ELSEVIER Journal of Banking & Finance 20 (1996) 1447–1461

Journal of BANKING & FINANCE

The impact of firm specific news on implied volatilities

Monique W.M. Donders [a], Ton C.F. Vorst [b,*]

[a] Institute for Research and Investment Services, Coolsingel 120, 3011 AG Rotterdam, The Netherlands
[b] Department of Finance and Erasmus Centre for Financial Research, Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Received 10 February 1995; accepted 1 February 1996

Abstract

We study the implied volatility behavior of call options around scheduled news announcement days. Implied volatilities increase significantly during the pre-event period and reach a maximum on the eve of the news announcement. After the news release, implied volatility drops sharply and gradually moves back to its long-run level. Only on the event date are movements in the price of the underlying significantly larger than expected. These results confirm the theoretical results of Merton (1973).

JEL classification: G13; G14

Keywords: Event study; Implied volatility

1. Introduction

This paper examines the behavior of implied volatilities of call option prices around scheduled news announcements. In an efficient securities market we expect

* Corresponding author. Tel.: (31) 10 408 1270; fax: (31) 10 452 7746; e-mail: vorst@opres.few.eur.nl.

PII S0378-4266(96)00011-8

stock prices to adjust to new information very quickly. Patell and Wolfson (1984) find an initial price reaction to earnings and dividend announcements within a few minutes on the NYSE, but disturbances in the variance persist for several hours. Ederington and Lee (1993) examine the impact of scheduled macroeconomic news on interest rate and foreign exchange futures. They find that volatility is substantially higher than normal for roughly fifteen minutes after the news release and slightly elevated for several hours.

There are different ways to model a volatility process that incorporates large movements in the asset price. The most well-known models describing changing volatilities over time are the autoregressive conditional heteroskedasticity (ARCH) models, originally introduced by Engle (1982) and extended to Generalized ARCH (GARCH) models by Bollerslev (1986) and Exponential GARCH (EGARCH) models by Nelson (1990). Bollerslev et al. (1992) provide an excellent survey on the theory and evidence of these models for financial time series. One important characteristic of this family of models, however, is that a period of high volatility is assumed to *follow* a large movement in the price of the asset, while in an efficient market we expect uncertainty to decrease after new information is released. Although predictable volatility models like the GARCH family are useful in valuing derivative instruments (see e.g. Kuwahara and Marsh, 1993), it is not clear that the implications on option values of these one shot increases in volatility due to scheduled news announcements are covered by the standard forms of these models. However, we show that GARCH models with an additional structural break in volatility might better describe these implications. [1]

The shock in the price of the underlying stock on the earnings announcement day might be better described by a jump diffusion process. Merton (1976) and Amin (1993) consider the valuation of options under this kind of process. However, in a jump diffusion process the exact time of a jump in the stock price is unknown, while scheduled news announcement dates are, by definition, known in advance.

Market's expectations of future volatility are reflected in the implied volatilities of option prices. According to Merton (1973), the implied volatility of a European option is equal to the average volatility over the remaining life of the option if volatility is a deterministic function of time. Heynen et al. (1994) show that the same result approximately holds for at-the-money options if volatility is stochastic or follows an (E)GARCH process. For these models, the implied volatility is equal to the average expected volatility of the asset over the remaining life of the option.

If investors sense more uncertainty about the stock price on the eve of a news release than on other days, the average expected volatility, and hence the implied

[1] This approach was suggested by an anonymous referee.

volatility, should increase during the pre-announcement period. After the stock price adjusts to the new information, volatility will drop to its normal level.

We use the event study methodology to study the influence of scheduled news announcements on implied volatilities. We find that implied volatility rises during the pre-announcement period, reaches a maximum just before the news release and drops sharply afterward. This result holds both for the raw data as well as for implied volatilities that are adjusted for general market movements in volatility. Furthermore, we examine the behavior of the variance of returns of the underlying assets and conclude that, except for the event day itself, this volatility is not significantly different from its normal level in the period surrounding an announcement. This means that the increase in implied volatility cannot be explained by higher variance of stock returns before the news announcement.

We find one important difference between the theoretical and empirical patterns in implied volatilities: the implied volatility ten days before an announcement is lower than expected. We test a trading strategy using call options and the underlying asset that may profit from this 'overreaction' in implied volatilities and find that, when taking transactions costs into account, it does not yield significant returns. The market seems to be efficient in the sense that large shocks in the stock price coincide with large reductions in future volatility.

This paper is organized as follows. Section 2 presents a model on the behavior of implied volatilities around scheduled news announcements. Section 3 discusses the test methodology and the data. Section 4 contains the empirical analysis and discusses the empirical results. Section 5 concludes.

2. Theory

In this paper we concentrate on 'scheduled news', for which the disclosure date is known in advance, but the information content is not. In the Black and Scholes (1973) pricing model for European type options, it is assumed that the volatility of the underlying stock is constant over time. However, if volatility is a deterministic function of time, Merton (1973) shows that the Black–Scholes formula still holds if we replace the volatility by the average volatility until expiration. Daily stock price returns are random variables that might be independently and identically distributed on normal days. During scheduled news announcement days, however, a higher volatility is expected. If volatility on a normal day is σ^2_{normal} and on unexpected news announcement day is equal to σ^2_{high}, then average volatility AV_x over the remaining life of the option if the announcement has not occurred yet is defined as

$$AV_x = \sqrt{\frac{(x-1)}{x}\,\sigma^2_{\text{normal}} + \frac{1}{x}\,\sigma^2_{\text{high}}} \tag{1}$$

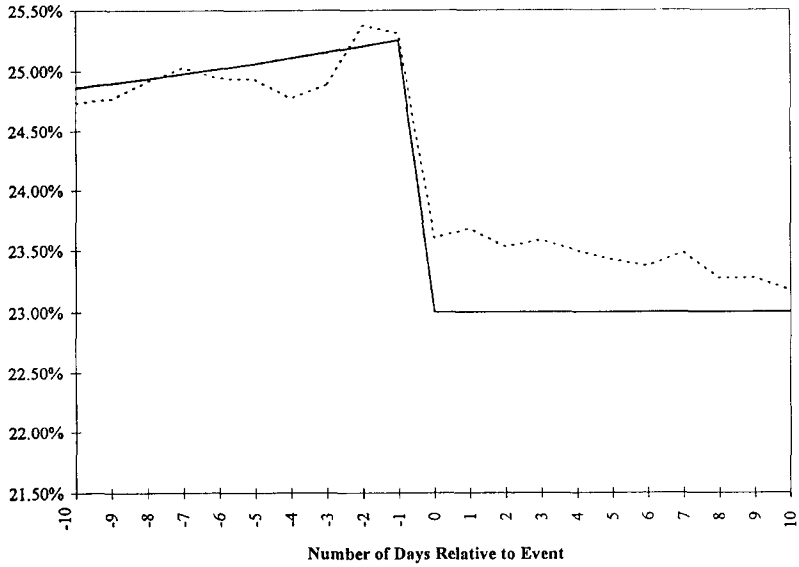

Fig. 1. Average volatility and GARCH volatility around announcement days (– average volatility; ... GARCH volatility).

Parameter values for the simulated GARCH process are $\alpha_0 = 0.000015$, $\alpha_1 = 0.8$, $\alpha_2 = 0.1$, $\kappa = 0.5$, $r = 0.05$, $\lambda = 0.05$ and $\beta_u = 5/4$, $\beta_d = 5/9$. Simulation for asset prices starts at 110 days before the event day at a fixed asset price and the stationary volatility, which is 0.231869 for these parameter values. At-the-money options with maturity date 40 days after the event have been used. 10,000 Monte Carlo simulations are used to compute the GARCH option prices and from these prices implied volatilities are calculated for the Black-Scholes formula. This simulation is repeated 50 times to compute the average implied volatilities around the event day.

where x is the number of days until the expiration date of the option. After the news announcement day, the average volatility drops to σ^2_{normal} (assuming there are no other scheduled information releases before expiration). With this simple model the implied volatility, as a function of the time until and after the expected news announcement, can be described by the function depicted by the solid line in Fig. 1.

It is well known that volatilities are not constant over time and the assumption that the volatility is constant except for scheduled news announcement days does not seem very realistic. However, Heynen et al. (1994) show that for the option pricing model of Hull and White (1987) which is based on a stochastic volatility process, the implied volatility in option prices is approximately equal to average expected volatility until the expiration date. They also show that the same result holds under the assumption that stock prices follow a GARCH or EGARCH process and option prices are based on the model developed by Duan (1995). If the volatility during non-announcement days can be described by for example a

GARCH model and if the peaks in volatility on scheduled news announcement days are large compared to the fluctuations due to the GARCH specification, we can use the following extension, which exhibits a structural break in volatility in the news announcement date:

$$\ln\left(\frac{X_{t+1}}{X_t}\right) = r + \lambda\sigma_{t+1} - \frac{1}{2}\sigma_{t+1}^2 + \sigma_{t+1}\epsilon_{t+1}, \tag{2}$$

$$\sigma_{t+1}^2 = \left[\alpha_0 + \alpha_1\sigma_t^2 + \alpha_2\sigma_t^2(\epsilon_t - \kappa)^2\right]\left[1 + \chi_{t+1=t^*}\beta_u - \chi_{t+1=t^*+1}\beta_d\right]^2 \tag{3}$$

with X_t the underlying stock price, r the riskless interest rate, λ the market price of risk, σ_t the volatility, ϵ_{t+1}, conditional on the time t information, a standard normal random variable, t^* the event date and χ an indicator function. α_0, α_1, α_2, β_u and β_d are constants. If we set $\beta_u = \beta_d = 0$ in Eq. (3), the non-linear asymmetrical GARCH model of Engle and Ng (1993) results. The parameters β_u and β_d determine the magnitude of the increase and decrease in volatility around the event day.

Invoking the local risk-neutralization principle of Duan (1995) theoretical option prices and hence implied volatilities can be calculated with respect to the risk-neutral probability measure.

To illustrate the properties of this model, the behavior of implied volatility around event days for this model is depicted by the dotted line in Fig. 1 for a certain set of parameter values. It is clear that the extended GARCH model and our simple average volatility model yield basically the same patterns. Although there are some differences, both lines basically reveal the same pattern in implied volatilities.

The solid line in Fig. 1 describes the hypothesis about implied volatilities that is tested in this paper. In addition to the pattern in implied volatilities we also test the assumption that the underlying asset has a significantly higher volatility only on the scheduled news announcement day and not in the pre- and post-event periods. Since changes in returns and (implied) volatilities of individual stocks might not only result from firm specific circumstances, but can also be caused by general market trends, we correct for these market wide changes in (implied) volatility.

3. Methodology

3.1. Data

3.1.1. Announcements

The announcement sample contains 96 scheduled news disclosures released by 23 firms during the period June 1991–December 1992. During this period, stocks of these firms were listed on the Amsterdam Stock Exchange (ASE) and options on the stocks were listed on the European Options Exchange (EOE).

We used Beursplein 5 (a weekly information bulletin of the ASE) to identify scheduled news announcement dates. Using the Financieele Dagblad (FD), we checked whether on these dates a news item on the particular event did indeed occur. Of the initial 143 announcements, 47 did not appear in FD. [2] The average number of days between the official publication of the announcement day and the actual information release is 78.5 days. Since the announcements frequently concern quarterly- (28), semi-annual- (43) or annual (22) earnings disclosures, and since most listed firms have a fiscal year ending in December, the announcements are clustered in time.

3.1.2. Implied volatilities

Daily records for each traded call option were obtained from the EOE for the period June 18, 1991 through December 30, 1992. We only include records that satisfy the following criteria:

1. the option has at least ten days to maturity;
2. the option has a bid- and an ask price that are both larger than zero
3. the average of the bid- and the ask price is greater than or equal to $f\,0.20$

We impose these criteria for the following reasons. De Jong et al. (1992) find a slight reduction in volatility around expiration days. To avoid this effect, we eliminate very short maturity calls. The minimum tick size on the EOE is $f\,0.10$. Because the influence of rounding errors in option prices smaller than $f\,0.20$ is too large to make accurate estimates of implied volatilities, we eliminate these. To avoid the effect of bid–ask bouncing on consecutive days, we use the average of the bid- and the ask price as an estimate for the 'true' price of an option. Therefore, both prices must be larger than zero. For every day and each option class we follow Beckers (1981) by selecting only the option with the shortest time to maturity and the smallest absolute moneyness to calculate the implied volatility. We define moneyness as the ratio of the stock price net of the present value of dividends to be paid during the life of the option to the present value of the exercise price minus 1. To incorporate the early exercise premium due to dividend payments we apply the Roll-Geske-Whaley formula (Hull, 1993) for American options on dividend paying stocks. This formula is valid because our sample does not contain any options with more than one dividend payment during the remaining life of the option. We use a Newton Raphson iterative search to calculate the implied volatilities.

[2] Additionally, there were 45 news items in Financieele Dagblad that did not appear in Beursplein 5 as announcements. These news items most frequently concern annual stockholder meetings and publication of annual reports. Since we cannot be sure that these news items are 'scheduled news', we do not include these in our sample.

The interest rate used is the Amsterdam Interbank Offered Rate with maturity closest to the maturity of the option. These AIBOR rates are obtained from Datastream. Dividend data are collected from EOE publications, the Effectengids and the Officiële Prijscourant. We use actual dividends instead of market expectations. Since we only use options with a relative short time to maturity, we assume that there is no uncertainty regarding the amount of dividend paid. The recorded stock price is the price from the last transaction in that stock on the ASE.

We control for market wide changes in volatility in two different ways. First, for every implied volatility, we subtract the implied volatility of the EOE-index option with the same time to maturity as the stock option. This index consists of 25 stocks, including the 23 stocks studied in this paper. A second way to correct for market-wide changes in volatility, following Sheikh (1989), is using control stocks. For every trading day betas for all stocks are estimated by an OLS regression of the foregoing 120 daily total returns on the returns of the EOE-index. Five groups of stocks [3] of comparable betas and standard deviations are formed. For every control group we calculate the implied volatility on each day as being the average of the implied volatilities of stocks that do not have an announcement in the ten days preceding or following this day. The control groups are formed in such a way that for every day there is at least one such stock for every control group.

3.2. Tests

We are interested in comparing cross-sectional volatilities and implied volatilities during the periods around the scheduled news announcements to the volatilities and implied volatilities during normal periods. To see whether these differences are significant, a number of tests are used that are described below and are based on sctn. 16.5 of Davidson and MacKinnon (1993). Let

$$r_{jt} = \ln S_{jt} - \ln S_{jt-1} \tag{4}$$

be the logarithmic return for stock j during day t and define the extra return R_{jt} as

$$R_{jt} = r_{jt} - \frac{1}{T}\sum_{t=1}^{T} r_{jt} \tag{5}$$

[3] As discussed before the events are strongly clustered in time. Because pre-event and post-event periods of the announcement stock and the control stock must be non-overlapping, in many cases it is not possible to match an announcement stock with a single control stock.

where T is the length of the total observation period. Let $V_{jt} = \sqrt{R_{jt}^2}$ be the so-called 'volatility on day t'. For each event we distinguish four periods:

I_{j1}: the period from T_1 days to 1 day before the event date – the pre-event period;
I_{j2}: the event-day;
I_{j3}: the period from 1 day to T_3 days after the event – the post-event period;
I_{j4}: the control period, all days not included in an I_{j1}, I_{j2} or I_{j3} period, T_4 days

We define the excess volatility of stock j on day t as

$$EV_{jt} = V_{jt} - \frac{1}{T_4} \sum_{t \in I_{j4}} V_{jt} \tag{6}$$

and dummy-variables x_1, x_2, and x_3:

$$x_{j1t} = \begin{cases} 1 & \text{if } t \in I_{j1} \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

$$x_{j2t} = \begin{cases} 1 & \text{if } t \in I_{j2} \\ 0 & \text{otherwise} \end{cases}$$

$$x_{j3t} = \begin{cases} 1 & \text{if } t \in I_{j3} \\ 0 & \text{otherwise} \end{cases}$$

To test for heteroskedasticity we use a GMM regression with Newey–West standard errors to estimate

$$EV_{jt} = c + \alpha_1 x_{j1t} + \alpha_2 x_{j2t} + \alpha_3 x_{j3t} + \epsilon_{jt} \ \forall\, j, t. \tag{8}$$

The t-values of the coefficients α_1, α_2 and α_3 indicate whether the volatilities in the three sub-periods are statistically different from those in the control period.

We use the same regression equation to test for heteroskedasticity in implied volatilities by replacing V_{jt} with the implied volatilities. EIV_{jt} is the excess implied volatility.

4. Results

4.1. Stock return volatilities

For the simple model of implied volatilities in (1) to hold, the standard deviation of stock returns should only be different from their control-period level on the event day and not in the pre- and post-event period. To test this hypothesis,

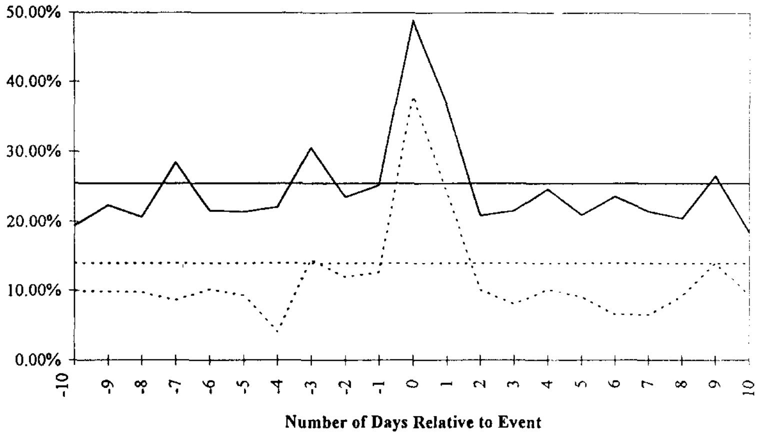

Fig. 2. Stock return volatilities around event days (– raw return volatilities; ... EOE-index adjusted return volatilities).

we apply regression Eq (8) to both raw stock return volatilities (Eq. (9a)) and EOE-index adjusted returns (Eq. (9b)), t-values are in brackets:

$$EV_{jt} = 0.00015 + \underset{(2.62)}{0.00108}x_{j1t} + \underset{(10.16)}{0.01256}x_{j2t} + \underset{(1.42)}{0.00058}x_{j3t} + \epsilon_{jt}, \tag{9a}$$

$$EV_{jt} = 0.00012 + \underset{(1.28)}{0.00046}x_{j1t} + \underset{(12.11)}{0.01312}x_{j2t} + \underset{(1.98)}{0.00071}x_{j3t} + \epsilon_{jt}. \tag{9b}$$

The pre-event and post-event periods consist of ten days.

Although the coefficient for x_{j1t} in Eq. (9a) is slightly significant, it is of a different order of magnitude than the coefficient for x_{j2t}. Both Eq. (9b) and Fig. 2 show that the (EOE-index adjusted) volatility of stock returns during the pre-event period and the post-event period does not differ significantly from the volatility during periods without scheduled news announcements. It is clear, however, that the volatility on the event day differs significantly from the volatility during other periods. Results for control group adjusted volatilities do not differ significantly from those for EOE-index adjusted volatilities and are available from the author upon request. We conclude that the underlying assumptions of the models as described in Section 2 seem to hold.

4.2. Implied volatilities

Fig. 3 gives the (EOE-index adjusted) average implied volatilities (IV) in the period from ten days before the event until ten days after the event. On each day this average is calculated as the sum of the implied volatilities of the individual announcements divided by the number of announcements. The horizontal solid lines in the graphs give the long-run level of the implied volatility. This long-run

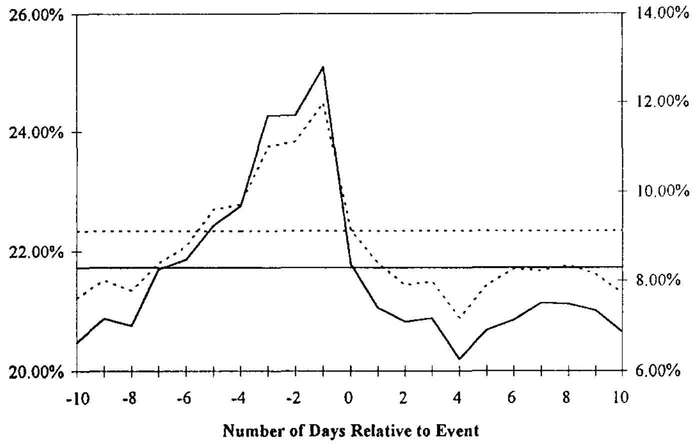

Fig. 3. Implied volatilities around event days (– raw implied volatility (left axis); ... EOE-index adjusted implied volatility (right axis))

level of the implied volatility is the weighted average of implied volatilities of the 23 stocks over the control period. Each stock is weighted with the number of events for that stock.

Both graphs peak just before the scheduled news announcement date. After the announcement, the implied volatility is sharply lower.

The decline in implied volatilities from approximately 25% to 22% means a price drop of 13% for a typical option in our sample. For implied volatilities, we slightly adjust Eq. (8). Instead of one dummy variable for the pre-event period, we introduce a trend variable $x_{j1t}^* = (T_l + 1 + t)\ \forall t \in [-T_l, \dots -1]$, which increases as the number of days till the event declines. The regression equation is given by (t-values in brackets):

$$EIV_{jt} = -0.00405 + \underset{(6.03)}{0.00509} x_{j1t}^* + \underset{(1.30)}{0.00987} x_{j2t} + \underset{(0.30)}{0.00079} x_{j3t} + \epsilon_{jt}, \quad (10a)$$

$$EIV_{jt} = -0.00379 + \underset{(6.25)}{0.00468} x_{j1t}^* + \underset{(1.32)}{0.00889} x_{j2t} - \underset{(-0.44)}{0.00100} x_{j3t} + \epsilon_{jt}. \quad (10b)$$

Note to Table 1:

[a] The changes are for 95 announcements of 'scheduled' news over the period June 18, 1991 through December 30, 1992. In Panel A we use IVs calculated from call option prices. In Panels B and C we subtract the IVs of the EOE-index and the average implied volatility of the group of control stocks respectively. 'Days' is the number of days in the pre-event and post-event period. N^- (N^+) denotes the number of events for which the average IV in the post-event period was smaller (larger) than the average IV in the pre-event period. 'Median %Ch.' gives the median percentage change between the post-event period IV and pre-event period IV, averaged over all events. 'Events' gives the number of events for which IVs could be calculated for every day in the pre-event and post-event period. 'Wilcoxon' is the Wilcoxon signed-rank statistic for the hypothesis that the sum of the plus ranks equals the sum of the minus ranks.

As in (9a) and (9b), the pre-event and post-event periods consist of 10 days. Eq. (10a) and (10b) show that during the pre-event period the rise in (EOE-index adjusted) implied volatility is highly significant, while in the post event period the implied volatility is not significantly different from that in the control period. Control-group adjusted implieds yield similar results. These results are comparable to those found by Patell and Wolfson (1979).

Table 1
Changes in implied volatilities of option prices around announcement dates. [a]

Days	Events	N^+	N^-	Median %Ch.	Wilcoxon
Panel A. Post-event period average IV as compared to pre-event period average IV					
10	85	23	62	− 6.99	−4.879
9	87	21	66	− 7.67	−5.350
8	87	21	66	− 8.54	−5.896
7	87	19	68	− 9.74	−6.222
6	88	18	70	− 9.77	−6.453
5	88	17	71	−10.45	−6.612
4	89	15	74	− 9.87	−6.499
3	90	13	77	− 9.66	−6.460
2	92	21	71	− 8.23	−5.506
1	92	20	72	− 8.51	−4.984
Panel B. Post-event period average IV as compared to pre-event period average IV, using EOE-index adjusted IVs ($IV_i - IV_{\text{EOE-index}}$)					
10	85	26	59	− 5.94	−3.227
9	87	27	60	−18.55	−3.496
8	87	25	62	−21.16	−4.101
7	87	24	63	−21.71	−4.160
6	88	24	64	−24.65	−4.148
5	88	25	63	−26.15	−4.369
4	89	23	66	−27.89	−4.433
3	90	19	71	−30.68	−4.912
2	92	27	65	−27.75	−3.773
1	92	26	66	−18.80	−2.566
Panel C. Post-event period average IV as compared to pre-event period average IV, using control-group adjusted IVs ($IV_i - IV_{\text{control-group}}$)					
10	85	30	55	−35.74	−2.890
9	87	33	54	−29.14	−2.607
8	87	35	52	−23.71	−1.955
7	87	32	55	−33.32	−2.476
6	88	33	55	−33.29	−2.401
5	88	32	56	−31.21	−1.789
4	89	32	57	−40.02	−2.318
3	90	33	57	−33.82	−1.978
2	92	32	60	−32.23	−2.083
1	92	32	60	−25.17	−2.403

In Panel A of Table 1 the implied volatilities of individual stocks are compared for the pre- and post-event period. For each stock the average implied volatility over a fixed number of days before the event is compared with the average implied volatility of the same number of days, including the event date, after the event. N^- gives the number of events for which the post-event average implied volatility is lower and N^+ gives the number of events for which the post-event average implied volatility is higher. For example, in the period of three days before the event the implied volatility was higher than for the period of three days after the event in 77 out of 90 cases. Also, the median percentage change in implied volatilities for all stocks is given. In Panel B and C the same numbers are given for the index adjusted and control-group adjusted implied volatilities respectively.

From Table 1 and the cross-sectional regression Eqs. (10a) and (10b) we may conclude that Fig. 3 is in many respects similar to the theoretical Fig. 1. The drop in implied volatilities on the event day is significant (the last row of every panel in Table 1). Before the news announcement the implied volatility is significantly higher than the implied volatility in the period after the event.

Fig. 1 and Fig. 3 differ with respect to the first days in the pre-event period. Since the exact dates of the news releases are known well in advance, the empirical implied volatility for these days are lower than theoretically expected. Furthermore, the implied volatility drops below its long-run level in the four days (t-statistic: -1.987) after the event, which is not explained by the model.

4.3. Event-day returns and implied volatilities

To determine whether a large shock in the price of the underlying asset at the event day coincides with a large reduction in uncertainty about future stock price returns, we regress the absolute value of market model residual returns on an event day on the differences in implied volatility on the event day and the pre-event day. The following equation results:

$$|R_{j0}| = 0.0169 - \underset{(-2.001)}{0.0024} DIV_j + \epsilon_j \tag{11}$$

with $|R_{j0}|$ the absolute value of the market model adjusted return for stock j on the event day and $DIV_j = IV_{j,0} - IV_{j,-1}$ the difference between the implied volatility on the event day and the implied volatility on the pre-event day. Since the estimated coefficient is significant, we conclude that large price shocks are correlated with large reductions of uncertainty. Alternatively, the market predicts large shocks by significantly increasing the implied volatility before an announcement.

4.4. Trading strategy

As noticed before Fig. 3 exhibits roughly the same shape as Fig. 1. Hence, the hypothesis in (1) seems to hold. Options have a higher implied volatility before the event, because average expected volatility is higher and hence the replication costs increase. There is one remarkable difference between Fig. 3 and Fig. 1, however. Both raw implied volatilities and adjusted implied volatilities are approximately at their long-run level ten days before and ten days after the event. For ten days after the event, this was expected in Fig. 1. For ten days before the event, however, the average implied volatility is too low as compared to the model in Fig. 1. Although the hypothesis of Section 2 seems to hold, there might also be a secondary effect in the market which can explain the difference with Fig. 1. Just before a scheduled news announcement there are market participants who speculate on a large increase in the stock price. Instead of buying stocks investors buy call options because of the higher leverage of these instruments and the limited down-side risk in case the stock price falls in reaction to the new information. Hence, there is a higher demand for call options. This increases the price and the implied volatility. Thus, the observed implied volatility might be higher than the theoretical implied volatility in Fig. 1.

We consider a trading strategy that profits from this overreaction. Ten days before the event we buy a call option and short the underlying stock to obtain a delta-neutral portfolio. Each trading day the stock position is adjusted to changes in the hedge ratio of the option. Daily trading profits and losses are assumed to yield the AIBOR-rate. The portfolio is liquidated on the day before the event-day when the implied volatility, and with that the call-option value, reach a maximum.

The average return over all events for this strategy is 49.69% (Wilcoxon signed-rank statistic 3.27) of the initial option value. Of course, this strategy involves both implicit and explicit transaction costs. Since explicit transaction costs such as commissions and fees differ dramatically across investors, we only consider implicit transaction costs due to the bid–ask spread. Quoted bid and ask prices for all options are in the dataset. Unfortunately, there are no quoted spreads available for ASE stocks. Interviews with brokers, however, indicate that, on a normal trading day, these spreads vary from f0.10 (the minimum ticksize) for the stocks with the highest liquidity to f0.50 for smaller stocks. Taking the bid–ask spread for both options and stocks into account, we find a return on the delta-neutral position of -2.34% (Wilcoxon signed-rank statistic 1.12). Hence, this trading strategy does not yield economically significant returns, indicating that the European Options Exchange functions efficiently. [4]

5. Conclusion

In this paper the implied volatility of options around scheduled news announcement days of the underlying stock are studied. Implied volatilities increase when

[4] We also tried a similar trading strategy to profit from the drop of the implied volatility below the long run level four days after the event. This strategy was not profitable even without transaction costs.

the event day approaches. After the news announcement, the implied volatility drops sharply for a few days. It even drops below its long-run level. The volatility of the underlying stock is not different from its control period level during both the pre-event and the post-event period.

On the event-date itself, however, volatility seems to be higher. If we use a standard (E)GARCH or a stochastic volatility process to model the behavior of the underlying stocks, we do not detect this one-day increase in volatility. Also, jump diffusion option pricing models are not suited to describe this kind of phenomena, since these models usually assume that the jumps appear at random. However, the adjusted GARCH model described in Eq. (2) and (3) indicates a direction for further research to combine a GARCH model with exogenous shocks in volatility due to earnings announcements.

The simple model for the volatility of the underlying asset around scheduled news announcements produces a time-dependence in implied volatilities that seems to agree with the pattern we find in the options market. However, the market seems to react too strong. This overreaction might be caused by trading by investors who speculate on changes in the price of the underlying stock, causing an excess demand for call options before the event. We find a trading strategy that profit from these overreactions. This strategy yields economically insignificant returns when transaction costs are taken into account, indicating that the European Options Exchange functions efficiently.

Acknowledgements

We acknowledge the helpful comments and suggestions of Werner de Bondt, Frank de Jong, Frank Kleibergen, Teun Kloek, Eva Liljeblom, Michael Rockinger and two anonymous referees as well as participants in the May 1994 Workshop on Empirical Research in Securities Markets at EIASM, the 1994 AFFI, EFA and NFA conferences, the INQUIRE 1994 meeting, the 15th AMEX Options and Dervatives Colloquium and the Fifth Annual Derivatives Conference at Cornell University. We thank the EOE-Optiebeurs for providing the data and Marco Lavooi for excellent research assistance. The usual disclaimer applies.

References

Amin, K., 1993, Jump diffusion option valuation in discrete time, Journal of Finance 48, 1833–1863.

Beckers, S., 1981, Standard deviations in option prices as predictors of future stock price variability, Journal of Banking and Finance 5, 363–382.

Black, F. and M.J. Scholes, 1973, The pricing of options on corporate liabilities, Journal of Political Economy 81, 637–659.

Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics 31, 307–327.

Bollerslev, T., R. Chou and K. Kroner, 1992, ARCH modelling in finance: A review of theory and empirical evidence, Journal of Econometrics 52, 5–59.

Davidson, R. and J.G. MacKinnon, 1993, Estimation and inference in Econometrics (Oxford University Press, New York).

De Jong, F., A.G.Z. Kemna and T. Kloek, 1992, A contribution to event study methodology with an application to the Dutch stock market, Journal of Banking and Finance 16, 11–36

Duan, J.C., 1995, The GARCH option pricing model, Mathematical Finance 5, 13–32.

Ederington, L.H. and J.H. Lee, 1993, How markets process information: News releases and volatility, Journal of Finance 48, 1161–1191.

Engle, R., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K inflation, Econometrica 50, 987–1008.

Engle, R. and V.K. Ng, 1993, Measuring and testing the impact of news on volatility, Journal of Finance 48, 1749–1778.

Heynen, R., A. Kemna and T. Vorst, 1994, Analysis of the term structure of implied volatilities, Journal of Financial and Quantitative Analysis 29, 31–56.

Hull, J., 1993, Options, futures and other derivative securities (Prentice Hall, Englewood Cliffs, NJ).

Hull, J. and A. White, 1987, The pricing of options on assets with stochastic volatilities, Journal of Finance 42, 281–300.

Kuwahara, H. and T. Marsh, 1993, The pricing of Japanese equity warrants, Management Science, forthcoming.

Merton, R.C., 1973, The theory of rational option pricing, Bell Journal of Economics and Management Science 4, 141–183

Merton, R.C., 1976, Option pricing when underlying stock returns are discontinuous, Journal of Financial Economics 3, 125–144

Nelson, D., 1990, Conditional heteroskedasticity in asset returns: A new approach, Econometrica 59, 347–370.

Patell, J.M. and M.A. Wolfson, 1979, Anticipated information release reflected in call option prices, Journal of Accounting and Economics, 107–140.

Patell, J.M. and M.A. Wolfson, 1984, The intraday speed of adjustment of stock prices to earnings and dividend announcements, Journal of Financial Economics 13, 223–252.

Sheikh, A.M., 1989, Stock splits, volatility increases and implied volatilities, Journal of Finance 44, 1361–1372.