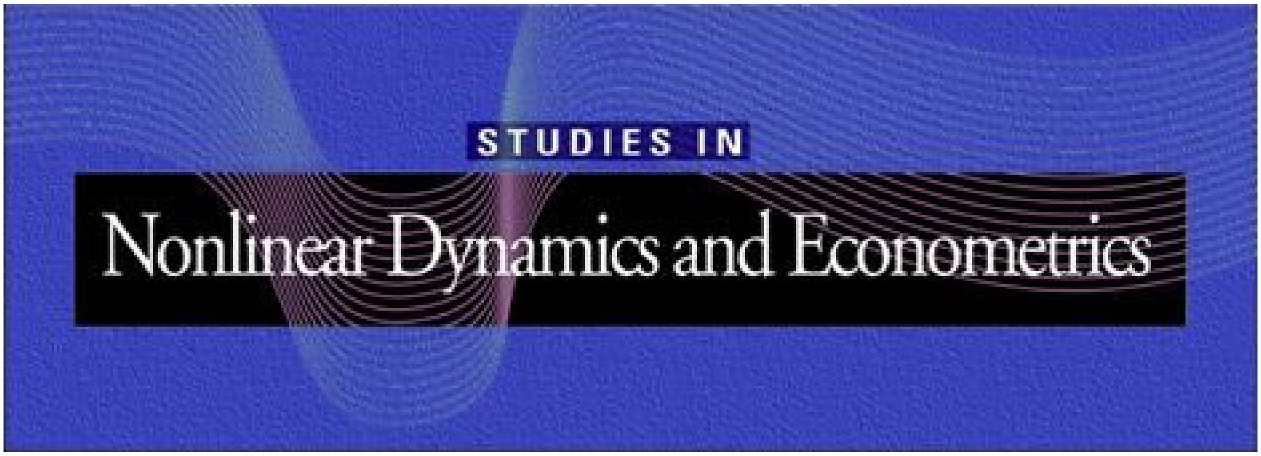

Studies in Nonlinear Dynamics and Econometrics

Quarterly Journal
January 1998, Volume 2, Number 4
The MIT Press

Studies in Nonlinear Dynamics and Econometrics (ISSN 1081-1826) is a quarterly journal published electronically on the Internet by The MIT Press, Cambridge, Massachusetts, 02142. Subscriptions and address changes should be addressed to MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142; tel.: (617) 253-2889; fax: (617) 577-1545; e-mail: journals-orders@mit.edu. Subscription rates are: Individuals $40.00, Institutions $130.00. Canadians add 7% GST. Prices subject to change without notice.

Early News is Good News: The Effects of Market Opening on Market Volatility

Giampiero M. Gallo
Dip. di Statistica "G. Parenti"
Università di Firenze
Firenze, Italy
and
Economics Department
European University Institute
GALLOG@stat.ds.unifi.it

Barbara Pacini
Dip. di Statistica "G. Parenti"
Università di Firenze
Firenze, Italy
PACINI@stat.ds.unifi.it

Abstract. *In this paper, we examine the characteristics of market opening news and its impact on the estimated coefficients of the conditional volatility models of the GARCH class. We find that the differences between the opening price of one day and the closing price of the day before have different characteristics when considering various stock-market indices on which options are actively traded. The impact of a suitable positive-valued transformation of these differences has the effects of modifying the direct impact of daily innovations on volatility and reducing the estimated overall persistence of such innovations. The overall contribution of the variable is evaluated in an out-of-sample forecasting exercise, where we obtain significant improvements above the simple GARCH model.*

Acknowledgments. Thanks are due to an anonymous referee for helping us to better focus certain issues, and to Norm Swanson for his encouragement in pursuing this line of research. We would also like to thank Lucia Buzzigoli for helpful discussion. Financial support from the Italian MURST and CNR is gratefully acknowledged.

Keywords. volatility, GARCH models, news, persistence, forecasting

1 The Issues

Among the possible nonlinearities exhibited by financial time series, volatility is the one aspect that has received special attention in the literature, with a number of statistical models suggested, starting with the seminal paper by Engle (1982), to explain the behavior of second moments of returns. In particular, the success of these models is to be found in their relative simplicity and their good performance in forecasting variances conditional on an information set. Apart from the possibility of making the proper inference by deriving the correct standard errors, the specific interest in second-moment behavior lies in the usefulness of time-varying measures of volatility in derivative products (such as options) pricing, and in risk evaluation. From an empirical point of view, the strategy usually followed (cf. Engle and Ng, 1993) is to transform prices into returns (since, most commonly, unit-root tests performed on the logarithms of prices justify working in logarithmic first differences), and to extract systematic components, such as the so-called anomalies (week-end and January effects), and linear-autoregressive components (usually a low order suffices). According to the efficient market hypothesis, what is left should be a series of realizations of independent random variables. Yet, the non–Gaussianity of the empirical distribution of such a series implies that serial uncorrelation is not sufficient for independence, and, in fact, the results by Ding, Granger, and Engle (1993) show that positive powers of absolute returns exhibit a remarkable serial correlation, the highest being for an exponent of 1.25. The efforts to model such a dependence seem to have been mostly devoted to squared returns in view of the economic interpretability of the results.

The most popular models for conditional volatility are GARCH-type processes (Bollerslev, Engle, and Nelson 1994), which are designed to reproduce various stylized facts of volatility clustering. The conditional variance is modeled as a function of past conditional variances and innovations, and possible asymmetric behavior is accommodated by inserting a different impact of positive and negative innovations. Such an impact is well described by the so-called news-impact curve introduced by Engle and Ng (1993), which measures graphically the estimated effect of an arbitrary grid of past innovations (the "news") on the conditional volatility.

Such a device points out a relationship between the news in the markets and its effect on the volatility, although, from a theoretical point of view, it is still not clear which mechanisms are at work on the market which cause volatility clustering (the sign of the innovations notwithstanding). As a matter of fact, market efficiency would require an immediate adjustment to the news reaching the market by a proper discounting of its effects, and no further adjustment ensuing. This is in apparent contrast with the excessive persistence in the estimated effects of the impact of news to volatility, measured as the coefficient in a companion first-order difference equation for the conditional volatility. Almost invariably, financial returns (be they on equities, exchange rates, stock indices, and so on) on such an estimated persistence is higher than 0.9, which is too high a value to be reconciled with the historical behavior of squared returns.

In this paper we concentrate on GARCH and EGARCH models, and we are interested in four questions; namely, whether the difference between the opening price of one day and the closing price of the previous trading day:

1. alters the shape of the news impact curve,

2. is responsible for the high estimated persistence of GARCH and EGARCH models,

3. gives account of the leverage effects detected on stock markets, and/or

4. helps in improving the forecasting performance of GARCH and EGARCH models.

Since volatility forecasting is of interest in pricing derivative products such as options, we will focus on five stock indices of the U.S. market for which options are actively traded and the Dow Jones Industrial Average as a benchmark for comparing our approach with other approaches in the literature.

The paper is organized as follows. In Section 2 we present the characteristics of the stock indices in their coverage of different sectors and different market conditions. In Section 3 we refer to the model suggested by Clark (1973), and we discuss the modification of the basic assumption of the daily return being the sum of i.i.d. innovations, by keeping separate the portion of the innovation due to a different opening price from the previous day's closing price. In Section 4 we present the various models of conditional volatility that are used throughout the paper, and the estimation results. In Section 5 we analyze the impact of the overnight surprise on the news-impact curve by extending the suggestion by Engle and Ng (1993), and showing graphically the impact of a range of values of past innovations on the conditional volatility. In Section 6 we document the reduction in the estimated volatility. The forecast performance of the various models is evaluated in Section 7 by means of Diebold and Mariano's (1995) test, which shows the significant gains obtained in the present context. Concluding remarks follow.

2 The Stock Market Indices

We have chosen six stock market indices: the Dow Jones Industrial Average, the NASDAQ Composite, the Russell 2000 Index, the Standard & Poor 500, the AMEX Major Market, and the Philadelphia KBW Bank Index. We use daily data from January 4, 1993 to December 31, 1996 (1,042 observations) for estimation, and from January 2, 1997 to January 30, 1997 for forecasting comparison.

These indices vary in coverage of the various segments of the market, and with the exception of the Dow Jones Industrial Average, share the interesting aspect of having options traded on them. As a reference, therefore, we give a brief account[1] of the main characteristics of each index.

1. The Dow Jones Industrial Average (DJIA) is the "oldest continuous price measure in the United States" (Berlin 1990, p. 15) and is based on the prices of 30 widely traded blue chips. No options are traded on this index.

2. The NASDAQ Composite Index (NASDAQ) reflects the behavior of the over-the-counter (OTC) market, thus tracking the performance of smaller and more dynamic companies. The average is made over 4,300 stock prices. Options are traded on this index.

3. The Russell 2000 Index tracks the 2,000 most actively traded shares relative to medium-sized companies with market values between US$ 20M and US$ 300M, or about 9% of the total market (cf. Berlin 1990, p. 56). Options are traded on this index.

4. The Standard & Poor 500 Index (S&P 500) includes stocks of 500 relatively large companies which are traded in the NYSE, the AMEX, and the OTC, and comprises various industry segments (industrial, transportation, utilities, and financial). The S&P 500 index option is among the most actively traded index options.

5. The AMEX Major Market Index (AMEX) is designed to track the behavior of industrial companies without duplicating the DJIA. Its main interest, therefore, is that it constitutes an underlying price for option contracts that are not allowed on the DJIA.

6. The Philadelphia Stock Exchange KBW Bank Sector Index (Phlx KBW Bank) is computed based on the capitalization of 24 banking sector institutions scattered throughout the U.S. The interesting aspect for the purposes of this paper is the availability of option contracts and that the exercise-settlement value of the index is based on the opening price of the component stocks on the last trading day prior to expiration (cf., http://www.phlx.com/bkx.html).

To give an idea of the importance of the various indices in options trading, in Table 1 we reproduce the volumes and open interest of call-and-put options on these indices in one day of trading (January 3, 1997).

[1]Cf., also, http://www.cnbc.com/tickerguide/indices.html, the financial page of CNBC, which gives a description of these indices.

Table 1
Market Activity on Stock Index Options[a]

	Total Call Volume	Total Put Volume	Call Open Interest	Put Open Interest
NASDAQ Composite	200	17	375	451
Russell 2000	785	128	27,920	24,172
S&P 500	26,741	47,794	501,199	834,913
AMEX Major Market	231	86	2,719	4,042
Phlx KBW Bank	751	754	6,240	6,317

[a]Data referred to January 3, 1997 (from *Investor's Business Daily*, January 7, 1997).

3 Time-Varying Volatility and Overnight News

The stylized facts on asset returns point to the randomness of their volatility, with a certain degree of persistence exhibited by the data, giving rise to leptokurtic empirical distributions. The phenomenon that seems more closely correlated with volatility seems to be information arrival, as documented, for example, by Berry and Howe (1994), who measure the impact of news wire per unit of time on volume and volatility. From a statistical point of view, Clark (1973) shows how the consideration of the daily return as a sum of identical and independently distributed within-the-day price movements δ_{it} with a zero mean and a constant variance σ^2, that is,

$$\epsilon_t = \sum_{i=0}^{n_t} \delta_{it}, \tag{1}$$

gives rise to a stochastic variance model, since n_t in Equation 1 is a random variable indicating the number of trades in day t. By a central-limit theorem argument, Clark concludes that, conditional on n_t, ϵ_t is normal with variance equal to $n_t\sigma^2$. Leptokurtosis arises as a result for the unconditional distribution, since we are considering mixtures of normal distributions with n_t as the mixing variable.

The variable n_t is itself a latent random variable and, as recalled by Andersen (1995), it can be seen as driving both returns and volumes [as in (Tauchen and Pitts 1983)] in a bivariate model giving rise to a stochastic volatility model. More simply, it can be adopted within a univariate GARCH framework as an additional exogenous variable in need of being approximated, as in (Lamoureux and Lastrapes 1990a), where contemporaneous volume was chosen.

We will follow the second route, and suggest an alternative proxy for information arrival, starting from the consideration that the daily return measured from successive closing prices can be decomposed as:

$$\epsilon_t = \delta_{0t} + \sum_{i=1}^{n_t} \delta_{it}, \tag{2}$$

where δ_{0t} refers to the return between the opening price of one day and the closing price of the previous day. The assumption that δ_{0t} is assumed to be independently and identically distributed from the within-the-day price movements $\sum_{i=1}^{n_t} \delta_{it}$ seems very strong. News often do accumulate during market closing time, and are reflected in a market opening at a different level than the night before, with a discrete change of a considerable size.

To substantiate the claim of the lack of independence, we present graphical evidence in the form of a plot of δ_{0t} against $\sum_{i=1}^{n_t} \delta_{it}$, which shows a tendency of the scatter to be negatively sloped (with the exception of the NASDAQ). This impression is confirmed by the regression results, which signal a significantly negative relationship between the two variables (with the exception of NASDAQ and the Phila KBW indices). To take the analysis a step further, a nonparametric contingency table χ^2 test on the relative frequency of sign agreement between the two variables signals a rejection of the null hypothesis of stochastic independence between the two characters in all cases except the NASDAQ.[2]

The lack of independence prompts two questions on the nature of δ_{0t}: its importance relative to the daily return, and the shape of the frequency distribution of such opening surprises. To assess its importance, we standardize the δ_{0t} to the highest daily return in the sample for each index. By so doing we are aiming at characterizing both the frequency at which opening prices are different from the previous day's closing prices, and their relative size (Figure 2). Second, we estimate the probability of the opening news being close to zero and the threshold values for "abnormal" news measured, respectively, as the 5th and 95th percentiles of the probability density function estimated nonparametrically (Table 2).

Looking at Figure 2, we see how the dynamics of δ_{0t} vary by index. The Dow Jones Industrial Average exhibits the highest variability of this measure, and the surprises are noticeable since they are mostly included in the interval (-.3,+.3) and are quite symmetric about zero. The incidence of surprises on the other indices is less remarkable, although an increase in the volatility of the measure for the NASDAQ and the Philadelphia KBW Bank Sector Index is noticeable, starting from the end of 1995. The coincidence of the date of the largest surprises is to be stressed: for example, the surprise of February 16, 1993 is common to all indices (except the Dow Jones Industrial Average, which had a negative opening that day), and for three of them is even higher than the highest return in the sample period. By the same token, other surprises are more idiosyncratic (cf., the positive peaks for the NASDAQ around May and June, 1995).

As mentioned, in Table 2 we analyze in detail the distribution of δ_{0t} whose density we estimated nonparametrically by kernel smoothing.[3] The first column reports the estimated probability that the opening return is in a neighborhood of zero.[4] The second and third columns report the estimated 5% and 95% quantiles as annualized rates.

As the evidence shows, the grouping that we had seen graphically in Figure 2 is confirmed, with the Dow Jones Industrial Average, NASDAQ Composite, and Philadelphia KBW Bank Sector Indices exhibiting the highest quantiles and the lowest probabilities to open the market at the same level as the previous close. At the other end of the spectrum, the Russell 2000 and the AMEX Major Market indices show a very high probability of being around zero and low estimated quantiles. The Standard & Poor 500 Index has a 50% chance of opening at a different level than the previous close. The symmetry of the distribution (comparison of the values of the percentiles) signals no systematic pattern of good or bad news accumulating during the night.

All this evidence points, on the one hand, to some peculiarities of the overnight accumulation of news which make it different from normal market activity. By the same token, we can emphasize the fact that each market has different sizes and working mechanisms, so that the behavior of the opening news is quite different across indices. Our working hypothesis is that this difference is likely to affect the trading during the day, resulting in an impact on the daily volatility which remains to be investigated. This is done in the next section.

4 Volatility Models

From Engle's (1982) seminal paper on the possibility of modeling volatility clustering by a simple autoregressive scheme on past-squared innovations, many models have been suggested that capture various aspects of volatility pointed out by empirical regularities. We will formally define innovations as the portion of returns that is not explained by a linear-autoregressive component or by deterministic components that

[2] Results available upon request.

[3] We adopted a Gaussian kernel with optimal bandwidth selection, according to Silverman's (1986, p. 45) heuristic formula.

[4] For numerical reasons the interval around zero is of a length variable by index: the largest is equal to 5.34E-05.

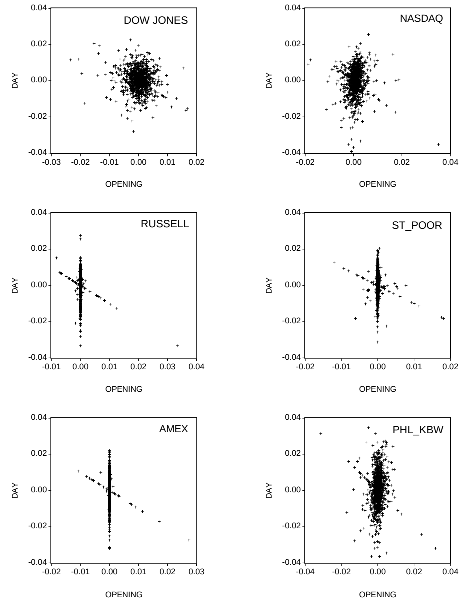

Figure 1
Scatterplot of opening return δ_{0t} against the net return during the day $\sum_{i=1}^{n_t} \delta_{it}$.

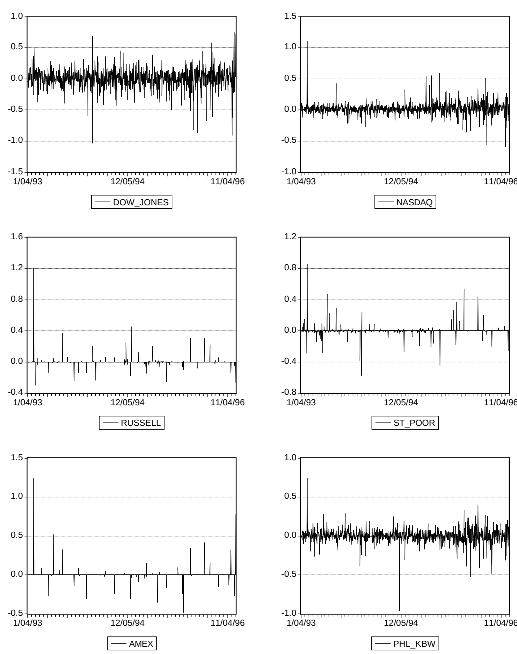

Figure 2
Opening returns standardized to the highest daily return in the sample.

Table 2
Distribution Characteristics of Opening Returns (Annual Rates)

	$P[\delta_{0t} \in I_{\varepsilon(0)}]$	5% Quantile	95% Quantile
DJIA	0.0060	−1.802	1.804
NASDAQ	0.0102	−1.254	1.698
Russell 2000	0.8459	−0.007	0.024
S&P 500	0.5016	−0.102	0.106
AMEX	0.9659	−0.001	0.011
Phlx KBW Bank	0.0133	−1.334	1.469

represent the so-called anomalies (day-of-the-week, or January effects):

$$\epsilon_t = r_t - A(L) r_{t-1} - d_t' \phi,$$

where the order of the polynomial $A(L)$ in the lag operator L is typically low (and its roots close to zero), and the vector d_t collects all relevant dummy variables for the anomalies.

The GARCH class of models starts from a basic assumption that these innovations follow a conditionally heteroskedastic process, namely, conditional on an information set I_{t-1},

$$\epsilon_t | I_{t-1} \sim (0, \sigma_t^2),$$

where σ_t^2 is a time-varying variance whose dynamics can be written as (cf., Campbell, Lo, and MacKinlay 1997, p. 488):

$$\begin{aligned} \sigma_t^2 &= \omega + \beta \sigma_{t-1}^2 + \alpha \epsilon_{t-1}^2 + \psi X_t \\ &= \omega + (\alpha + \beta) \sigma_{t-1}^2 + \alpha (\epsilon_{t-1}^2 - \sigma_{t-1}^2) + \psi X_t. \end{aligned} \tag{3}$$

As noted by Campbell, Lo, and MacKinlay (1997), the addition of explanatory variables known at time t, X_t, is "straightforward," and for reasons which will be clarified later, we prefer such a representation, which shows that the term $\epsilon_{t-1}^2 - \sigma_{t-1}^2$ (the shock to volatility) has a conditional mean equal to zero, and that the conditional volatility feeds on its past and on the impulses provided by the (positive valued) X_t. The sum $\alpha + \beta$ measures the rate at which the shocks to volatility instantaneously incorporated in σ_t^2 are transmitted to future volatility.

The other model considered is the EGARCH derived by Nelson (1991) to avoid some of the problems with the simple model, namely, the possibility of obtaining negative estimated variances (and hence the need for imposing non-negativity constraints on the parameters), and the symmetric treatment of negative and positive shocks. In detail, the EGARCH(1,1) model can be written as:

$$\log(\sigma_t^2) = \omega + \beta \log(\sigma_{t-1}^2) + \alpha \left(\frac{|\epsilon_{t-1}|}{\sigma_{t-1}} - \sqrt{2/\pi} \right) + \gamma \frac{\epsilon_{t-1}}{\sigma_{t-1}} + \psi X_t. \tag{4}$$

The parameter β measures the persistence effect in the dynamics of the variance, while the γ parameter measures the impact of negative innovations on the variance (with an expected negative sign on the basis of the so-called leverage effect).

In our case, since we focus on the overnight surprise, we turn it into a suitable X_t variable for Equations 3 and 4 by taking the absolute value of δ_{0t} and dubbing it ONI (overnight indicator). With this variable in hand, which is always available, we are equipped for analyzing the effect of such an indicator on the estimates of the GARCH (Table 3) and EGARCH (Table 4) models.

The estimated coefficient on the ONI is always significant (an exception being the Standard & Poor 500 Index) and positive (with the exception of Russell 2000 and AMEX Major Market indices—in such cases, though, the estimated variance never turns negative). Note that the insertion of the ONI has the effect of sharply decreasing the estimated β in the GARCH model, and of β in the EGARCH model.

Table 3
GARCH(1,1) and GARCH(1,1) with ONI[b]

	α	β	α	β	ψ(ONI)
DJIA	0.056*	0.871**	0.014	0.360**	0.006**
	(0.023)	(0.051)	(0.037)	(0.127)	(0.001)
NASDAQ	0.118**	0.775**	0.116**	0.619**	0.004**
	(0.035)	(0.058)	(0.044)	(0.092)	(0.001)
Russell 2000	0.242**	0.475**	0.325**	0.297**	−0.004**
	(0.071)	(0.112)	(0.068)	(0.110)	(0.0004)
S&P 500	0.044*	0.909**	0.040*	0.922**	0.0008
	(0.020)	(0.049)	(0.016)	(0.039)	(0.0009)
AMEX	0.043*	0.903**	0.149**	0.600**	−0.003**
	(0.019)	(0.046)	(0.035)	(0.102)	(0.0003)
Phlx KBW Bank	0.103**	0.740**	0.160**	0.070	0.017**
	(0.037)	(0.103)	(0.045)	(0.080)	(0.003)

[b] *significant at 5% level; **significant at 1% level.

Table 4
EGARCH(1,1) and EGARCH(1,1) with ONI[c]

	β	α	γ	β	α	γ	ψ(ONI)
DJIA	0.852**	0.121**	−0.119**	0.763**	0.042	0.150**	0.132**
	(0.056)	(0.044)	(0.033)	(0.063)	(0.045)	(0.033)	(0.038)
NASDAQ	0.830**	0.178**	−0.195**	0.746*	0.130	−0.237**	0.096*
	(0.050)	(0.068)	(0.044)	(0.053)	(0.093)	(0.050)	(0.043)
Russell 2000	0.650**	0.214*	−0.268**	0.568*	0.252*	−0.304**	−0.072
	(0.093)	(0.097)	(0.062)	(0.104)	(0.103)	(0.065)	(0.038)
S&P 500	0.898**	0.106*	−0.123**	0.895*	0.103*	−0.129**	−0.007
	(0.041)	(0.045)	(0.031)	(0.042)	(0.047)	(0.029)	(0.013)
AMEX	0.896**	0.111*	−0.086**	0.002	0.021	−0.128*	−0.370**
	(0.049)	(0.051)	(0.027)	(0.181)	(0.077)	(0.056)	(0.072)
Phlx KBW Bank	0.885**	0.128*	−0.088*	0.485*	0.242**	−0.131**	0.141**
	(0.053)	(0.057)	(0.034)	(0.129)	(0.075)	(0.044)	(0.046)

[c] *significant at 5% level; **significant at 1% level.

Lamoureux and Lastrapes (1990a) obtain similar results by including a variable measuring the total volume of stocks traded during the same day. Although for the single stocks the volume exchanged is usually available, for the market indices and averages, such data are not. When they are, it is possible to insert this variable as an additional variable in the conditional volatility expression. In our case, we had available just the data for the Dow Jones Industrial Average and the NASDAQ Composite Index. The results shown in Table 5 confirm on the one hand the results obtained for the stocks; and on the other hand, the fact that also the EGARCH model (not analyzed by Lamoureaux and Lastrapes) exhibits the same property.

5 Direct Impact on Volatility

The importance and relevance of the ONI variable can be investigated by first analyzing its direct impact on volatility. This can be suitably done by modifying the news-impact curve proposed by Engle and Ng (1993). The goal of such a tool is to show graphically the response of the conditional volatility (estimated by one of

Table 5
GARCH(1,1) and EGARCH(1,1) with Volume[d]

	α	β	ψ(Volume) ($\times 10^{-4}$)	β	α	γ	ψ(Volume)
DJIA	0.149**	0.600**	0.332	0.025	−0.101	−0.237**	1.661**
	(0.041)	(0.088)	(0.178)	(0.133)	(0.080)	(0.047)	(0.278)
NASDAQ	0.155**	0.547**	5.040**	0.468**	0.115	−0.355**	0.652**
	(0.052)	(0.130)	(1.760)	(0.102)	(0.069)	(0.056)	(0.123)

[d] *significant at 5% level; **significant at 1% level.

the GARCH-class models) to a range of values (positive and negative) of innovations. Evaluating the steepness of such a response, one can get an idea of the immediate impact of the innovation, thus complementing the characteristics of the process in terms of persistence and asymmetry.

In our case, we are interested in investigating how this response changes owing to the consideration of the ONI variable in the estimated model. In the GARCH specification, the curve is a function of the α parameter, i.e.,

$$\sigma_t^2 = \text{constant} + \alpha\epsilon_{t-1}^2,$$

where the constant is equal to $\omega + \beta\sigma^2 + \psi$ ONI, and σ^2 is the unconditional volatility. In what follows, to favor comparisons, we have estimated the constant with ONI set to its modal value of zero for all indices. This can be justified by the fact that we are interested in examining the modification of the shape of the *instantaneous* impact of the news after the ONI variable is inserted. In fact, the presence of ONI acts in two directions:

1. through a modification of the constant term (especially a change in the estimated β, usually a smaller one relative to the standard case), which determines a scale change in the conditional variance similar to the scale change in the Markov switching ARCH model of Kim (1993—which imposes a discrete change in the constant of the conditional variance); and

2. through a change in the estimated α, which determines in practice whether the curve's steepness increases or decreases. We will defer the analysis of ONI's practical impact in forecasting to the next section.

Analogously, for the EGARCH case, the news-impact curve is built by isolating the effects of ϵ_{t-1} from the rest, obtaining:

$$\sigma_t^2 = \begin{cases} \text{constant} * \exp\left(\frac{\gamma+\alpha}{\sigma}\epsilon_{t-1}\right) & \text{if } \epsilon_{t-1} > 0, \\ \text{constant} * \exp\left(\frac{\gamma-\alpha}{\sigma}\epsilon_{t-1}\right) & \text{if } \epsilon_{t-1} < 0. \end{cases}$$

In this case, the constant term is equal to $\sigma^{2\beta} * \exp(\omega - \alpha\sqrt{2/\pi} + \text{ONI}^{\psi})$, but again, it is evaluated in what follows by setting ONI to zero.

The results are presented graphically in Figures 3–6. We get an array of different results, which we summarize as follows:

1. The impact of ϵ_{t-1}^2 on the volatility of the NASDAQ Composite and of the Standard & Poor 500 indices changes very little as a result of the insertion of the ONI, both in the GARCH and EGARCH models. While this is expected for the latter, given the lack of significance of the corresponding coefficient, for the former the decrease in the impact is due to the decrease in the β coefficient, while the α stays the same.

2. The impact of the news on the Dow Jones Industrial Average decreases once the effect of the ONI is taken into account; for all the others, the insertion of this variable accounts for an increased steepness of the curve.

3. The positive portion of the curve for the EGARCH model is very flat and close to zero, owing to the closeness (in absolute value) of the coefficients α and γ.

4. The EGARCH news-impact curve is more "robust" to the insertion of the ONI variable, and hence the small difference between the two curves for the first four indices (Figures 3, 4 and 5A). For the AMEX

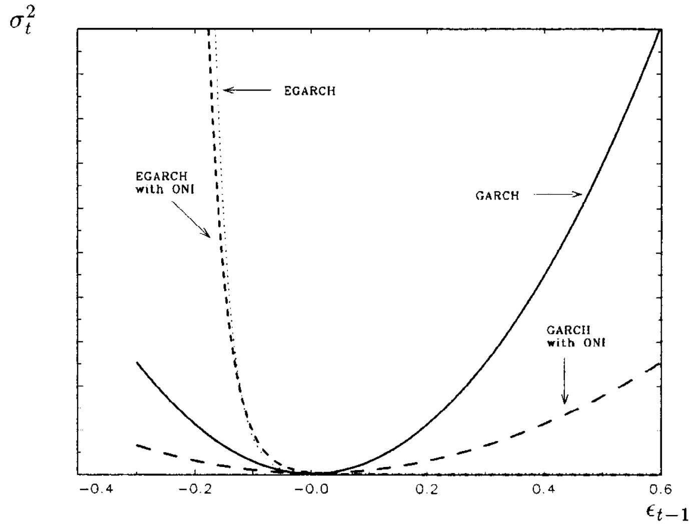

Figure 3
Dow Industrials. News impact curves of GARCH(1,1), EGARCH(1,1), GARCH(1,1) with ONI and EGARCH(1,1) with ONI.

Major Market Index, the insertion of ONI flattens the curve in the negative portion, while for the Philadelphia KBW Bank Sector Index, it makes the curve steeper on either side.[5]

It seems that the GARCH model is more sensitive overall to the presence of an overnight indicator than the EGARCH model. This should not be confused with the presence of asymmetric effects, since the insertion of ONI still leaves the need for the consideration of leverage effects. By the same token, the variable ONI is significant in most EGARCH models for the indices at hand [and also in the case of stocks, as in (Gallo and Pacini 1997)], and hence captures phenomena that are different from a leverage effect. Moreover, the graphical evidence presented in Figure 2 shows that the occurrence of overnight surprises might be a rare event and yet be significant in the conditional volatility models examined here.

6 Overall Impact on Volatility

The main characteristic of GARCH-type processes is that a squared current innovation higher (lower) than the previous volatility estimate increases (decreases) the current estimate of the volatility.

We are interested in measuring the persistence of the conditional variance estimate, that is, the rate at which the conditional variance reverts to its unconditional counterpart $\frac{\omega+\psi E(X)}{1-\alpha-\beta}$. For the GARCH model, standard measures of persistence are based on the sum $\alpha + \beta$ in Equation 3 (Lamont, Lumsdaine, and Jones 1996):

- the first is the sum of the direct effect α and indirect effect β as the coefficient of the associated first-order difference equation for the conditional variance;
- the so–called “half-life” index, namely $\frac{ln(2)}{ln(\alpha+\beta)}$, which represents the time that the conditional variance takes to revert halfway to its unconditional value; and
- a different definition of persistence can be derived from the decay of shocks to conditional volatility,

[5]This does not say anything about the forecasting performance of this model, nor about the actual contribution of ONI to forecasting, since what we are examining here is just the direct impact of news.

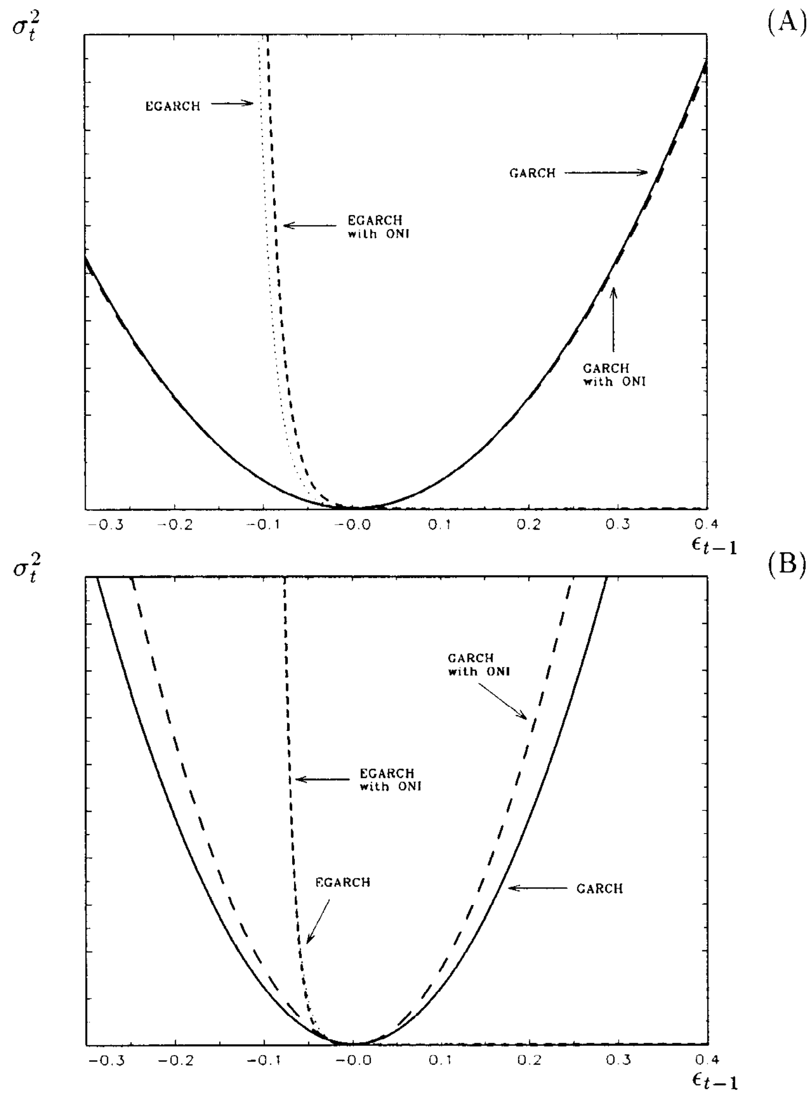

Figure 4
Nasdaq Composite Index and Russell 2000 Index, respectively. News impact curves of GARCH(1,1), EGARCH(1,1), GARCH(1,1) with ONI and EGARCH(1,1) with ONI.

considering the total effect of a shock ϵ_t, $\frac{\alpha}{1-\beta}$, and determining the rate of the decay to the half of the effect of ϵ_t as $-\frac{ln(2)}{\ln(\beta)} - 1$.

We will focus on the sum $\alpha + \beta$ as a direct measure of persistence of the conditional variance estimate; by the same token, the coefficient β in Equation 4 plays the same role, this time with the asymmetric effects being taken into account.

Weiss (1985) derives formal conditions for the conditional volatility to converge to a finite value in a general GARCH-type model, where he considers past innovations, lagged dependent variable, exogenous variables, and a forecast of the dependent variable. In his framework, he derives a measure of the speed of convergence to this finite value, which can therefore be taken as a measure of persistence of shocks.

The model we adopted here is much simpler and satisfies Weiss's conditions concerning the existence of the unconditional expectation of the variance equation and the convergence of the process to a finite value.

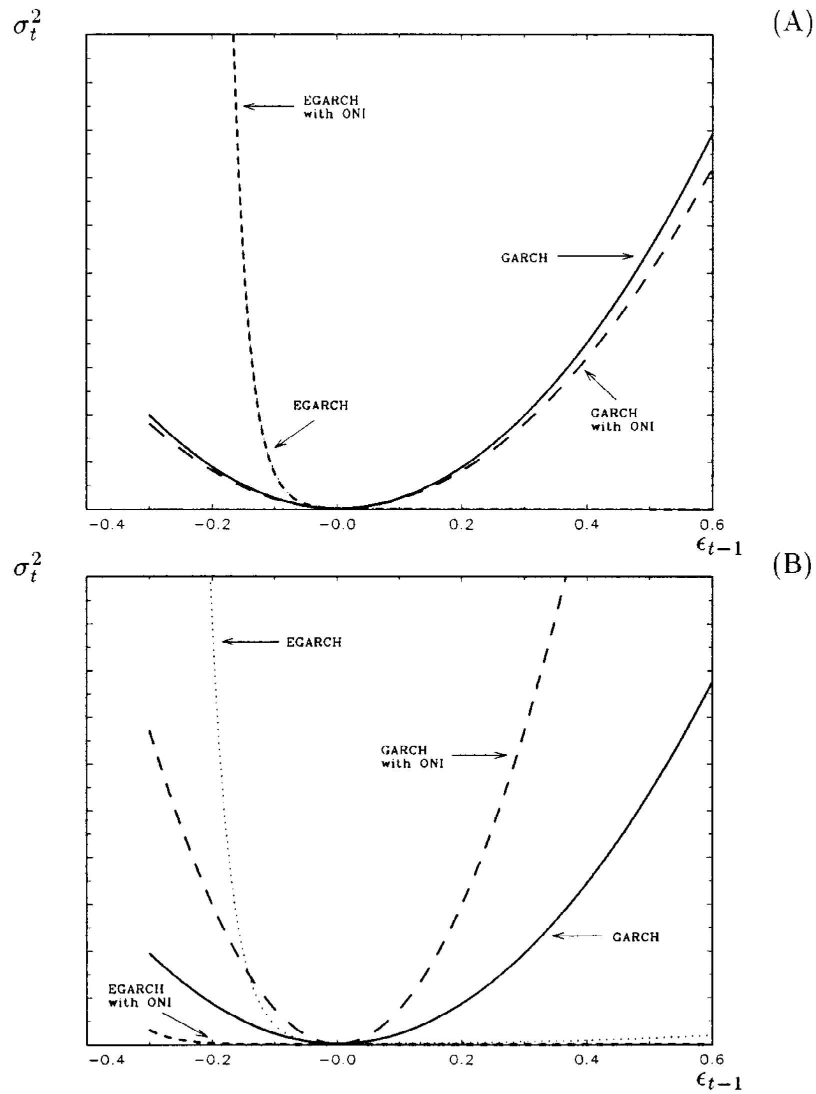

Figure 5
S&P 500 Stock Index and Amex Major Market Index, respectively. News impact curves of GARCH(1,1), EGARCH(1,1), GARCH(1,1) with ONI and EGARCH(1,1) with ONI.

In our case, in fact, it is enough to have $\omega + \psi E(\text{ONI}) > 0$, and $\alpha + \beta < 1$ in order to obtain finite-constant variance for ϵ_t in Equation 3.

The interest in estimating this persistence accurately revolves around the forecastability of volatility: when the empirical persistence is close to one (a common occurrence in practice), future values of the conditional variance tend to be similar to the most recently estimated value. Since the (E)GARCH variance reacts to the current news, an abnormal return tends to generate high forecasts of the variance for a long period, whereas, historically, the news are absorbed much faster.

Among the attempts to account for the high estimated persistence, a few authors have investigated alternatives to the standard GARCH-type model.

Based on the idea of the existence of periods characterized by low or high volatility, Hamilton and Susmel (1994) propose to classify sizeable innovations not as exceptionally high realizations of a single process, but as being generated by a high-volatility process ruled by a time-varying variance which is a multiple of the

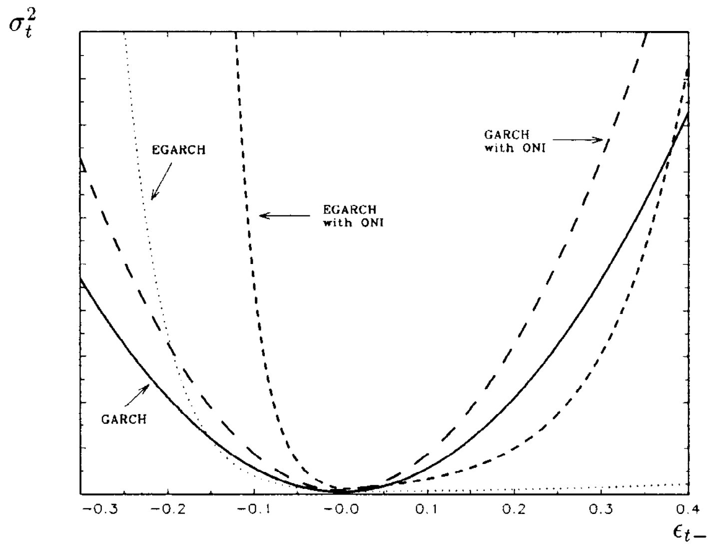

Figure 6
Phlx Kbw Bank Index. News impact curves of GARCH(1,1), EGARCH(1,1), GARCH(1,1) with ONI and EGARCH(1,1) with ONI.

variance in the low-volatility regime. This Markov-switching ARCH model is capable of reducing the volatility persistence, but still lacks an investigation of its consequences for market efficiency and a theoretically sound explanation for why some innovations are so important as to give rise to the switch between regimes.

Within a framework of GARCH-type models with exogenous variables, Lamont, Lumsdaine, and Jones (1996) study the effect of news released at specific days in the year by allowing dummies in both the mean equation and the conditional volatility equation to detect whether a day of "news" significantly alters the behavior of the volatility. Although their answer is positive (in that the coefficients of the step function for announcement days are all significant), the estimated persistence is still high.

Lamoreaux and Lastrapes (1990a) obtain a remarkable reduction in persistence within a GARCH(1,1) model with traded volume as the exogenous variable, but it is not clear, as the authors themselves admit, whether their results are free from a simultaneity problem between returns and volume that would impair the capability of volume to act as a mixing variable for the total movements of stock prices.

In this context, our contribution is to examine the impact of the ONI on the estimated volatility, and to compare the reduction achieved with what can be obtained with the insertion of traded volume, where available. For the three cases (simple, with ONI, and with volume), we summarize persistence results in Table 6. For ease of reference, we reported the percentage change in the estimated persistence relative to the simple case (first two columns). While a reduction (sometimes very substantial) is achieved for all indices (with the exception of the Standard & Poor 500 Index), it is remarkable that it is achieved even for some indices (the Russell 2000 and the AMEX Major Market) for which the ONI variable showed infrequent movements (cf., Figure 2). This seems to be in line with the explanation proposed by Lamoureux and Lastrapes (1990b) that a few events might be responsible for the high persistence of estimated conditional variance. A comparison with the results obtained with traded volume as an exogenous variable shows that the order of reduction is very similar overall.

7 Forecasting Evaluation

The overall evaluation of the importance of the variable ONI in explaining volatility behavior will be complete by performing a one-step-ahead forecasting exercise carried out on the parameter estimates for the sample

Table 6
Estimated Persistence and Percent Reduction[e]

	GARCH	EGARCH	With ONI GARCH	With ONI EGARCH	With Volume GARCH	With Volume EGARCH
DJIA	0.927	0.852	−59.6	−10.4	−19.2	−97.1
NASDAQ	0.893	0.830	−17.7	−10.1	−20.2	−43.6
Russell 2000	0.717	0.650	−13.2	−12.6	—	—
S&P 500	0.953	0.898	0.94	−0.33	—	—
AMEX	0.946	0.896	−20.8	−99.7	—	—
Phlx KBW Bank	0.843	0.885	−72.7	−45.1	—	—

[e]The columns for ONI and Volume report the percentage reduction (if negative) or increase (if positive) in estimated persistence relative to the standard estimation.

period considered, and using the following month (January 1997) as the out-of-sample forecasting period with a horizon of 23 periods.

The most appropriate way to compare forecasting performances is currently under debate in the literature. In a recent paper, Diebold and Mariano (1995) have suggested that the sheer comparison of mean-squared prediction errors (MSPE) as such is flawed by the lack of an indication of the statistical significance of such differences. Moreover, considering just the squares of the errors is limiting when asymmetry in the direction of the forecast error makes a difference (overprediction may have different consequences than underprediction, especially in finance). The possibility that the forecast errors are not normally distributed or are serially correlated increases the difficulty of deriving a test for comparing different forecasts.

Diebold and Mariano (1995) suggest a way to overcome some of the problems arising with other tests in the literature, by considering a generic function $g(\cdot)$ of the forecast errors from two competing models i and j; namely, $g(e_{i,t})$ and $g(e_{j,t})$ for a horizon $t = 1, \ldots, T$. The comparison is a simple difference between the two, that is, $d_t \equiv g(e_{i,t}) - g(e_{i,t})$. They show that:

$$\sqrt{T}(\bar{d} - \mu) \to \mathcal{N}(0, 2\pi f_d(0)),$$

where $\bar{d} \equiv \frac{1}{T}\sum_{t=1}^{T} d_t$ and $f_d(0)$ is the spectral density of $\bar{d}$ evaluated at frequency 0. They propose a test statistic S_1:

$$S_1 = \frac{\widehat{\bar{d}}}{\sqrt{\frac{\widehat{2\pi f_d(0)}}{T}}} \overset{a}{\sim} \mathcal{N}(0, 1).$$

In the present context, we wish to assess the improvement (if any) provided by ONI to the volatility forecast with the two parametric models (GARCH and EGARCH) considered here. We will adopt the GARCH results as the benchmark against which the comparison is performed. Since we are interested in detecting whether ONI manages to reduce persistence in predicted volatility, we will suggest, in the spirit of Granger and Pesaran (1996) or Christoffersen and Diebold (1996), a simple nonlinear function $g(\cdot)$ that penalizes overprediction more than underprediction, namely:

$$g(e_{i,t}) = \begin{cases} |e_{i,t}| & \text{if } e_{i,t} \geq 0 \\ e_{i,t}^2 & \text{if } e_{i,t} < 0 \end{cases}$$

(recall that forecast errors are smaller than 1 in modulus in our case).[6]

[6]Note that this function is different from what was considered in the illustrative example of Diebold and Mariano, who use in their paper symmetric functions such as the absolute value and the square of forecast errors, therefore testing for the significance of the difference between mean-absolute errors and between mean-squared errors.

Table 7
Diebold and Mariano Test for Volatility Forecast Performance[f]

	GARCH with ONI	EGARCH	EGARCH with ONI	GARCH with ONI	EGARCH	EGARCH with ONI
	Dow Jones Industrial Average			NASDAQ Composite Index		
Test Statistic	**−1.733**	*17.690*	1.379	−1.174	*24.892*	1.369
p-values	0.042	0.000	0.084	0.120	0.000	0.086
	Russell 2000 Index			Standard & Poor 500		
Test Statistic	−1.014	−1.293	**−2.309**	**−2.319**	**−6.602**	**−3.439**
p-values	0.155	0.098	0.010	0.010	0.000	0.000
	Amex Major Market Index			Philadelphia KBW Bank Sector Index		
Test Statistic	−1.445	−1.365	**−3.704**	−1.307	*8.212*	*4.498*
p-values	0.074	0.086	0.000	0.096	0.000	0.000

[f]**Boldface** type indicates model significantly better than GARCH(1,1); *italic* type indicates model significantly worse than GARCH(1,1).

The results are arranged in Table 7, where the sign of the test statistics signals a better forecast of the model in the column relative to GARCH if it is negative, worse otherwise.[7] Therefore, we are considering two different one-sided alternative hypotheses to detect whether the model in each column of this table performed significantly better (as shown in boldface type) or significantly worse (shown in italics) than GARCH at the 5% significance level.

The remarkable result is that volatility forecasts improve (according to the criterion chosen) when the ONI variable is inserted in the model, with the following qualifications.

1. When ONI is inserted into the GARCH model, it always provides an improvement in the forecasting performance (significantly so for the Dow Jones Industrial Average and the Standard & Poor 500 Index).

2. EGARCH as such is not better than GARCH across all indices: indeed, in three out of six indices (Dow Jones Industrial Average, NASDAQ Composite Index, and Philadelphia KBW Bank Sector Index) it is significantly worse.

3. When ONI is inserted in the EGARCH model, it confirms the good performance for the Standard & Poor 500, it improves the situation for the Dow Jones Industrial Average and NASDAQ Composite Index, which are not significantly worse than the GARCH model, and provides better forecasts for the Russell 2000 and the AMEX Major Market indices.

8 Conclusions

Options on indices are actively traded on the markets, and for these indices an accurate estimate of volatility is of interest. We have examined five of these indices and the Dow Jones Industrial Average as a benchmark, focusing on the excessive estimated persistence in the conditional volatility, which determines a very slow absorption of sizeable innovations relative to what is observed in practice.

We have concentrated on a measure of market opening news obtained by decomposing the innovation from the closing price of one day to the closing price of the following day as the sum of two innovations that are not identically distributed, as in Clark's (1973) model. We have provided evidence to show that what we call an overnight surprise (the difference between the opening price of one day and the closing price of the previous day) has a significant impact on conditional volatility (the only exception being the Standard & Poor 500 Index). A nonparametric estimation of its distribution provides a useful tool in determining a confidence interval for overnight surprises which could be used ex ante to evaluate an expected range for future volatility.

[7]The *p*-values are inserted for the asymptotic distribution with a word of warning needed to remind the reader that the simulations of Diebold and Mariano showed a small sample (slight) tendency to over-reject the null hypothesis of no difference.

The results are similar to those obtained by Lamoureux and Lastrapes (1990a), which rested on the crucial assumption that the volume is weakly exogenous with respect to the returns. In practice, data on the volume exchanged on the indices are seldom available, and this makes it impossible to use this variable as a proxy for number of trades.

By means of our overnight-surprise indicator (which is always available), we achieve a modification in the news-impact curves for the indices and a reduction in persistence. The forecasting performance of the GARCH and EGARCH models is always improved upon when the ONI variable is inserted into the model. In most cases, the improvement is statistically significant when we compare forecasts using Diebold and Mariano's (1995) test.

References

Andersen, T. (1995). "Return volatility and trading volume: an information flow interpretation of stochastic volatility." Working Paper 170, National Bureau of Economic Research, Cambridge, MA.

Berlin, H. M. (1990). *The Handbook of Financial Market Indexes, Averages, and Indicators*. Homewood, IL: Dow Jones-Irwin.

Berry, T. D., and K. M. Howe (1994). "Public information arrival." *The Journal of Finance*, 4:1331–1346.

Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). "ARCH models." In R. F. Engle and D. L. McFadden (eds.), *Handbook of Econometrics*, vol. 4. Amsterdam: North-Holland, ch. 49, pp. 2961–3038.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

Christoffersen, P. F., and F. Diebold (1996). "Optimal prediction under asymmetric loss." Unpublished Manuscript, Department of Economics, University of Pennsylvania.

Clark, P. (1973). "A subordinated stochastic process model with finite variance for speculative prices." *Econometrica*, 41:135–156.

Diebold, F., and R. S. Mariano (1995). "Comparing predictive accuracy." *Journal of Business and Economic Statistics*, 3:253–263.

Ding, Z., C. W. J. Granger, and R. F. Engle (1993). "A long memory property of stock market returns and a new model." *Journal of Empirical Finance*, 1:83–106.

Engle, R. F. (1982). "Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation." *Econometrica*, 55:987–1008.

Engle, R. F., and V. K. Ng (1993). "Measuring and testing the impact of news on volatility." *Journal of Finance*, 48:1749–1778.

Gallo, G. M., and B. Pacini (1997). "The effect of trading activity on market volatility." In *Proceedings of the BNP—Imperial College Conference on Forecasting Financial Markets*. London: Banque Nationale de Paris and Imperial College.

Granger, C. W. J., and M. H. Pesaran (1996). "A decision-theoretic approach to forecast evaluation." Unpublished Manuscript, Trinity College, Cambridge.

Hamilton, J. D., and R. Susmel (1994). "Autoregressive conditional heteroskedasticity and changes in regime." *Journal of Econometrics*, 64:307–333.

Kim, C. J. (1993). "Unobserved-component time series models with Markov-switching heteroskedasticity: Changes in regime and the link between inflation rates and inflation uncertainty." *Journal of Business and Economic Statistics*, 11:341–349.

Lamont, O., R. Lumsdaine, and C. M. Jones (1996). "Public information and the persistence of bond market volatility." Working Paper 5446, National Bureau of Economic Research, Cambridge, MA.

Lamoureux, C. G., and W. D. Lastrapes (1990a). "Heteroskedasticity in stock return data: Volume versus GARCH effects." *The Journal of Finance*, 45:221–229.

Lamoureux, C. G., and W. D. Lastrapes (1990b). "Persistence in variance, structural change, and the GARCH model." *Journal of Business and Economic Statistics*, 8:225–234.

Nelson, D. B. (1991). "Conditional heteroskedasticity in asset returns: A new approach." *Econometrica*, 59:347–370.

Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. New York: Chapman and Hall.

Tauchen, G., and M. Pitts (1983). "The price variability-volume relationship on speculative markets." *Econometrica*, 51:485–505.

Weiss, A. A. (1985). "Stability of heteroskedastic process." *Journal of Time Series Analysis*, 4:303–310.

Editorial Policy

The *SNDE* is formed in recognition that advances in statistics and dynamical systems theory may increase our understanding of economic and financial markets. The journal will seek both theoretical and applied papers that characterize and motivate nonlinear phenomena. Researchers will be encouraged to assist replication of empirical results by providing copies of data and programs online. Algorithms and rapid communications will also be published.

ISSN 1081-1826