SHORT-TERM DYNAMICS IN THE CYPRUS STOCK EXCHANGE

Gregory Koutmos*

Andreas Pericli**

Lenos Trigeorgis***

This Version: December 2004.

** Gerald M. Levin Professor of Finance, Charles F. Dolan School of Business, Fairfield University, CT 06430-5195, USA, Telephone 203-254 4000, Email: gkoutmos@mail.fairfield.edu.

** Chairman and CEO, Euclid Financial Group, Inc., 1737 H Street, NW Suite 100, Washington DC 20006, USA, Telephone 202-822 2139, Email: info@euclidgroup.com.

*** Professor of Finance, Department of Public and Business Administration, School of Economics and Management, P.O. Box 20537, CY-1678 Nicosia, Cyprus, Telephone:+357-22892476, Email:lenos@ucy.ac.cy

SHORT-TERM DYNAMICS IN THE CYPRUS STOCK EXCHANGE

ABSTRACT

This paper investigates the short-term dynamics of stock returns in an emerging stock market namely, the Cyprus Stock Exchange (CYSE). Stock returns are modeled as conditionally heteroskedastic processes with time-dependent serial correlation. The conditional variance follows an EGARCH process, while for the conditional mean three nonlinear specifications are tested namely, a) the LeBaron (1992) exponential autoregressive model, b) the positive feedback trading model, due to Sentana and Wadhwani (1992) and finally, c) a model that nests both (a) and (b). There is an inverse relationship between volatility and autocorrelation consistent with the findings from several other stock markets, including the U.S. This pattern could be the manifestation of a certain form of noise trading namely positive feedback trading or, momentum trading strategies. There is little evidence that market declines are followed with higher volatility than market advances, the so-called "leverage effect", that has been observed in almost all developed stock markets. In out of sample forecasts, the nonlinear specifications provide better results in terms of forecasting both first and second moments of the distribution of returns.

Keywords: Cyprus Stock Exchange, positive feedback trading, stock return dynamics, EGARCH, GED, Exponential Autoregression, Forecasting.

SHORT-TERM DYNAMICS IN THE CYPRUS STOCK EXCHANGE

I. INTRODUCTION

A number of studies have examined the stochastic behavior of major national stock markets, uncovering several stylized empirical facts: First, the empirical distribution of stock returns appears to be excessively leptokurtic vis-a-vis the normal distribution (e.g., Mandelbrot 1963, Fama 1965, Nelson 1991, and Booth et al. 2000, 1992, among others). Second, short-term stock returns exhibit volatility clustering. Mandelbrot (1963, p. 418) noted that "large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes". This type of behavior has been modeled quite successfully with ARCH-type models (e.g., Engle 1982, Bollerslev et al. 1994, Koutmos and Knif 2002, Moschini and Myers 2002, Scruggs and Glabanidis 2003 and Zhou 2002, to mention but a few). Third, changes in stock prices tend to be negatively related to changes in volatility (e.g., Black 1976, Christie 1982 and Bekaert and Wu 2000), i.e., negative returns (lower prices) are followed by greater volatility than positive returns (higher prices) of an equal magnitude. Finally, the first-order autocorrelation of stock index returns is negatively related to the level of volatility (e.g., LeBaron 1992, Campbell et al. 1993 and Sentana and Wadhwani 1992).

Most of the work thus far has focused on developed capital markets with few exceptions (see for example, Koutmos 1997). For investors seeking opportunities in emerging markets however, a better understanding of the stock return behavior is of great interest. One such promising market that has received no attention (to our knowledge) is the Cyprus Stock Exchange. Cyprus has a small but dynamic economy and its capital markets have grown substantially in recent years. Moreover, Cyprus has recently become a full member of the European Union. Consequently, a better understanding of its stock market should be of interest to international investors. More specifically, it would be useful to examine whether stylized facts which are present in developed markets, also characterize the behavior of the Cyprus Stock Exchange (CYSE). During the period 1985-2002 the average annual stock return to investors has been approximately 15.7%. More importantly, returns on the CYSE have been found to have low (in some years even negative) correlation with the returns of major developed markets, providing opportunities for substantial risk reduction through diversification to international investors.[1] Thus, understanding the stochastic behavior of the CYSE is vital for both international as well as local investors and may point out interesting similarities and differences between the CYSE and developed stock markets. Thus far, there has been no empirical study that methodically investigates the short-term dynamics of stock returns in the CYSE despite the potential importance of this market to investors.

The main objective of this paper is to investigate whether the stylized facts found in major developed markets are also present in the Cyprus Stock Exchange. Confirmation of such empirical findings would be indirect evidence of market integration. Several models are employed in order to investigate empirically which one explains better the stochastic behavior of the CYSE.

As in other developing markets, the CYSE returns are highly leptokurtic. To accommodate the high kurtosis observed in the data, the Generalized Error Distribution (GED), which allows for deviation from normality (fat tails and peakedness) is used. The GED density function subsumes the normal as well as several other density functions as special cases. The conditional standard deviation is modeled via an Exponential GARCH (EGARCH). An array of diagnostics, including recently proposed variance specification tests by Engle and Ng (1993), are also used to test whether such models can explain the stochastic behavior of the Cyprus stock market returns.

The CYSE return distribution seems to be closer to the double exponential than to the normal. There is a negative relationship between volatility and autocorrelation consistent with the view that some traders follow positive feedback trading or, momentum strategies. There is little evidence that market declines are followed with higher volatility than market advances, the so-called “leverage effect”. In out of sample forecasts, the nonlinear specifications used in this paper appear to provide better results in terms of forecasting both first and second moments of the return distributions.

The paper is organized as follows. The next section provides some institutional background on the CYSE. Section III discusses the data and presents some preliminary findings. Section IV describes the econometric models used in this study. Section V discusses the estimation procedure and the empirical findings. Section VI concludes.

II. A BRIEF DESCRIPTION OF THE CYPRUS STOCK MARKET

During the period under study, transactions in the Cyprus stock market took place primarily through a decentralized network of dealers/brokers in an over-the-counter market. The absence of a formal stock exchange was compensated in part through sponsorship and monitoring of this over-the-counter market by the Cyprus Chamber of Commerce and Industry (CCCI or KEVE). CCCI has been publishing daily quotations for bid and ask prices supplied by individual brokers, which were binding only for a specified minimum block of shares. Regular, centralized, auction-type meetings have been taking place at CCCI´s premises where all the brokers convened to arrive at a single market price for traded securities. (There are no specialists or official market makers.) Typically these market prices set at each centralized meeting served as benchmarks for market price levels until the next such meeting. However, the absence of a continuous, high-volume auction market and of a regulated competitive environment left open the possibility that quoted prices might deviate from the underlying fundamental value for many securities. Rumors, manipulations by certain traders, overreaction by mainly unsophisticated investors, and a herd-mentality, may have exacerbated momentum trading strategies and resulted in positive autocorrelation in stock market prices, especially in the earlier period. Despite substantial inefficiencies in the market in the earlier period, taking advantage of arbitrage opportunities was limited by relatively high transaction costs (commission fees in excess of 1% each way) and limits on short selling. A spot settlement system was used.

Over this period the Cyprus Stock Exchange experienced several significant structural changes. First, there has been a substantial increase in the number of auction-type meetings at CCCI´s premises which constituted a market trading “floor”. This increase in the frequency of centralized, auction-type meetings was a definite sign of market progress. The prices arrived at the closing of each meeting were a much better indication of the underlying supply and demand than the price quotes published by a number of brokerage firms in the press before December 1990 when the number of meetings was small. Cyprus improved preferential relationship with the European Union after 1990 may also have contributed to a potential structural change around this period.

Parallel to the increase in the frequency of auction-type meetings, the number of brokerage firms also increased substantially. The biggest players were investment branches of the major banks and some individual brokers who entered early on. The number of public companies also increased significantly, especially between 1985-1991, with the total number of public companies reaching 39 in 1994 (up from 13 in 1985). Thus a more active market for securities was developing in the early 1990´s. A large increase in the daily price variance was observed in the early 1990´s as compared to the late 1980´s, consistent with a substantial improvement in market efficiency.

A law to provide for the development of the securities market in the Republic of Cyprus, the establishment and operation of a Cyprus Stock Exchange, the operation of a stock exchange Council and other related matters was passed by parliament in 1995. After its formation, the Cyprus Stock Exchange council proceeded to formulate and propose a comprehensive set of regulations to govern the operation of the formal Cyprus Stock Exchange, which became effective in March 1996. The drafting and passage of an appropriate legal/institutional framework (e.g., governing information disclosure, new securities registration, credential requirements for brokers, etc.) was deemed crucial for the proper functioning and development of the capital market in Cyprus and the efficient allocation of economic resources.

III. DATA AND PRELIMINARY FINDINGS

The data used in this study consists of daily stock returns from the CYSE. The sample begins on January 21, 1985, and ends on November 22, 2002 for a total of 4,080 observations. The returns are continuously-compounded and they are based on the closing prices of the Cyprus Stock Exchange Index (CYSE), which is a value-weighted market index.

Table 1 reports several diagnostic test statistics on the distributional properties of the daily return data. These include the standard deviation (σ = 1.3387), measures for skewness (S = 1.8517) and excess kurtosis (K = 39.2598), the Kolmogorov-Smirnov D-statistic testing for normality, and the Ljung-Box statistic testing for autocorrelation and the returns and the squared returns. The Kolmogorov-Smirnov (D) statistic rejects normality as do measures for skewness and excess kurtosis.[2] Non-normality may be caused in part by temporal dependencies in the returns, especially second-moment temporal dependencies. The presence of such dependencies is tested by means of the Ljung-Box (LB) statistic calculated for ten lags and applied to returns (testing for linear or, first moment dependencies) as well as to squared returns (testing for nonlinear or, second moment dependencies). The hypothesis that all autocorrelations up to the 10th lag are jointly zero is soundly rejected for both the returns and the squared returns. A possible reason for the autocorrelation in the returns is nonsynchronous trading (e.g., Fisher 1996, Scholes and Williams 1977). Another reason may be time-varying short-term expected returns (e.g., Conrad and Kaul 1988, and Fama and French 1988). Autocorrelation of the squared returns provides evidence of time-varying second moments and a justification for the subsequent use of ARCH-type specification for the variance.

Even though the LB statistic provides evidence for second-moment time dependencies, it cannot be used to test for the empirically observed phenomenon in developed markets of asymmetric volatility. To investigate whether the return shocks in the CYSE have an asymmetric effect on volatility, we use some diagnostics recently proposed by Engle and Ng (1993). These tests are based on the news impact curve implied by the particular ARCH model used. The premise is that if the volatility process is correctly specified, then the squared standardized residuals should not be predictable on the basis of observed variables. These tests are a) the Sign Bias Test, b) the Negative Size Bias Test, c) the Positive Size Bias Test, and d) the Joint Test. The first test examines the impact of positive and negative innovations on volatility not predicted by the model. The squared residuals are regressed against a constant and a dummy S^{-}_{t} that takes the value of unity if ε_{t-1} is negative and zero otherwise. The test is based on the t-statistic for S^{-}_{t}. The Negative Size Bias Test examines how well the model captures the impact of large and small negative innovations. It is based on the regression of the standardized residuals against a constant and $S^-_t \varepsilon_{t-1}$. The calculated t-statistic for $S^-_t \varepsilon_{t-1}$ is used in this test. The Positive Sign Bias Test examines possible biases associated with large and small positive innovations. Here, the standardized filtered residuals are regressed against a constant and $(1-S^-_t)\varepsilon_{t-1}$. Again, the t - statistic for $(1- S^-_t)\varepsilon_{t-1}$ is used to test for possible biases. Finally, the joint test is an F-test based on a regression that includes all three variables, i.e., S^-_t, $S^-_t \varepsilon_{t-1}$ and $(1-S^-_t)\varepsilon_{t-1}$. The calculated t-statistics as well as the F-statistics based on these regressions are reported in Table 2.

The results document significant Negative Size Bias, Positive Size Bias and a significant joint F-test, suggesting the presence of asymmetries in the conditional variance. Overall, the preliminary evidence supports the inclusion of asymmetric components in the volatility specification in order to model adequately the stock market volatility of the CYSE.

IV. MODELING SHORT-TERM DYNAMICS

This section describes the econometric models used in this study: These models combine the EGARCH volatility process with the following specifications for the conditional mean: a) the exponential autoregressive model due to LeBaron (1992), b) the positive feedback trading model due to Sentana and Wadhwani (1992) and c) a model that nests both (a) and (b). The three models are designed to test several hypotheses about the short-term behavior of CYSE daily returns.

Let r_t be the continuously compounded rate of return. The conditional density of stock returns can be described as,

$$r_t \mid I_{t-1} \sim f(\mu_t,\sigma^2_t,\nu) \quad (1)$$

where, μ_t, σ^2_t and ν are the conditional mean, the conditional variance and a constant scale parameter (indicating deviation from normality) respectively. I_{t-1} represents the information set at time t-1. The return at time t will be equal to its conditional expectation plus the innovation (error), i.e.,

$$r_t = \mu_t + \sigma_t z_t \tag{2}$$

where, $z_t = \varepsilon_t/\sigma_t$ is a zero mean and unit variance variate and ε_t is the residual or innovation at time t, i.e., $\varepsilon_t = r_t - \mu_t$. The conditional variance of ε_t given by

$$\ln(\sigma^2_t) = \alpha_0 + \Sigma^p_{i=1} \alpha_i\big(|z_{t-i}| - E|z_{t-i}| + \delta z_{t-i}\big) + \Sigma^q_{j=1}\gamma_j \ln(\sigma^2_{t-j}) \tag{3}$$

The conditional variance in (3) follows an Exponential Generalized Autoregressive Conditionally Heteroskedastic process of orders p, q (EGARCH[p,q]). In most applications a lag structure of p = q = 1 is adequate. With longer time series, however, it is sometimes necessary to increase the number of lags. The EGARCH was suggested by Nelson (1991) and has several advantages over Bollerslev's (1986) linear GARCH model. First, the natural log formulation ensures positive variances, thus dispensing with the need for parameter restrictions; and second, volatility at time t depends on both the size and the sign of past normalized errors. Equivalently, $\ln(\sigma^2_t)$ is allowed to respond asymmetrically to positive and negative normalized residuals. Depending on the sign of δ in (4), the sign effect may be reinforcing or partially offsetting the size effect.[3] Existence of the unconditional variance requires that $\Sigma^q_{j=1}\gamma_j < 1$. If this condition holds, then the log of the unconditional variance will be given by

$$\ln(\sigma^2) = \alpha_0 \, (1 - \Sigma^{q}_{j=1}\gamma_j)^{-1} \qquad (4)$$

IV. A. Exponential Autoregressive Model

The exponential autoregressive model assumes that stock returns are related to their past in a nonlinear fashion. The particular functional form suggested by LeBaron (1992) is as follows:

$$\mu_t = \beta_0 + \beta_1\sigma^2_t + (\varphi_0 + \varphi_1 \exp\{-\sigma^2_{t-1}\})r_{t-1} \,. \qquad (5)$$

The conditional mean, given by equation (5), is an exponential autoregressive process of order one (EAR[1]) along with an EGARCH-M effect (captured via parameter β_1). The autocorrelation of the returns is an exponential functions of the conditional variance. It can be seen from (5) that autocorrelation is high during calm periods and low during volatile periods. When the conditional variance tends to infinity, autoregression tends to φ_0 and when the conditional variance tends to zero, the first-order autoregression tends to $\varphi_0 + \varphi_1$.

IV. B. Positive Feedback Trading Model

The positive feedback-trading model of Sentana and Wadhwani (1992) assumes that traders consist of two heterogeneous groups namely, risk averse expected utility maximizers and positive feedback traders. The demand for shares by the first group is based on the conditional CAPM and it is given by

$$Q_{1,t} = (E_{t-1}(r_t) \; - \; \beta_0)/ \, \beta_1\sigma^2_t, \qquad (6)$$

where $Q_{1,t}$ is the fraction of shares demanded at time t, r_t is the ex-post return at t, E_{t-1} is the expectation as of time t -1, β_0 is the rate of return on the risk-free asset, σ^2_t is the conditional variance (risk) at t, and β_1 is the coefficient of risk aversion.[4] The second group of investors follows a positive feedback strategy, i.e., they buy (sell) after price increases (decreases). Thus, their demand function is given by

$$Q_{2,t} = \rho r_{t-1}, \tag{7}$$

where, $\rho > 0$.[5] It should be noted that positive feedback trading is not necessarily irrational. It may be a result of certain portfolio insurance strategies and the use of stop-loss orders. Portfolio insurance is rational if risk aversion declines rapidly with wealth (e.g., Sentana and Wadhwani 1992). In equilibrium, all shares must be held, i.e., $Q_{1,t} + Q_{2,t} = 1$. It follows from (6) and (7) that

$$E_{t-1}(r_t) = \beta_0 + \beta_1\sigma^2_t - \beta_1\rho\sigma^2_t r_{t-1} . \tag{8}$$

The term $-\beta_1\rho\sigma^2_t$ in (8) implies that the presence of positive feedback trading will induce negative autocorrelation in returns, and the higher the volatility the more negative the autocorrelation. The first group of investors will not necessarily exploit the higher autocorrelation, in absolute terms, that is induced by positive feedback trading because the risk is higher. Equation (8) can be converted into a regression equation with a stochastic error term by setting $r_t = E_{t-1}(r_t) + \varepsilon_t$, yielding:

$$r_t = \beta_0 + \beta_1\sigma^2_t + (\varphi_0 + \varphi_1\sigma^2_t)\, r_{t-1} + \varepsilon_t \tag{9}$$

where, $-\beta_1 \rho = \varphi_1$, and φ_0 has been added to pick up possible constant autocorrelation. The presence of positive feedback trading would imply that φ_1 is negative and statistically significant.

IV. C. Nested Models

The conditional mean specification can be expanded to nest both the LeBaron and the positive feedback trading models, as follows:

$$r_t = \beta_0 + \beta_1\sigma^2_t + (\varphi_0 + \varphi_1\sigma^2_t + \varphi_2\exp\{-\sigma^2_t\})r_{t-1} + \varepsilon_t \qquad (10)$$

This model implies that as volatility approaches zero return autocorrelation approaches $\varphi_0 + \varphi_2$ whereas, for high volatility values the autocorrelation becomes $\varphi_0 + \varphi_1\sigma^2_t$. If, $\varphi_2 = 0$ and $\varphi_1 < 0$ then it can be concluded that positive feedback trading is driving the time-varying autocorrelation.[6]

V. ESTIMATION AND EMPIRICAL FINDINGS

Estimation of the models outlined above requires that the functional form of the conditional density function $f(\mu_t, \sigma^2_t, \nu)$ be specified. In most applications the normal density function is used. However, the normalized residuals obtained from ARCH-type models that assume normality appear to be leptokurtic, thereby rendering standard t-tests unreliable. To deal with this problem, distributions that allow for fatter tails such as the student's t or the Generalized Error Distribution (GED) can be used instead. In this paper we employ the GED distribution[7]. Estimates of the parameters are obtained by maximizing the log-likelihood function over the sample period[8]. Given that the log-likelihood function is highly nonlinear in the parameters, numerical maximization techniques are used to obtain estimates of the parameter vector. The method of estimation used is based on the Berndt, Hall, Hall, and Hausman (1974) algorithm[9].

Table 3, reports the maximum likelihood estimates for the three models using daily returns. Likelihood ratio tests suggest that the optimal lag structure for the volatility process is $p = 1$, $q = 1$. The focus of this investigation is the possible interaction between volatility and autocorrelation. The constant part of the autocorrelation, φ_0, is positive and significant across all models. Focusing on the LeBaron model, it can be seen that the coefficient φ_1 that links autocorrelation to volatility is positive and statistically significant. Thus, stock returns in the Cyprus Stock Exchange appear to be linked to their past history in a nonlinear fashion meaning that past returns as well as past volatilities can be used to improve one-step ahead forecasts. Specifically, during periods of high (low) volatility the magnitude of the autocorrelation is low (high). In this respect our findings agree with LeBaron's evidence for the U.S. stock market. The estimated parameters suggest that during low volatility periods the first-order autocorrelation approaches 0.2160 (i.e., $\varphi_0 + \varphi_1$). During high volatility periods, autocorrelation approaches 0.1195. Thus, there is an inverse relationship between the absolute of the autocorrelation and the level of volatility.

The coefficients describing the conditional variance process, α_0, α_1, and γ_1, are highly significant, while δ is marginally significant (at the 10% level). The results for the volatility parameters are similar across all three models suggesting that the conditional variance is not affected significantly by the particular specification used for the conditional mean. Thus, in all instances current volatility is a function of the magnitude of last period's standardized innovation and last period's volatility. The autoregressive nature of volatility is important in situations where forecasts of future volatility are needed.[10]

Turning to the positive feedback trading model (Table 3, column 2) it can be seen that the parameter of interest that is governing the autocorrelation of returns, i.e., φ_1 is negative, as predicted by the model and statistically significant. As pointed out earlier, nonsynchronous trading and/or time variation in ex-ante returns can cause positive autocorrelation in ex-post returns, at least in high frequency data. Positive feedback trading, on the other hand, causes negative autocorrelation that rises, in absolute terms, with the level of volatility. The implication is that positive feedback trading is an important determinant of short-term movements in the CYSE, in line with the findings of Sentana and Wadhwani (1992) for the U.S. stock market. Thus, it appears that certain aspects of stock return dynamics are similar across national stock markets irrespective of size or stage of development. This should not be surprising, given the growing interdependence of stock markets around the world. The greater predictability (negative autocorrelation) that is induced by feedback traders may not necessarily produce arbitrage opportunities for rational risk- averse investors because volatility also rises. This is extremely important if risk averse investors have short holding horizons and are concerned about liquidating mispriced assets. The extent of potential mispricing is likely to increase during periods of higher volatility because during those periods positive feedback traders have a greater influence on prices. Comparing the LeBaron model and the feedback model it appears that the latter provides a slightly better fit judging by the higher value of the likelihood function. Since the models are not nested however, we cannot say with a high degree of certainty which model is better. This question can be addressed by examining the results of the model that nests both the LeBaron and the feedback models, reported in Table 3, column 3. Interestingly, the feedback parameter φ_1 is negative and statistically significant, as predicted by the model, whereas, the LeBaron parameter φ_2 while having the right sign, is in fact statistically insignificant. The clear implication from this finding is that the inverse relation between volatility and autocorrelation is due to positive feedback trading pursued by some investors in the market.

It is typical for stock returns, especially from emerging markets, to exhibit substantial unconditional kurtosis. The unconditional kurtosis in Table 1 is 39.2598. The kurtosis measures of the standardized residuals obtained from the three models are 16.0324, 16.16340 and 16.3198 respectively. This implies that successful modeling of nonlinearities in the conditional mean and the conditional variance can go a long way toward reducing kurtosis, in this case an approximate 60% reduction. For all three models, the estimated parameter ν is well below 2, indicating substantial deviation from normality. Conventional t-tests reject the hypothesis that $\nu = 2$. This confirms the earlier assertion, that departures from normality observed in the raw return series cannot be entirely attributed to nonlinearities in the first- and second-moments.

The validity of these empirical findings, of course, depends on the correct specification of the model. A minimum requirement is that the standardized residuals are zero-mean, unit-variance i.i.d. processes. As can be seen from Table 3, the means and variances of the standardized residuals have zero mean and unit variance. The estimated Ljung-Box statistics for five and ten lags reject the hypothesis of nonlinear dependence in the squared normalized residuals. Thus, all three models appears to be successful in terms of describing linear and nonlinear dependencies in the returns.

As mentioned earlier, correct specification of the conditional variance is extremely important when the latter enters the conditional mean equation. Diagnostics, such as the LB statistics for the squared normalized residuals, are useful in terms of detecting any remaining nonlinear structure. However, they are not designed to test how well the model captures the asymmetry in the conditional variance, or the impact of the magnitude of positive and negative innovations on volatility. Results of these tests are reported in Table 4. The t-statistics for the Sign Bias Test are significant in all three cases. However, the remaining individual tests and the F-test fail to reject the hypothesis that the normalized estimated residuals are i.i.d. Thus we conclude that all three models are well specified.

It is quite common however, to have well-specified models that do very poorly in out of sample forecasting. To test the forecasting performance of the three models we perform one-step ahead out-of-sample forecasting of the conditional mean and the conditional variance. Specifically, we use the first 2,000 observations to estimate the three models as well as benchmark models. For the conditional mean, the benchmark or, null model is a simple autoregressive process of order one with constant variance (i.e. AR1). In other words returns are assumed to follow the process given by: $r_t = a + b\, r_{t-1} + e_t$. For the conditional variance estimation the benchmark model is the Exponentially-Weighted-Moving-Average (EWMA) used by J.P. Morgan which assumes that the conditional variance follows the process given by: $\sigma^2_t = (1-\lambda)r^2_{t-1} + \lambda\sigma^2_{t-1}$. The AR1 parameters along with parameter λ in the EWMA model are estimated using maximum likelihood.[11] Subsequently the models are updated by adding one more observation until the hold-out-sample of 2,080 is depleted. The one step ahead forecast is retained in each step and a forecast metric based on the Root-Mean-Square-Error (RMSE) is estimated. The results for the three models as well as the benchmark model are reported in Table 5. First, in terms of the conditional mean, the RMSE produced by the three nonlinear models is smaller than the RMSE produced by the benchmark model. On average, forecasting performance can be improved by 7% over and above the simple AR(1) model with EWMA variances. In this respect, a test statistic for forecasting accuracy due to Diebold and Mariano (1995) confirms that the three models outperform that benchmark model. The relevant estimated statistics are 2.7202, 2.1854 and 2.1732 respectively.[12] In large samples this statistic follows standard normal distribution. In terms of the conditional variance the results are similar in the sense that the RMSE is smaller for the three models compared to the EWMA. In terms of forecasting improvement however, the gains for the variance are smaller, approximately 2.5%. Indeed, the same test fails to confirm significant improvements in the forecasting ability of the three models versus the benchmark. In should be noted though that the EWMA model is a very sophisticated model capable of capturing time varying variances. In fact the EWMA is a GARCH(1,1) one model with zero intercept and infinite persistence.

Overall, the nonlinear models used in this paper are capable of describing the dynamics of stock returns quite satisfactorily. Moreover, by nesting the LeBaron and the feedback model we were able to determine that the inverse relationship between volatility and autocorrelation is due to positive feedback trading followed by some traders.

VI. SUMMARY AND CONCLUSIONS

This paper has examined the short-term return dynamics, in particular the interaction between volatility and autocorrelation in index stock returns in the Cyprus Stock Market. The empirical evidence shows that first-order autocorrelation is linked to the conditional variance in a nonlinear fashion. Specifically, periods of high (low) volatility are associated with low (high) autocorrelations. Nonsynchronous trading or time variation in expected returns cannot be the cause of this phenomenon, because typically they give rise to linear, time-invariant autocorrelation. Nested hypothesis testing reveals that the inverse relationship between volatility and autocorrelation can be attributed to positive feedback trading strategies. The impact of feedback trading is to produce negative first order autocorrelation in stock returns, which becomes more negative as the level of volatility rises.

Stock returns are highly leptokurtic. The Generalized Error Distribution (GED) with endogenously estimated degrees of freedom captures the excess leptokurtosis in the standardized residuals quite satisfactorily. An extensive array of specification tests as well as out-of-sample forecasting, show that the tested models explain nonlinearities in the first and second moments quite well.

ENDNOTES

1. At the same time, this emerging market is not yet as liquid and efficient as other developed markets, with prices potentially influenced by feedback trading strategies, a “herd” mentality or rumors.

2. All hypotheses testing in this paper is done at the 5% level for simplicity and uniformity.

3. A negative δ is consistent with the so-called "leverage hypothesis," whereby negative returns increase the debt/equity ratio. As a result, negative returns are followed by higher volatility than positive returns of an equal size (see also Nelson 1991).

4. If all investors had the same demand function given by (6), then in equilibrium $E_{t-1}(r_t) - \beta_0 = \beta_1\sigma^2_t$, which is the dynamic Capital Asset Pricing Model proposed by Merton (1973).

5. If $\rho < 0$ there would be negative feedback trading.

6. It should be noted that it is the interaction of risk-averse utility maximizers and positive-feedback traders that produces negative autocorrelation that rises in magnitude as volatility rises. Negative feedback trading would produce positive time-varying autocorrelation instead. However, if all investors followed positive-feedback trading strategies then returns would exhibit positive autocorrelation.

7. The density function of the GED is given by

$$f(\mu_t, \sigma_t, \nu) = \nu/2\ [\Gamma(3/\nu)]^{1/2}[\Gamma(1/\nu)]^{-3/2}(1/\sigma_t)\exp\{-[\Gamma(3/\nu)/\Gamma(1/\nu)]^{\nu/2}\left|\varepsilon_t/\sigma_t\right|^{\nu}\}$$

where $\Gamma(.)$ is the gamma function, and ν is a scale parameter capturing the degree of deviation from normality (fat tails or excess kurtosis) to be estimated endogenously. This parameter controls the shape of the distribution allowing the GED to nest several other densities. For example, if $\nu = 2$, $f(.)$ yields the normal distribution, while for $\nu = 1$ it yields the Laplace or, double exponential distribution (see also Nelson 1991).

8. The sample log-likelihood can be expressed as follows: $L(\Theta) = \Sigma^{T}_{t=1}\log f(\mu_t, \sigma^2_\tau, \nu)$.

9. Several sets of initial values were used to ensure that the final estimates are robust to the choice of initial values.

10. This is obviously true for the evaluation of derivative securities, such as options and options on futures, where ex-ante volatility measures are critical inputs.

11. The choice of the AR1 specification for the benchmark model is based on the need to have nested models so that meaningful statistical inferences can be applied. For example, with suitable parameter restrictions the three models used for the conditional mean can be reduced to the AR1 model. Likewise, the choice of the EWMA specification for the benchmark conditional

variance was made due to the popularity and wide use of the model as well as it being nested within the GARCH(1,1). It should be noted that the results were almost identical when the GARCH(1,1) model was used.

12. The version of the Diebold-Mariano test used here is the sign test based on the residuals obtained from the null (i.e., benchmark) and the alternative models.

REFERENCES

Bekeart, G. and Wu, G., 2000. “Asymmetric Volatility and Risk in Equity Markets”, The Review of Financial Studies, 13, 1-42.

Berndt, E.K., Hall, H.B., Hall, R.E., and Hausman J.A., 1974. “Estimation and Inference in Nonlinear Structural Models”. Annals of Economic and Social Measurement, 4, pp. 653-666.

Black, F., 1976. "Studies in Stock Price Volatility Changes." Proceedings of the 1976 Business Meeting of the Business and Economic Statistics Section. American Statistical Association, 177-181.

Bollerslev, T., 1986. “Generalized Autoregressive Conditional Heteroskedasticity.” Journal of Econometrics, 31, 307-327.

Bollerslev, T., Engle, R.F. and Nelson D.B. 1994. "ARCH Models." in Engle R.F. and McFadden, D.L., eds., Handbook of Econometrics, Vol. 4, pp. 2959-3038., Elsevier Science B.V.

Booth, G.G., Kaen F.R., Koutmos G. and Sherman H.C. 2000. “Bundesbank Intervention Effects Through Interest Rate Policy”, Journal of International Financial Markets Institutions and Money, 10, 263-274.

Campbell, J.Y., Grossman, S.J. and Wang, J. 1993. "Trading Volume and Serial Correlation in Stock Returns." Quarterly Journal of Economics, Vol. CVIII, 4, 905-939.

Christie, A.A., 1982. "The Stochastic Behavior of Common Stock Variance: Value Leverage and Interest Rate Effects." Journal of Financial Economics, 10, 407-432.

Conrad, J. and Kaul, G. 1988. “Time-Variation in Expected Returns.” Journal of Business, 61, 409-425.

Diebold, F.X., and Mariano, R.S., 1995. “Comparing Predictive Accuracy”, Journal of Business and Economic Statistics, 13, 253-263.

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

Engle, R.F., and Ng, V.K., 1993. "Measuring and Testing the Impact of News on Volatility." The Journal of Finance, 48, 1749-1778.

Fama, E.F. 1965. "The Behavior of Stock Market Prices." Journal of Business, 38, 34-105.

Fama, E. F. and French, K. R. 1988. “Permanent and Transitory Components of Stock Prices.” Journal of Political Economy, 96, 246-273.

Fisher, L., 1966. “Some New Stock Market Indices.” Journal of Business, 39, 191-225.

Merton, R. C., 1973. "An Intertemporal Capital Asset Pricing Model," Econometrica, 41, 867-888.
Koutmos G. and Knif J., 2002. "Estimating Systematic Risk Using Time-Varying Distributions", European Financial Management, 8, 59-73.

Koutmos G., 1997. "Do Emerging and Developed Markets Behave Alike? Evidence from Six Pacific-Basin Stock Markets", Journal of International Financial Markets Institutions and Money, 7, 221-234.

LeBaron, B., 1992. "Some Relations Between Volatility and Serial Correlations in Stock Market Returns", The Journal of Business, 65, 199-219.

Mandelbrot, B. 1963. "The Variance of Certain Speculative Price." Journal of Business, 36, 394-419.

Moschini G.C. and Myers R.J., 2002. "Testing for Constant Hedge Ratios in Commodity Markets: A Multivariate GARCH Approach", Journal of Empirical Finance, 9, 589-603.

Nelson, D., March 1991. "Conditional Heteroskedasticity in asset Returns: A New Approach." Econometrica, 59, 347-370.

Scholes M. and Williams, J. 1977. "Estimating Betas from Nonsynchronous Data." Journal of Financial Economics, 5, 309-327.

Scruggs J.T. and Glabadanidis P., 2003. "Risk Premia and the Dynamic Covariance Between Stock and Bond Returns", Journal of Financial and Quantitative Analysis, 38, 295-316.

Sentana, E. and Wadhwani, S. 1992. "Feedback Traders and Stock Return Autocorrelations: Evidence from a Century of Daily Data." The Economic Journal, 102, 415-425.

Zhou A. 2002. "Modeling the Volatility of the Heath-Jarrow-Morton Model: A Multifactor GARCH Analysis", Journal of Empirical Finance, 9, 35-56.

Table 1. Diagnostic Checks and Descriptive Statistics on CYSE Returns

Mean (μ)	0.0404
Standard Deviation (σ)	1.3387
Skewness (S)	1.8517*
Kurtosis (K)	39.2598*
Kolmogorov-Smirnov (D)	0.1141*
LB(10)	260.8222*
$LB^2(10)$	378.5154*

Notes: (*) denotes statistical significance at the 5% level (at least). The sample period for the daily CYSE stock return series extends from January 21, 1985 to November 22, 2002 for a total of 4,080 observations. μ, σ, S, and K are the sample estimated mean, standard deviation, skewness and excess kurtosis, respectively. KS D-stat is the Kolmogorov-Smirnov statistic testing for normality. Its sample critical value at the 5% level is $1.36/\sqrt{n}$, where n is the sample size. LB(10) and $LB^2(10)$ are the Ljung-Box χ^2 statistics for 10 lags, calculated for returns and squared returns, respectively. The null hypothesis is that all autocorrelations up the 10^{th} lag are jointly zero.

Table 2. Volatility Specification Tests for Filtered Returns

Sign Bias (t-tests)	-0.7034
Negative Size Bias (t-tests)	-8.2620*
Positive Size Bias (t-tests)	12.0202*
Joint Test F(3,4072)	87.7218*

Notes: (*) denotes statistical significance at the 5% level (at least). This table reports some recent tests proposed by Engle and Ng (1993). These tests are specified as follows:

Sign Bias: $z^2_t = a + b\, S^-_t + e_t$ (i)

Negative Sign Bias: $z^2_t = a + b\, S^-_t\, \varepsilon_{t-1} + e_t$ (ii)

Positive Sign Bias: $z^2_t = a + b\, (1 - S^-_t)\, \varepsilon_{t-1} + e_t$ (iii)

Joint Test: $z^2_t = a + b_1\, S^-_t + b_2\, S^-_t\, \varepsilon_{t-1} + b_3\, (1 - S^-_t)\, \varepsilon_{t-1} + e_t$ (iv)

where S^-_t is a dummy variable that takes the value of unity if ε_{t-1} is negative, and zero otherwise. All t-statistics refer to the coefficient b in regressions (i), (ii), and (iii), while the joint test F(3,4072) is referring to regression (iv). $z_t = \varepsilon_t/\sigma_t$. The normalized residuals are based on an AR1 model applied to returns.

Table 3. Maximum Likelihood Estimates of Three Volatility Model Specifications

	LeBaron Model	Feedback Model	Nested Model
β_0	0.0091 (2.069)*	0.0088 (2.011)*	0.0098 (2.222)*
β_1	-0.0060 (-0.800)	-0.0015 (-0.212)*	-0.0026 (-0.342)
φ_0	0.1195 (6.348)*	0.1869 (14.515)*	0.1557 (6.977)*
φ_1	0.0965 (2.882)*	-0.0050 (-2.448)*	-0.0032 (-2.202)*
φ_2			0.0472 (1.276)
α_0	-0.0070 (-1.525)	-0.0068 (-1.495)	-0.0071 (-1.566)
α_1	0.03281 (15.717)*	0.3228 (15.543)*	0.3241 (15.725)*
γ	0.9853 (380.250)*	0.9858 (383.265)*	0.9858 (385.449)*
δ	-0.0661 (-1.878)	-0.0674 (-1.919)	-0.0683 (-1.955)*
ν	0.9080 (46.202)*	0.9074 (46.235)*	0.9122 (46.563)*
Log L	-4206.22	-4203.92	-4203.49
Kurtosis (K)	16.0324*	16.1634*	16.3198*
$E(z_t)$	0.0608	0.0584	0.0584
$E(z^2_t)$	1.0772	1.0783	1.0783
$LB^2(5)$	2.2744	2.3048	2.3048
$LB^2(10)$	5.3369	5.2326	5.2326

Notes: (*) denotes significance at the 5% level (at least). Numbers in parentheses are the t-statistics. z_t is the model normalized residual. $LB^2(n)$ is the Ljung-Box statistics for z^2_t, using 5 and 10 lags respectively. See also Table 1 notes.

Table 4. Volatility Specification Tests for Normalized Residuals

	LeBaron Model	Feedback Model	Nested Model
Sign Bias (t-test)	-2.0384*	-2.0426*	-2.0557*
Negative Size Bias (t-test)	0.0322	0.0686	0.0747
Positive Size Bias (t-test)	1.0148	1.0175	0.9954
Joint Test (F-test)	1.7112	1.6979	1.7067

Notes: (*) denotes statistical significance at the 5% level (at least). The tests are applied to the estimated normalized residuals from the three models. See also Table 2 notes.

Table 5. Out of Sample Forecasting

	Benchmark Model*	LeBaron Model	Feedback Model	Nested Model
First Moment RMSE	1.4187	1.3181	1.3187	1.3188
Second Moment RMSE	11.0890	10.8041	10.7792	10.7717

Notes: * The benchmark model used for first-moment (mean) forecasting is based on a linear AR(1) model applied to returns. The benchmark model used for second-moment (volatility) forecasting is based on the Exponentially Weighted Moving Average (EWMA) used by JP Morgan. RMSE is the Root Mean Square Error based on the one-step ahead forecast errors.