Analysing Trends and Volatility in Atmospheric Carbon Dioxide Concentration Levels

Felix Chan

Michael McAleer

Analysing Trends and Volatility in Atmospheric Carbon Dioxide Concentration Levels

Felix Chan[a] and Michael McAleer[b]

[a]*School of Economics and Commerce, University of Western Australia (Felix.Chan@uwa.edu.au)*
[b]*School of Economics and Commerce, University of Western Australia*

Abstract: Atmospheric carbon dioxide concentration (ACDC) is a crucial variable for many environmental simulation models, and is regarded as an important factor for predicting temperature and climate changes. However, the conditional variance of ACDC levels has not previously been examined. This paper analyses the trends and volatility in ACDC levels using monthly data from January 1965 to December 2002. The data are a subset of the well known Mauna Loa atmosphere carbon dioxide record obtained through the Carbon Dioxide Information Analysis Center. The conditional variance of ACDC levels is modelled using the generalised autoregressive conditional heteroscedasticity (GARCH) model and its asymmetric variations, namely the GJR and EGARCH models. These models are shown to be able to capture the dynamics in the conditional variance in ACDC levels and to improve the out-of-sample forecast accuracy of ACDC.

Keywords: Atmospheric Carbon Dioxide Concentration, Conditional Volatility, Forecasting, GARCH, GJR, EGARCH.

1. Introduction

Atmospheric carbon dioxide concentration (ACDC) is a crucial variable for many environmental simulation models, and is regarded as an important factor for predicting temperature and climate changes (Glaser (2000)). Many studies in environmental modelling have focused on the application of ACDC as an indictor of the status of the environment (see, for example, Phillips et al. (1998)), while other studies have been interested in the impacts of rising ACDC on the ecological system (see, for example, Jones et al. (1998)). However, these studies have seldom modelled the level of ACDC directly, while the conditional variance of ACDC has not previously been investigated. Although there are mathematical models that are designed to estimate the level of ACDC based on Carbon Dioxide (CO2) emissions from the environment (Phillips et al. (1998)), these simulation models are often complicated and computationally intensive. Moreover, they do not generally provide a simple description of the dynamics in the level of ACDC, and it is difficult to evaluate their forecast performance.

This paper investigates the trends and volatility in ACDC levels using the well known Mauno Lao data set. There are two motivations for modelling the conditional variance of ACDC. First, modelling the conditional variance of ACDC would allow a more accurate confidence interval to be constructed for the one-period ahead forecast. Consider the general regression model given by

$$y_t = E(y_t \mid x_t) + \varepsilon_t,$$

for which the variance of the forecast error, $(\hat{y}_T - y_T)$, is given by

$$\mathrm{Var}(\hat{y}_T - y_T \mid x_0) = \sigma^2\left(1 + \frac{1}{T} + \frac{(x_0 - \bar{x})^2}{\sum_{t=1}^{T}(x_t - \bar{x})^2}\right),$$

where the variance of the innovation, σ^2, is typically assumed to be constant. However, if σ^2 is time varying, the forecast variance can be reduced by accommodating the conditional variance of the time series to permit a more accurate confidence interval to be constructed for the one-period ahead forecast.

The second motivation for modelling the conditional variance of ACDC is related to the pricing of carbon dioxide emission quotas. In financial markets, the risk associated with a stock return is typically measured by its (possibly time-varying) volatility. Therefore, the volatility of ACDC should be an important indicator of the risk in selling or buying emission rights, and would also be an important factor in determining the market value of such quotas. Further details of emissions trading can be found at http://www.ieta.org.

Modelling the conditional variance, or volatility, of a time series has been a popular topic in the financial econometrics literature. Three of the most popular models to capture the time-varying volatility in financial time series are the Generalised Autoregressive Conditional Heteroscedasticity (GARCH) model of

Engle (1982) and Bollerslev (1986), the Glosten, Jagannathan and Runkle (1992) GJR model, and Nelson's (1991) Exponential GARCH (EGARCH) model. This paper examines the dynamics of the conditional variance in the level of ACDC using the GARCH, GJR and EGARCH models. The forecast performance of each model will also be investigated, and the standard errors of the one-day ahead forecasts arising from each model compared.

The plan of the paper is as follows. Section 2 describes the data used. The structural and statistical properties of the three conditional variance models, namely GARCH, GJR and EGARCH, are given in Section 3. The empirical results are presented in Section 4, and Section 5 contains some concluding remarks.

2. Data

The level of ACDC has been closely monitored and documented for over 30 years. The data used in this paper are a subset of the famous Mauna Loa monthly data set, which can be downloaded from http://cdiac.esd.ornl.gov/trends/co2/sio-mlo.htm. The scientific details regarding the measurement of the ACDC level can be found in Keeling, Bacastow and Whorf (1982). Due to missing observations in 1958 and 1964, only the data from January 1965 to December 2002 are used in this paper, giving a total of 456 observations.

Figure 1 contains the plots of ACDC levels from January 1965 to December 2002. The data exhibit cyclical patterns around a time trend. Furthermore, the autocorrelation function of ACDC suggests that it is highly correlated with its past and is highly persistent, as shown in Table 1. The high first-order autocorrelation coefficient might suggest that the series are non-stationary, but the Phillips-Perron (1988) (PP) test for non-stationarity shows that the ACDC level is trend stationary. Using the EViews 4 econometric software package with a wide range of lags, the choice of the truncated lag order did not seem to affect the test results. The motivation for using the PP test over the conventional Augmented Dickey-Fuller (ADF) test is to accommodate the possible presence of ARCH/GARCH errors. While the ADF test accommodates serial correlation by specifying explicitly the structure of serial correlation in the errors, the PP test does not assume the specific type of serial correlation or heteroscedasticity in the disturbances, and can have higher power than the ADF test under a wide range of circumstances.

The sample volatility, v_t, of a time series, y_t, with a non-constant conditional mean is typically calculated as follows:

$$v_t = (y_t - E(y_t \mid \Im_{t-1}))^2 = \varepsilon_t^2, \qquad (1)$$

where $\Im_t$ denotes the information set available to time t. Since the level of ACDC exhibited cyclical patterns, a time trend, and strong autocorrelation, it is reasonable to specify the conditional mean to be

$$E(y_t \mid \Im_{t-1}) = \phi_0 y_{t-1} + \phi_1 t + \theta' D_t \qquad (2)$$

where $\theta = (\theta_1, \theta_2, ..., \theta_{12})'$ and $D_t = (D_1, D_2, ..., D_{12})'$ is the vector of seasonal dummy variables, such that $D_i = 1$ in month i, otherwise $D_i = 0$, $\forall i = 1,...,12$. The plot of the volatility of ACDC can be found in Figure 2.

Table 1: Autocorrelation of the ACDC level.

Lag	Autocorrelation
1	0.991
2	0.978
3	0.964
4	0.952
5	0.942
6	0.934
7	0.929
8	0.926
9	0.925
10	0.926
11	0.926
12	0.922

Figure 1. Atmospheric Carbon Dioxide Concentration, January 1965 – December 2002

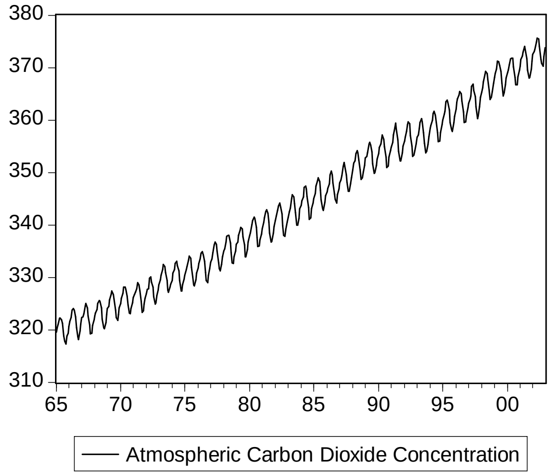

The descriptive statistics of the level, y_t, the estimated residuals from (1), ε_t, and the volatility, v_t, of ACDC are given in Table 2.

As shown in Figure 1 and Table 2, the level of ACDC grew steadily over the last 35 years. The descriptive statistics of the estimated residuals, as given in equations (1) and (2), indicate that the error term, ε_t, is normally distributed. In fact, the Lagrange multiplier test for normality, LM(N), is 1.446 with a p-value 0.485, suggesting that normality cannot be rejected. The p-values of both the F and LM test statistics for the null hypothesis of no ARCH effects with one lag are 0.001, suggesting that the null hypothesis can be rejected at the 1% level of significance. Therefore, there is considerable evidence to suggest that the conditional variance of ACDC is not constant over time, so that conditional volatility models would seem to be an appropriate choice for capturing the time-varying volatility in the level of ACDC.

Figure 2. Volatility of Atmospheric Carbon Dioxide Concentration,
January 1965 – December 2002

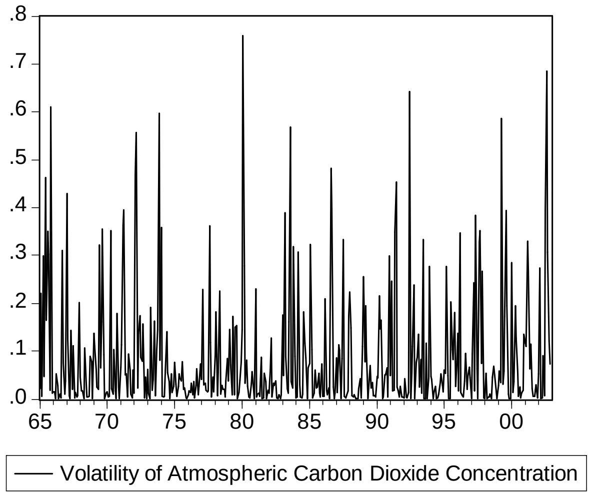

Table 2. Descriptive statistics of the level, estimated residuals and volatility of ACDC

Statistics	y_t	ε_t	v_t
Mean	344.5	0.000	0.085
Median	343.5	-0.005	0.033
Maximum	375.6	0.801	0.76
Minimum	317.3	-0.872	0
SD	16.00	0.292	0.125
Skewness	0.147	0.117	2.353
Kurtosis	1.809	3.148	9.054

3. Models Specifications

The primary empirical purpose of the paper is to model the volatility in the level of ACDC. This approach is based on Engle's (1982) idea of capturing time-varying volatility (or uncertainty) using the autoregressive conditional heteroskedasticity (ARCH) model, and subsequent developments forming the ARCH family of models (see, for example, the recent survey by Li, Ling and McAleer (2002)). Of these models, the most popular has been the symmetric generalised ARCH (GARCH) model of Bollerslev (1986) and the asymmetric Glosten, Jagannathan and Runkle (1992) (GJR) model, especially for the analysis of financial data. A number of further theoretical developments has been suggested by Wong and Li (1997) and Ling and McAleer (2002a, 2002b, 2003).

Consider a GARCH(p,q) model for the level of ACDC, y_t:

$$y_t = E(y_t \mid \Im_{t-1}) + \varepsilon_t, \qquad (3)$$

where $\Im_t$ denotes the information set available to time *t*, and the shocks (or variations in the level of ACDC) are given by

$$\begin{aligned} &\varepsilon_t = \eta_t \sqrt{h_t}, \quad \eta_t \sim iid(0,1) \\ &h_t = \omega + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{q} \beta_i h_{t-i}, \end{aligned} \qquad (4)$$

and $\omega > 0, \alpha_i \geq 0$ (i = 1,…,p) and $\beta_i \geq 0$ (i = 1,…,q) are sufficient conditions to ensure that the conditional variance $h_t > 0$. The ARCH (or α) effect captures the short run persistence of shocks, while the GARCH (or β) effect captures the contribution of shocks to long run persistence (namely, $\alpha + \beta$ for *p=q=1*). Using results from Ling and Li (1997) and Ling and McAleer (2002a, 2002b) (see also Bollerslev (1986) and Nelson (1990)), the necessary and sufficient condition for the existence of the second moment of ε_t, or $E(\varepsilon_t^2) < \infty$, for GARCH(1,1) is $\alpha + \beta < 1$.

Equation (2) assumes that a positive shock ($\varepsilon_t > 0$) has the same impact on the conditional variance, h_t, as a negative shock ($\varepsilon_t < 0$), but this assumption is often violated in practice. In order to accommodate the possible differential impact on the conditional variance between positive and negative shocks, Glosten, Jagannathan and Runkle (1992) proposed the following asymmetric GJR specification for h_t:

$$h_t = \omega + \sum_{i=1}^{p} (\alpha_i + \gamma_i I(\varepsilon_{t-i}))\varepsilon_{t-i}^2 + \sum_{i=1}^{q} \beta h_{t-i}, \qquad (5)$$

where $I(\varepsilon_t)$ is an indicator function such that

$$I(\varepsilon_t) = \begin{cases} 0, & \varepsilon_t \geq 0 \\ 1, & \varepsilon_t < 0. \end{cases}$$

When $\beta = 0$, GJR(1,1) is called the asymmetric ARCH(1), or AARCH(1), model. Furthermore, for GJR(1,1), $\omega > 0$, $\alpha + \gamma > 0$ and $\beta > 0$ are sufficient conditions to ensure that the conditional variance $h_t > 0$. The short run persistence of positive (negative) shocks is given by α ($\alpha + \gamma$). Under the assumption that the conditional shocks, η_t, follow a symmetric distribution, the average short run persistence is $\alpha + \frac{\gamma}{2}$, and the contribution of shocks to average long run persistence is $\alpha + \frac{\gamma}{2} + \beta$. Ling and McAleer (2002a) showed that the necessary and sufficient condition for $E(\varepsilon_t^2) < \infty$ is $\alpha + \frac{\gamma}{2} + \beta < 1$.

The parameters in equations (1), (2) and (3) are typically estimated by the maximum likelihood method to obtain Quasi-Maximum Likelihood Estimators (QMLE) in the absence of normality of η_t. The conditional log-likelihood function is given as follows:

$$\sum_t l_t = -\frac{1}{2}\sum_t \left(\log h_t + \frac{\varepsilon_t^2}{h_t} \right).$$

Ling and McAleer (2003) showed that the QMLE for GARCH(p,q) is consistent if the second moment is finite, that is, $E(\varepsilon_t^2) < \infty$. Furthermore, Jeantheau (1998) showed that, when $\beta \neq 0$, the following log-moment condition

$$E(\log(\alpha\eta_t^2 + \beta)) < 0 \qquad (6)$$

is sufficient for the QMLE to be consistent for GARCH(1,1), while Boussama (2000) showed that the QMLE is asymptotically normal for GARCH(1,1) under the same condition. It is important to note that (6) is a weaker condition than the second moment condition, namely $\alpha + \beta < 1$. However, the log-moment condition is more difficult to compute in practice as it is the expected value of a function of an unknown random variable and unknown parameters.

McAleer, Chan and Marinova (2002) established the log-moment condition for GJR(1,1) when $\beta \neq 0$, namely

$$E(\log((\alpha + \gamma I(\eta_t))\eta_t^2 + \beta)) < 0, \qquad (7)$$

and showed that it is sufficient for the consistency and asymptotic normality of the QMLE for GJR(1,1). Furthermore, using Jensen's inequality, they showed that the second moment condition, namely $\alpha + \frac{\gamma}{2} + \beta < 1$, is also a sufficient condition for consistency and asymptotic normality of the QMLE for GJR(1,1). Therefore, the structural and statistical properties of both GARCH(1,1) and GJR(1,1) have been established (see Chan, Hoti and McAleer (2002) for the structural and statistical properties of the multivariate GJR(p,q) model).

An alternative model to capture asymmetric behaviour in the conditional variance is the Exponential GARCH (EGARCH(1,1)) model of Nelson (1991), namely:

$$\log h_t = \omega + \alpha \,|\, \eta_{t-1} \,|\, + \gamma\eta_{t-1} + \beta \log h_{t-1}, \; |\, \beta \,| < 1. \qquad (8)$$

When $\beta = 0$, EGARCH(1,1) becomes EARCH(1). There are some distinct differences between EGARCH and the previous two GARCH models, as follows: (i) EGARCH is a model of the logarithm of the conditional variance, which implies that no restrictions on the parameters are required to ensure $h_t > 0$; (ii) Nelson (1991) showed that $|\beta| < 1$ ensures stationarity and ergodicity for EGARCH(1,1); (iii) Shephard (1996) observed that $|\beta| < 1$ is likely to be a sufficient condition for consistency of QMLE for EGARCH(1,1); (iv) as the conditional (or standardized) shocks appear in equation (4), McAleer et al. (2002) observed that is likely $|\beta| < 1$ is a sufficient condition for the existence of all moments, and hence also sufficient for asymptotic normality of the QMLE of EGARCH(1,1).

Furthermore, EGARCH captures asymmetries differently from GJR. The parameters α and γ in EGARCH(1,1) represent the magnitude (or size) and sign effects of the conditional (or standardized) shocks, respectively, on the conditional variance. However, α and $\alpha + \gamma$ represent the effects of positive and negative shocks, respectively, on the conditional variance in GJR(1,1).

As GARCH is nested within GJR, a standard asymptotic test of $H_0 : \gamma = 0$ can be used to test the two models against each other. However, as EGARCH is non-nested with regard to both GARCH and GJR, the non-nested models are not directly comparable. Ling and McAleer (2000) proposed a simple non-nested test to discriminate between GARCH and EGARCH. Denoting GARCH as the null hypothesis and EGARCH as the alternative, the optimal test statistic for $H_{GARCH} : \delta = 0$ is given by:

$$h_t = w + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1} + \delta\hat{g}_t \qquad (9)$$

where $\hat{g}_t$ is the generated one-period ahead conditional variance of EGARCH. For the reverse case, that is, denoting EGARCH as the null hypothesis and GARCH as the alternative, the optimal test statistic for $H_{EGARCH} : \delta = 0$ is given by:

$$\log g_t = w + \alpha|\eta_{t-1}| + \gamma\eta_{t-1} + \beta\log g_{t-1} + \delta\log\hat{h}_t \qquad (10)$$

where $\hat{h}_t$ is the generated one-period ahead conditional variance of GARCH. Ling and McAleer (2000) showed that the QMLE of δ in both (9) and (10) are asymptotically normal under the respective null hypotheses, and consistent under the respective alternative hypotheses. They also derived the power functions of both test statistics under the respective hypotheses. A similar non-nested test for testing GJR and EGARCH against each other was derived in McAleer et al. (2002).

4. Empirical Results

4.1 Full Sample Estimates

The parameter estimates and their Bollerslev-Wooldridge (1992) robust t-ratios of the ARCH(1), AARCH(1), EARCH(1), GARCH(1,1), GJR(1,1) and EGARCH(1,1) models, with conditional means as defined in (2), are available on request. These estimates were obtained from EViews 4.0 using the BHHH algorithm.

The parameter estimates in the conditional mean are not particularly sensitive to the specification of the conditional variance equation, which is due to the block-diagonality of the Hessian matrix of the log-likelihood function. Moreover, the log-moment conditions are satisfied for both GARCH(1,1) and GJR(1,1), and the second moment conditions are satisfied for the ARCH(1) and AARCH(1) models, thereby indicating that the QMLE are consistent and asymptotically normal for each of these models. Furthermore, $\hat{\beta}$ < 1 for EGARCH, and it is not significant in the other two cases, suggesting the absence of long run persistence. Interestingly, γ is not significant in either AARCH(1) or GJR(1,1), but it is significant in both EARCH(1) and GARCH(1,1), indicating the presence of asymmetric behaviour. Based on the significance of the parameter estimates, ARCH(1) and EARCH(1) are empirically superior to the other four specifications. Subsequently, non-nested tests based on (9) and (10), with β = 0 in both equations, are conducted in order to choose between the two remaining adequate specifications. The test statistics are given in Table 3.

As shown in Table 3, the test statistic rejects ARCH(1) in favour of EARCH(1) at the 10% level of significance, but does not reject EARCH(1) in favour of ARCH(1) at any reasonable significance level.

Table 3. Non-nested Tests between ARCH(1) and EARCH(1)

Null H_0 Alternative H_1	H_0 : ARCH(1) H_1 : EARCH(1)	H_0 : EARCH(1) H_1 : ARCH(1)
Test Statistics	1.764	0.180

4.2 Forecasting

This section examines the forecast performance and forecast variance for the model as defined in equation (2), with three different conditional variance specifications, namely the constant conditional variance, ARCH(1) and EARCH(1). The three models are re-estimated using the sub-sample from January 1965 to December 2001, and the out-of-sample one-period ahead forecast of ACDC is calculated for January 2002 to December 2002. Three standard forecast criteria, namely root mean square error (RMSE), mean absolute error (MAE) and mean absolute percentage error (MAPE), for each model are reported in Table 4.

Table 4. Forecast Performance of Three Conditional Variance Specifications

Performance criteria	Constant conditional variance	ARCH(1)	EARCH(1)
RMSE	0.701	0.680	0.458
MAE	0.517	0.504	0.377
MAPE	0.138	0.135	0.101

As shown in Table 4, EARCH(1) has the best forecast performance based on the three forecast criteria. More importantly, allowing dynamic conditional variances improves the accuracy of the parameter estimates and also the out-of-sample forecasts. Table 5 gives the standard errors of the one-period ahead forecasts for each month from the three models.

Table 5. Standard Errors of the One-Period Ahead Forecasts for Three Volatility Models

Month	Constant	ARCH(1)	EARCH(1)
January	0.298	0.299	0.280
February	0.413	0.412	0.397
March	0.496	0.493	0.480
April	0.561	0.558	0.546
May	0.615	0.611	0.600
June	0.661	0.657	0.647
July	0.701	0.696	0.687
August	0.735	0.731	0.722
September	0.766	0.761	0.753
October	0.793	0.789	0.781
November	0.817	0.813	0.806
December	0.838	0.836	0.828

Apart from having the best forecast performance, the one-day ahead forecasts produced by EARCH(1) also have the smallest standard errors, as shown in Table 5. This suggests that the one-day ahead forecast produced by EARCH(1) will have the smallest confident intervals, indicating EARCH(1) is superior in terms of forecasting accuracy for the levels of ACDC. Moreover, the standard errors of the one-day ahead forecasts produced by ARCH(1) are smaller than those from the constant conditional variance model for eleven of twelve months. These results show that the accuracy in forecasting ACDC levels can be improved substantially by accommodating time-varying conditional variance in modelling ACDC.

5. Concluding Remarks

This paper examined the trends and volatility in the level of ACDC. Six different specifications of the conditional variance, namely ARCH(1), AARCH(1), EARCH(1), GARCH(1,1), GJR(1,1) and EGARCH(1,1), have been estimated and tested against each other. The test statistics suggested that EARCH(1) was superior to the other five specifications, having the best out-of-sample forecast performance in terms of three different forecast criteria, namely root mean square error, mean absolute error and mean absolute percentage error. Moreover, the one-day ahead forecasts produced by EARCH(1) also had the smallest standard errors.

6. References

Bollerslev, T. (1986), Generalised autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 31, 307-327.

Bollerslev, T. and J.M. Wooldridge (1992), Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances, *Econometric Reviews*, 11, 143-173.

Boussama, F. (2000), Asymptotic normality for the quasi-maximum likelihood estimator of a GARCH model, *Comptes Rendus de l'Academie des Sciences*, Serie I, 331, 81-84 (in French).

Chan, F., S. Hoti and M. McAleer (2002), Structure and asymptotic theory for multivariate asymmetric volatility: Empirical evidence for country risk ratings, paper presented to the Australasian Meeting of the Econometric Society, Brisbane, Australia, July 2002.

Engle, R.F. (1982), Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica*, 50, 987-1007.

Glaser, K. (2000), Global warming: Correlation between atmospheric carbon dioxide concentration and the temperature, Technical Report, Roche Colorado Corporation, Boulder, Colorado, USA.

Glosten, L., R. Jagannathan and D. Runkle (1992), On the relation between the expected value and volatility of nominal excess return on stocks, *Journal of Finance*, 46, 1779-1801.

Jeantheau, T. (1998), Strong consistency of estimators for multivariate ARCH models, *Econometric Theory*, 14, 70-86.

Jones, T.H., L.J. Thompson, J.H. Lawton, T,M. Bezemer, R.D. Bardgett, T.M. Blackburn, K.D. Bruce, P.F. Cannon, G.S. Hall, S.E. Hartley, G. Howson, C.G. Jones, C. Kampichler, E. Kandeler and D.A. Ritchie (1998), Impacts of rising atmospheric carbon dioxide on model terrestrial ecosystems, *Science*, 280, 441-443.

Keeling, C.D., R.B. Bacastow and T.P. Whorf (1982) Measurements of the concentration of carbon dioxide at Mauna Loa Observatory, Hawaii, in W.C. Clark (ed.), *Carbon Dioxide Review: 1982*, Oxford University Press, New York.

Li, W.K., S. Ling and M. McAleer (2002), Recent theoretical results for time series models with GARCH errors, *Journal of Economic Surveys*, 16, 245-269. Reprinted in M. McAleer and L. Oxley (eds.), *Contributions to Financial Econometrics: Theoretical and Practical Issues*, Blackwell, Oxford, 2002, pp. 9-33.

Ling, S. and W.K. Li (1997), On fractionally integrated autoregressive moving-average models with conditional heteroskedasticity, *Journal of the American Statistical Association*, 92, 1184-1194.

Ling, S. and M. McAleer, (2000), Testing GARCH versus E-GARCH, in W.-S. Chan, W.K. Li and H. Tong (eds.), *Statistics and Finance: An Interface*, Imperial College Press, London, pp. 226-242.

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

Ling, S. and M. McAleer (2002b), Stationarity and the existence of moments of a family of GARCH processes, *Journal of Econometrics*, 106, 109-117.

Ling, S. and M. McAleer (2003), Asymptotic theory for a vector ARMA-GARCH model, *Econometric Theory*, 19, 278-308.

McAleer, M., F. Chan and D. Marinova (2002), An econometric analysis of asymmetric volatility: theory and application to patents, paper presented to the Australasian Meeting of the Econometric Society, Brisbane, July 2002, to appear in *Journal of Econometrics*.

Nelson, D.B. (1990), Stationarity and persistence in the GARCH(1,1) model, *Econometric Theory*, 6, 318-334.

Nelson, D.B. (1991), Conditional heteroscedasticity in asset returns: a new approach, *Econometrica*, 59, 347-370.

Phillips, O.L., Y. Malhi, N. Higuchi, W. F. Laurance, P. V. Nunez, R.M. Vasquez, S.G. Laurance, L. V. Ferreira, M. Stern, S. Brown and J. Grace (1998), Changes in the carbon balance of tropical forests: Evidence from long-term plots, *Science*, 282, 439-442.

Phillips, P. and P. Perron (1988), Testing for a unit root in time series regression, *Biometrika*, 75(2), 335–346.

Shephard, N. (1996), Statistical aspects of ARCH and stochastic volatility, in O.E. Barndorff-Nielsen, D.R. Cox and D.V. Hinkley (eds.), *Statistical Models in Econometrics, Finance and Other Fields*, Chapman & Hall, London, pp. 1 – 67.

Wong, H. and W.K. Li (1997), On a multivariate conditional heteroscedasticity model, *Biometrika*, 4, 111-123.