N- and the t-GARCH bands leave substantially fewer observations under their lower bounds than required, while leaving an excessive number of observations over their upper bounds.

When all 20 (80% to 99%) confidence bands are jointly considered, only 164 (51.4%) out of the theoretically required total observations,

$$\left(T \times \sum_{i=1}^{20} i/100 \right) / 2 = \left(304 \times \sum_{i=1}^{20} i/100 \right) / 2 = 319 \text{ (100\%)},$$

are found below the lower bounds of the N-GARCH, while 414 observations (129.7%) appear above the upper bounds (table 2). In the case of the t-GARCH, 169 observations (or 52.9% of the theoretically required) are found below the lower bounds, and 422 (132.2%) appear above the upper bounds. The assumption of error-term symmetry, which causes these confidence bands to be symmetric about the predictions, is clearly incompatible with the observed soybean price data. The 80% to 99% A-GARCH bands, in contrast, leave a total of 286 (89.6%) observations below and 379 (118.7%) observations above their respective boundaries. Although not perfect, the A-GARCH bands adhere more closely to what is theoretically expected.

The A-GARCH bands, being kurtotic and nonsymmetric about the predictions because of the leptokurtosis and right-skewness of the estimated error-term distribution, better reflect the statistical behavior of soybean prices. At the same time, the average width of the A-GARCH confidence bands is \$1.96/bushel, versus \$2.03/bushel and \$2.25/bushel in the cases of the N- and t-GARCH bands. As illustrated in figure 1, compatibility with the data is improved by lower bounds that are higher than their N- and t-GARCH counterparts, and thus closer to the mass of low price occurrences, combined with upper bounds similar to those of the N-GARCH at high α levels, but become relatively higher at reduced α levels.

As illustrated in figure 3, the A-GARCH confidence bands for the sorghum price occurrences are also more compatible with the observed data than the bands under the N- and t-GARCH models. The 80% to 99% N-GARCH bands leave a total of 100 (35.7%) and 273 (97.4%) out of the theoretically required observations,

$$\left(T \times \sum_{i=1}^{20} i/100 \right) / 2 = \left(267 \times \sum_{i=1}^{20} i/100 \right) / 2 = 280 \text{ (100\%)},$$

below and above their lower and upper bounds, respectively, while the t-GARCH bands leave 189 (67.4%) and 439 (156.6%) observations below and above their respective lower and upper bounds. In contrast, in the case of the A-GARCH, 253 (90.2%) and 284 (101.3%) observations are left below and above the lower and upper bounds, respectively (table 2).

It could be argued that the N-GARCH bands in this case could be narrower, leaving more observations below and above their boundaries. However, although the 80% band is much wider than theoretically desired, the 95% and 99% bands leave 3.4% and 2.2% of the observations above their respective upper bounds (table 2), i.e., they are much too narrow on the up side. The problem, then, is the normal distribution cannot produce bands that are asymmetric and narrower at high α values and wider at the lower α levels. Interestingly, the average width of the 80% to 99% confidence bands is less under the t-GARCH than under the A-GARCH (table 2). This result is of no advantage because the band locations are incompatible with the data, but it corroborates the finding of Wang et al. that when the underlying error term is asymmetric, the t-GARCH tends to overestimate the degree of kurtosis, fitting an excessively peaked distribution.

In the case of the wheat prices, where the error term appears to be symmetrically distributed, the t- and A-GARCH models provide fairly similar results, as expected, while

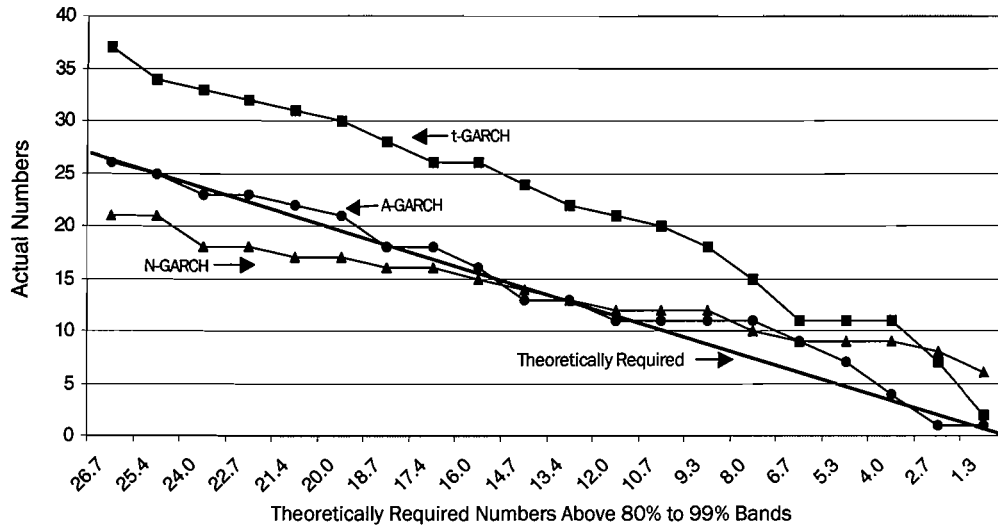


Figure 3. Theoretically required versus actual number of observations above the 80% to 99% confidence bands under the N-, t-, and A-GARCH models of sorghum prices

the N-GARCH still produces a less than adequate statistical representation of the data and relatively wide inefficient confidence bands for the price occurrences (table 2).

The previously discussed inconsistencies between the confidence bands implied by the symmetric-error N- and t-GARCH models and the data they are supposed to represent are important because a reliable, accurate confidence interval is an essential element of a good prediction. The demonstrated incompatibility of the N- and t-GARCH confidence bands with the data is also important because it illustrates the inadequacy of these models when the underlying error is asymmetrically distributed. Although the inadequacy of symmetric-error GARCH models when the true error-term distribution is asymmetric has been established in theory, the good R^2 s and RMSEs for the forecasts often associated with these misspecified models might deceive applied researchers into thinking that they provide reliable predictions of the data.

Concluding Remarks

The asymmetric-error GARCH model proposed and illustrated in this study represents an improved alternative for the forecasting of time-series variables and for producing reliable confidence intervals for these forecasts when the conditional probability distribution of the dependent variable is asymmetric. Researchers using GARCH models should test for error-term nonnormality as illustrated in this analysis. If the error term appears leptokurtic but not skewed, either the t- or the A-GARCH model discussed here should be utilized. If the tests suggest both positive kurtosis and right- or left-skewness, either the EGB2- or the A-GARCH could be suitable, subject to the previously addressed limitations in the levels of skewness and kurtosis allowed by the EGB2 family of distributions. Unfortunately, no theoretically suitable alternative to model the less common negative (platy) kurtosis is available.

[Received August 2002; final revision received December 2002.]

References

- Aptech Systems, Inc. *GAUSS-386i Mathematical and Statistical System*. Maple Valley WA: Aptech Systems, Inc., 1999.
- Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." *J. Econometrics* 31(1986): 307–27.
- . "A Conditionally Heteroskedastic Time-Series Model for Speculative Prices and Rates of Return." *Rev. Econ. and Statis.* 69(August 1987):542–47.
- D'Agostino, R. B., A. Belanger, and R. B. D'Agostino, Jr. "A Suggestion for Using Powerful and Informative Tests of Normality." *Amer. Statistician* 44,4(November 1990):316–21.
- Deb, P. "Finite Sample Properties of Maximum Likelihood and Quasi-Maximum Likelihood Estimators of EGARCH Models." *Econometric Rev.* 15(1996):51–68.
- Engle, R. F. (1982). "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation." *Econometrica* 50(1982):987–1008.
- Engle, R. F., and D. Kraft. "Multiperiod Forecast Error Variances of Inflation Estimated from ARCH Models." In *Applied Time-Series Analysis of Economic Data*, ed., A. Zellner, pp. 93–302. Washington DC: U.S. Bureau of the Census, 1993.
- Johnson, N. L., S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*. New York: John Wiley and Sons, 1994.
- Krinsky, I., and A. L. Robb. "On Approximating the Statistical Properties of Elasticities." *Rev. Econ. and Statis.* 68(1986):715–19.
- Lee, S., and B. E. Hansen. "Asymptotic Theory for the GARCH(1,1) Quasi-Maximum Likelihood Estimator." *Econometric Theory* 10(1994):29–52.
- McDonald, J. B., and W. K. Newey. "Partially Adaptive Estimation of Regression Models Via the Generalized *t*-Distribution." *Econometric Theory* 4(1988):428–57.
- McDonald, J. B., and S. B. White. "A Comparison of Some Robust, Adaptive, and Partially Adaptive Estimators of Regression Models." *Econometric Rev.* 12,1(1993):103–24.
- Mood, A. M., F. A. Graybill, and D. C. Boes. *Introduction to the Theory of Statistics*. New York: McGraw-Hill, 1974.
- Moss, C. B. "The Cost Price Squeeze in Agriculture: An Application of Cointegration." *Rev. Agr. Econ.* 14,1(1992):209–17.
- Moss, C. B., J. S. Shonkwiler, and S. A. Ford. "A Risk Endogenous Model of Aggregate Agricultural Debt." *Agr. Fin. Rev.* 50(1990):73–79.
- Pagan, A. R., and H. Sabau. "On the Inconsistency of the MLE in Certain Heteroskedastic Regression Models." Mimeo, University of Rochester, 1987.
- Ramírez, O. A., and J. S. Shonkwiler. "Autoregressive Conditional Heteroskedasticity Under Error-Term Non-Normality." CASNR Manu. No. 1-513, Texas Tech University, Lubbock, May 2001.
- U.S. Department of Agriculture, National Agricultural Statistics Service. NASS historical data (soybean, sorghum, and wheat historical prices). Online. Available at <http://www.usda.gov/nass/pubs/histdata.htm>. [Retrieved January 2001].
- U.S. Department of Labor, Bureau of Labor Statistics. Producer price index for non-processed agricultural products. Online. Available at <http://www.bls.gov/>. [Retrieved January 2001].
- Wang, K. L., C. Fawson, C. B. Barrett, and J. B. McDonald. "A Flexible Parametric GARCH Model with an Application to Exchange Rates." *J. Appl. Econometrics* 16(2002):521–36.
- Yang, S.-R., and B. W. Brorsen. "Nonlinear Dynamics of Daily Cash Prices." *Amer. J. Agr. Econ.* 74,3(August 1992):706–15.