A High-Frequency Investigation of the Interaction between Volatility and DAX Returns

Philippe Masset,* Martin Wallmeier†

Abstract

One of the most noticeable stylized facts in finance is that stock index returns are negatively correlated with changes in volatility. The economic rationale for the effect is still controversial. The competing explanations have different implications for the origin of the relationship: Are volatility changes induced by index movements, or inversely, does volatility drive index returns? To differentiate between the alternative hypotheses, we analyze the lead-lag relationship of option implied volatility and index return in Germany based on Granger causality tests and impulse-response functions. Our dataset consists of all transactions in DAX options and futures over the time period from 1995 to 2005. Analyzing returns over 5-minute intervals, we find that the relationship is return-driven in the sense that index returns Granger cause volatility changes. This causal relationship is statistically and economically significant and can be clearly separated from the contemporaneous correlation. The largest part of the implied volatility response occurs immediately, but we also observe a smaller retarded reaction for up to one hour. A volatility feedback effect is not discernible. If it exists, the stock market appears to correctly anticipate its importance for index returns.

JEL classification: G10; G12; G13
Keywords: Implied volatility; Granger causality; leverage effect; feed-back effect; asymmetric volatility.

*Philippe Masset, Chair of Finance and Accounting, University of Fribourg / Switzerland, Bd. de Pérolles 90, CH-1700 Fribourg, philippe.masset@unifr.ch.

†Prof. Dr. Martin Wallmeier, Chair of Finance and Accounting, University of Fribourg / Switzerland, Bd. de Pérolles 90, CH-1700 Fribourg, martin.wallmeier@unifr.ch (corresponding author).

1 Introduction

A well known stylized fact in finance is that stock index returns are negatively correlated with changes in volatility ("volatility returns"). (Black (1976)). The negative relationship is typically more pronounced in falling than in rising markets (Figlewski/Wang (2000)) and is stronger for indices than for individual stocks. The distinctive cross dependence pattern between return and volatility plays an essential role in the development of volatility as an asset class, in volatility modelling and in option pricing. Nevertheless, a fully consistent economic explanation for the effect has not yet been offered (see Bouchaud et al. (2001), Bollerslev/Zhou (2006)).

The first attempt to find an economic rationale for the negative return correlation relies on a corporate finance argument. Black (1976) and Christie (1982), among others, argue that a positive stock return increases the market value of the firm's equity, thereby diminishing its financial leverage ratio. The reduced leverage gear will result in a lower volatility of stock returns. The empirical observations, however, do not support this leverage hypothesis. First, it is not compatible with the observed asymmetry of the effect in falling and rising markets. Second, the leverage hypothesis predicts a stronger relationship on the individual firm level than the index level. This prediction is contrary to what is empirically observed (Bouchaud et al. (2001)).[1] In a US study, Figlewski/Wang (2000) conclude that the negative correlation on the index level is far too strong to be explained by the leverage hypothesis (see also Aydemir et al. (2006)).

The term "leverage effect" is sometimes used in a broader sense for the general hypothesis that the causality runs from stock return to volatility. In this paper, we call such a directional relationship "return-driven". In this terminology, the leverage effect is only one possibility to explain a return-driven negative correlation. Another explanation is that bad news might have different implications for future uncertainty than good news (see, e.g., Glosten et al. (1993) and Chen/Ghysels (2007)). For instance, price drops could induce more extensive portfolio adjustments of risk-averse agents than price increases. Bouchaud et al. (2001) suggest that the apparent return-driven relationship could be due to a retarded effect. In their framework, the scale for price updates does not depend on the instantaneous price but on a moving average of past prices which means that current returns lead subsequent volatility returns.

The hypothesis of a volatility-driven negative relationship is known as the "feedback effect" (see, e.g., Pindyck (1984), French et al. (1987), Campbell/Hentschel (1992)). It rests on the assumption that volatility is related to systematic risk and is therefore relevant for pricing. If new information gives rise to an unanticipated increase in volatility, this will also increase risk-adjusted discount rates. As long as cash flow expectations are not affected, stock prices will go down. However, the empirical evidence on the impact of volatility on expected returns is controversial. Some studies report a positive (French et al. (1987), Campbell/Hentschel (1992), Scruggs (1998), Ghysels et al. (2005), Lundblad (2007), Bae et al. (2007)), others a negative relationship (Campbell (1987), Nelson (1991)). Often, the link was found to be insignificant and unstable over time (Glosten et al. (1993), Turner et al. (1989), Harvey (2001)).

[1] This statement holds although new evidence suggests that the firm level effect might be stronger than previous work has documented (see Ericsson et al. (2007) and Chelley-Steeley/Steeley (2005)).

The return-driven and volatility-driven effects might well coexist. For instance, an initial price change could induce a volatility movement which in its turn amplifies the price change with yet another impulse on volatility (see the model of Bekaert/Wu (2000)). In efficient financial markets, the participants will try to anticipate these reactions. Therefore, the steps will evolve almost simultaneously. This makes it difficult to identify the different stages of the process. The higher the return frequency of the data, the better the chances to gain insight into the origin of the return-volatility correlation (see Bollerslev et al. (2006)).

Most empirical studies published during the last few years use a framework which incorporates return-driven as well as volatility-driven effects. The results are mixed. On the one hand, recent studies by Bollerslev et al. (2006), Giot (2005), and Dufour et al. (2006) report evidence of a return-driven relationship while the feedback effect is found to be negligible.[2] On the other hand, Bekaert/Wu (2000) and Dennis et al. (2006) find support for the volatility feedback argument.

To date, there is hardly any evidence for European countries. Within Europe, the German financial market appears to be particularly interesting for at least two reasons. First, during our sample period, DAX futures and options have represented the highest trading volume among all stock index derivatives in Europe. Thus, high-quality high-frequency transaction data are available, which is of crucial importance for this study. Second, the negative relationship between index and volatility returns has been particularly strong and stable at the German market over the last decade (see Hafner/Wallmeier (2007)).

Our main contribution is to go beyond a correlation analysis by studying causality in intervals sufficiently short to identify lead-lag-relationships. To the best of our knowledge, this is the first study which applies causality tests to the interaction of index return and volatility in high-frequency data. For our analysis it is important to accurately determine the point in time when changes in volatility occur. Therefore, in contrast to previous related work, our volatility measure is the at-the-money (ATM) implied volatility instead of realized volatility.[3] We take every effort to measure volatility returns as precisely as possible taking microstructure frictions into account. Our objectives are: (1) to analyze whether the lead-lag relationship is return-driven or volatility-driven, (2) to quantify the impact of an innovation in return or volatility and (3) to estimate how fast return-driven and feedback effects evolve and to draw conclusions for the information efficiency of the markets involved.

We find that a lead-lag relationship only exists for returns at the highest sampling frequency (5 minutes). This relationship is return-driven in the sense that index returns Granger cause volatility returns. Impulse-response functions show that a one-time innovation on index return has a significant impact on implied volatility. The largest part of the implied volatility response occurs immediately, but we also observe a smaller retarded reaction for up to one hour. A volatility feedback effect is not discernible. If it exists, it appears to be correctly anticipated by

[2] Earlier studies include Christie (1982), Duffee (1995), Figlewski/Wang (2000) and Chan et al. (2003) (based on regressions) and Nelson (1991), Cheung/Ng (1992), Engle/Ng (1993) and Glosten et al. (1993) (based on modelling of conditional volatility through various GARCH specifications).

[3] It is well known that ATM implied volatility is an upward-biased estimate of future realized volatility (see, e.g., Jackwerth/Rubinstein (1996)). We assume that the bias is approximately constant through time and therefore does not significantly influence volatility returns.

traders in the stock market, so that the initial DAX return already incorporates the feedback from the expected volatility reaction.

In the next section, we describe our data and explain how we account for microstructure frictions. Section 3 contains the causality tests, and section 4 concludes.

2 Data

2.1 Raw returns and the smile in option prices

Following the arguments of Jackwerth (2000), we infer DAX levels from DAX futures prices instead of using the observed index levels. The first reason is that DAX futures are the relevant instrument for hedging DAX option positions in practice. Other alternatives for trading the DAX portfolio are less flexible and more costly. Second, the observed index level rests on last trade prices of the index stocks. This means that the underlying stock prices are not simultaneous (see Bollen/Whaley (2004), p. 737).

To obtain the index level S_t corresponding to an observed futures market price F_t at time t, we solve the futures pricing model $F_t = S_t e^{r(T-t)}$ for S_t, where r is the risk-free rate of return and T the futures contract maturity date. We only consider the contract most actively traded on that day, which is normally the nearest available. Dividends do not have to be taken into account since the DAX is a performance index (see Kempf/Bühler (1995) for more details). The resulting "DAX futures implied" DAX values are used to calculate DAX returns and to estimate implied volatilities of DAX options.

Our data come from the joint German and Swiss options and futures exchange, Eurex.[4] The Eurex is one of the world's largest futures and options exchanges and is jointly operated by Deutsche Börse AG and SWX Swiss Exchange. Our database contains all reported transactions of DAX options and futures from January 1995 to December 2005. The average daily trading volume of DAX options (ODAX) and futures (FDAX) in December 2005 was 166,886 and 117,388 contracts. The options are European style. At any point in time during the sample period, at least eight option maturities were available. However, trading is heavily concentrated on the nearby maturities. Trading hours changed several times during our sample period, but both products were traded at least from 09:30 a.m. to 16:00.

As we use time-stamped tick-by-tick data, matching of option and futures prices is straightforward. We extract implied volatilities from DAX option prices using the Black/Scholes (1973) pricing model. We apply the method of Hafner/Wallmeier (2001) to ensure put-call-parity consistent estimates of implied volatilities and remove option prices which violate well-known arbitrage bounds.

Due to the smile in option prices, differences in implied volatilities of subsequent option prices can be due to different levels of moneyness defined as the quotient of strike price and forward price. To restrict the influence of the smile, we only keep ATM options with a moneyness

[4] We are very grateful to the Eurex for providing the data.

between 0.975 and 1.025. Since a small influence of moneyness might still exist, we estimate the smile structure each day following the cubic regression approach described in Hafner/Wallmeier (2001) and Hafner/Wallmeier (2007). We then use the fitted smile function to remove the impact of moneyness on implied volatilities in the relevant moneyness range of 0.975 to 1.025. More specifically, let K denote the strike price of an option with time to maturity $T - t$. Each trade is assigned a moneyness according to:

$$M(t, T, K) = \frac{\ln\left(\frac{K}{F_t(T)}\right)}{\sqrt{T - t}},$$

where $F_t(T)$ is the forward price at time t for maturity T. Thus, ATM options are characterized by a moneyness of 0. Suppressing the arguments of moneyness, we chose the cubic regression function:

$$\sigma = \beta_0 + \beta_1 M + \beta_2 M^2 + \beta_3 D \cdot M^3 + \varepsilon, \tag{1}$$

where σ is the implied volatility, $\beta_{i,}$ $i = 0, 1, 2, 3$ are regression coefficients, ε is a random error, and D is a dummy variable defined as:

$$D = \begin{cases} 0 & , \quad M \leq 0 \\ 1 & , \quad M > 0 \end{cases}.$$

The dummy variable accounts for an asymmetry of the pattern of implied volatilities around the ATM strike ($M = 0$).

Let $\sigma_{imp}(M, t)$ denote the implied volatility of an option with moneyness M traded at time t. Then, the corresponding ATM implied volatility $\sigma_{imp}^{ATM}(t)$ is calculated as

$$\sigma_{imp}^{ATM}(t) = \sigma_{imp}(M, t) - \left[\widehat{\beta}_1 M + \widehat{\beta}_2 M^2 + \widehat{\beta}_3 D \cdot M^3\right],$$

where $\widehat{\beta}_i$ are the estimated regression coefficients.

We classify all observations into two maturity groups. The first contains options with a time-to-maturity between 10 and 30 calendar days, the second contains all observations with an option's time-to-maturity between 31 and 60 days. Options with longer maturities are not considered due to thin trading. Very short maturities below 10 days are also excluded to leave out expiration-day effects and to avoid biases due to inaccurate estimates of implied volatilities.

Implied volatility returns (R_v) and raw returns of the underlying stock index (R_S) over the time period from t_i to t_j are calculated as:

$$R_v(t_i, t_j) = \ln\left[\sigma_{imp}^{ATM}(t_j)\right] - \ln\left[\sigma_{imp}^{ATM}(t_i)\right] \text{ and } R_S(t_i, t_j) = \ln S_{t_j} - \ln S_{t_i},$$

where S denotes the index level underlying the corresponding futures price. The values $\sigma_{imp}^{ATM}(t)$ and S_t are set equal to the last implied volatility and index level observed before t. If the last trade occurred more than 60 seconds before t, the return is not calculated. The results do not change if we further restrict the maximal distance to 30 seconds. We consider four different sampling frequencies $t_j - t_i$, namely 5 minutes, 15 minutes, hourly and daily. We are primarily interested in the high-frequency 5-minute intervals. Results for the longer intervals serve as a means of comparison.

2.2 Microstructure frictions

For data sampled at high frequency, market microstructure frictions have to be taken into account, particularly fluctuating trading activity, infrequent trading and the bid-ask bounce.

Intraday pattern of trading activity

Trading activity in DAX stocks and DAX options often follows an intraday pattern with three different periods: (1) in the morning, trading is typically at its maximum at the opening and then decreases until lunch; (2) the period after lunch is generally characterized by a peak of trading activity around the opening of the US stock exchanges; (3) during the last few hours, activity usually stays at a high level. This intraday pattern may have consequences for the modelling of conditional volatility. For example, it can produce biases in GARCH specifications (see Andersen et al. (1999)). Due to the varying conditional volatility, returns in low-activity intervals are not directly comparable to returns in periods with high trading activity.

In order to correct DAX returns for the intraday trading activity, we follow Andersen et al. (2001) and model the pattern with a Fourier Flexible Form (*FFF*). The main assumption underlying this approach is that an intraday return can be expressed as the product of a Gaussian white noise, a daily (or long-term) volatility component and an intraday pattern effect. Using this decomposition, it is then possible to estimate and filter out the intraday pattern effect using a *FFF* regression (see Taylor (2006) and Andersen et al. (2001) for more details).

Following Andersen et al. (2001), we only use the polynomial part of the *FFF* and break it up into three third order sub-polynomials to account for three different trading regimes during the day. Since the intraday pattern is not supposed to be constant over the 11-year period, we treat each year separately. We apply this filtering only to return data sampled at the highest frequencies (5 and 15 minutes). While conceptually relevant, a robustness check shows that the filtering is not important empirically. Filtered returns turn out to be very similar to raw returns and all results remain valid when the filtering is omitted.

Infrequent trading

The problem of infrequent trading arises if intervals without market transactions occur. There are several ways to handle this problem. Some authors simply set the missing value equal to the last transaction price (e.g. Stephan/Whaley (1990) or Dennis et al. (2006)), while others interpolate between the last and the next price to fill the gap (see, e.g., Corsi et al. (2001) for a discussion). However, the first method has the shortcoming that it leads to a bias in vector autoregressions, whereas the second approach generates spurious autocorrelations. We therefore restrain from generating fictitious values to replace non-available market prices. Instead, in each part of the study we control for the presence of a sufficient number of available lags and discard data which do not satisfy this requirement.

Bid-ask bounce

As is well known, the bid-ask bounce leads to a negative first order autocorrelation of returns (Roll (1984)). This spurious autocorrelation comes from successive trades, where one is executed at the bid, the other at the ask price. The bid-ask bounce effect is more pronounced for DAX options with their relatively large bid-ask spreads than for DAX futures and the underlying index. We use the standard method to remove spurious autocorrelation from returns, which consists in filtering the returns with a MA(1) process (see, e.g., Stephan/Whaley (1990), Easley et al. (1998) and Gwilym/Buckle (2001)):

$$R_t = \mu + e_t - \theta e_{t-1},$$

where R_t is the observed return, μ the unconditional mean of the return series, θ the moving average coefficient, and e_t the innovation of the process. Since the innovations from the MA(1) process are uncorrelated, we can use them as bid-ask bounce corrected returns. Through the FFF filtering and the MA(1) correction, raw returns R_v and R_S are transformed into adjusted r_v and r_S which are used in the empirical tests.

2.3 Descriptive return statistics

Descriptive statistics for DAX and volatility 5-minute log returns (raw returns R_S and R_v as well as adjusted returns r_S and r_v) are given in Table 1. The returns of implied volatilities are reported for a time to maturity of 31 to 60 days. Results for the shorter maturity are very similar. For better comparison, all statistics are given on a daily basis. We compute the summary statistics for different time-windows: each year, the whole 11-year period (All) and three subperiods corresponding to the long bullish market of the late nineties (P1: 01/01/1995 to 03/07/2000), the bearish market that followed the end of the tech bubble and the 9/11 attacks (P2: 03/08/2000 to 03/12/2003) and the bullish market that took place after the beginning of the Iraq War (P3: 03/13/2003 to 12/31/2005).

In most years, skewness of DAX returns is slightly negative and kurtosis widely exceeds three. P-values of the Jarque-Bera test (not reported in the table) unambiguously reject normality on the 1% level. The bearish market (P2) is characterized by a negative mean, high variance and kurtosis and strongly negative skewness. Volatility returns have a much higher variance than DAX returns. The sign of skewness of volatility return varies. Kurtosis typically exceeds three, so that the null of normality is rejected for all subsamples. The adjustments of raw returns to account for intraday patterns and microstructure effects have only marginal effects on DAX returns, whereas the MA(1) correction for variance returns noticeably modifies (unconditional) variance and skewness.

2.4 Correlation with lagged returns

To examine the lead-lag relationship between index and volatility returns, we calculate the correlation coefficient of DAX returns in a 5-minute interval t with implied volatility returns in 5-minute interval $t + j$, where $j \in \{-250, \ldots, 250\}$. As one trading day typically comprises

	R_S (5-minute)				r_S (5-minute)				R_v (5-minute)				r_v (5-minute)			
Year	Mean	Std.	S.	K.	Mean	Std.	S.	K.	Mean	Std.	S.	K.	Mean	Std.	S.	K.
1995	-0.0008	0.0066	-0.12	9.93	-0.0007	0.0065	-0.07	10.51	-0.0023	0.1690	0.50	9.98	-0.0057	0.1557	0.54	12.05
1996	0.0005	0.0058	-0.41	11.37	0.0004	0.0057	-0.57	14.13	0.0453	0.1853	-0.06	5.44	0.0408	0.1712	-0.07	5.05
1997	0.0001	0.0115	-0.53	18.65	0.0000	0.0111	-0.55	14.35	0.0158	0.2154	-0.22	16.00	0.0308	0.2006	-0.59	11.95
1998	-0.0007	0.0149	-0.24	9.66	-0.0005	0.0147	-0.21	9.22	0.0247	0.1537	-0.46	11.74	0.0190	0.1448	-0.74	11.05
1999	0.0004	0.0117	-0.21	7.39	0.0004	0.0116	-0.27	7.56	-0.0407	0.1399	-0.14	6.43	-0.0466	0.1328	-0.25	6.50
2000	-0.0005	0.0121	-0.18	6.79	-0.0004	0.0119	-0.20	6.18	-0.0116	0.1586	-0.40	10.77	-0.0119	0.1475	-0.68	8.01
2001	-0.0006	0.0172	-1.62	116.86	-0.0006	0.0175	-1.78	205.46	0.1044	0.2246	0.41	10.51	0.1191	0.2125	0.40	10.53
2002	-0.0018	0.0208	-0.02	8.76	-0.0016	0.0203	-0.11	7.11	0.0338	0.1826	0.29	11.05	0.0362	0.1733	0.39	11.91
2003	0.0003	0.0167	-0.05	7.61	0.0008	0.0164	-0.01	7.01	0.0145	0.1457	-0.23	9.54	0.0189	0.1392	-0.28	9.06
2004	-0.0005	0.0076	-0.07	8.94	-0.0004	0.0074	-0.04	11.20	0.0154	0.1191	0.28	8.95	0.0206	0.1137	0.36	8.44
2005	0.0004	0.0058	-1.10	34.57	0.0003	0.0057	-1.33	51.96	0.0145	0.1333	0.56	18.68	0.0237	0.1283	0.70	19.17
P1	0.0000	0.0111	-0.32	14.14	0.0001	0.0110	-0.31	12.69	0.0115	0.1762	-0.09	12.53	0.0098	0.1639	-0.29	10.48
P2	-0.0015	0.0180	-0.56	42.84	-0.0013	0.0178	-0.71	75.36	0.0398	0.1934	0.34	11.94	0.0475	0.1828	0.36	12.12
P3	0.0006	0.0103	0.03	14.88	0.0007	0.0101	0.11	13.86	0.0174	0.1307	0.29	14.19	0.0251	0.1255	0.34	14.39
All	-0.0003	0.0136	-0.53	53.10	-0.0002	0.0134	-0.65	91.86	0.0190	0.1673	0.15	13.89	0.0221	0.1572	0.08	13.07

Table 1: Summary statistics. We report Mean, Standard deviation (Std.), Skewness (S.) and Kurtosis (K.) for 5-minute stock index (DAX) returns and 5-minute implied volatility returns for options with the second nearest time-to-maturity. Mean and standard deviation are given on a daily basis. Periods P1 to P3 are 01/01/1995 - 03/07/2000 (P1), 03/08/2000 - 03/12/2003 (P2) and 03/13/2003 - 12/31/2005 (P3).

about 100 intervals of 5 minutes, the number of 250 leads and lags corresponds to 5 trading days around t. Calculations are based on all t during the total sample period. Figure 1 shows that the correlation coefficient is near zero for lagged volatility ($j < 0$). Thus, the DAX return does not seem to be systematically related to the preceding implied volatility return. However, we find a significantly negative correlation of DAX returns not only with contemporaneous volatility returns ($j = 0$), but also with volatility returns in the next few 5-minute periods. About 1 hour after the stock price shock ($j > 12$), the correlation goes back to zero. This observation supports the hypothesis that implied volatility is adjusted to changes in the index level, so that the relationship seems to be primarily return-driven. It is important to note that the retarded reaction of volatility cannot be explained by thin trading and missing volatility returns, because the return in period $t+j$ is calculated only if transaction prices at the beginning and the end of the period are available. If the return in $t+2$ is available while the $t+1$ return is missing, the implied volatility at the beginnning of period $t+2$ (which is available) should already reflect the price innovation in t. Thus, the return in $t+2$ should not be influenced regardless of whether the $t+1$ return is available or not.

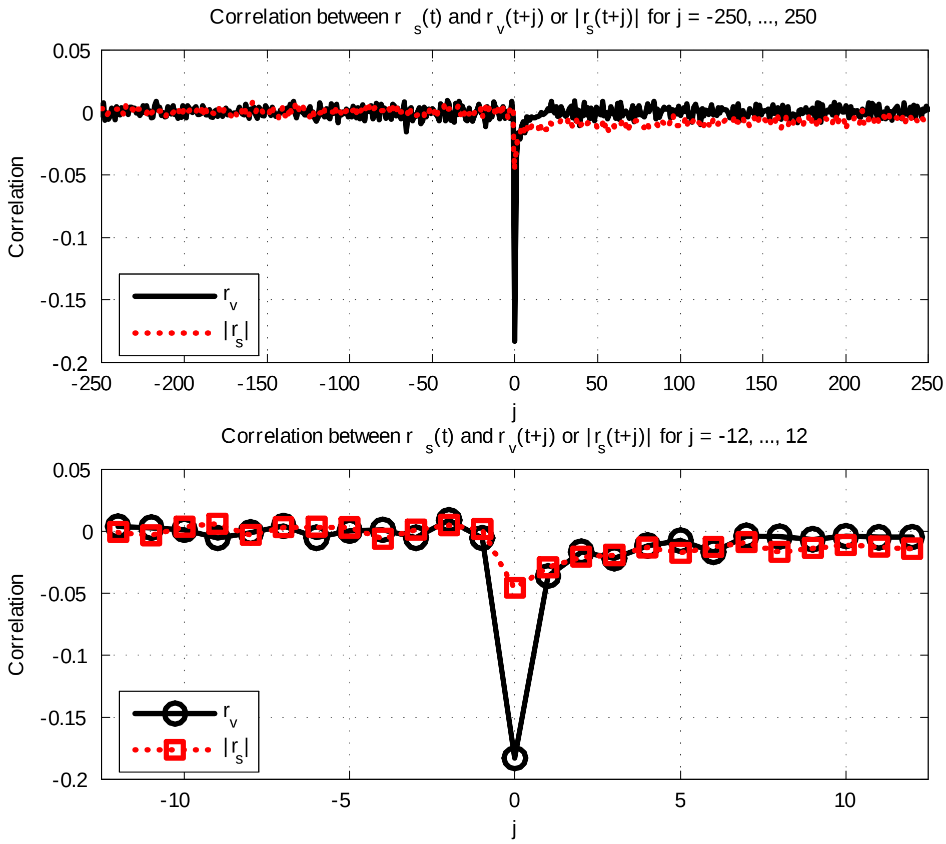

Figure 1: Correlations for the first maturity and 5-minute returns over the total sample period from 1995 to 2005. The upper panel reports correlations for $j = -250, \ldots, 250$ while the bottom panel focuses on correlations for $j = -12, \ldots, 12$.

Following Bollerslev et al. (2006), we also examine the correlation between $r_{S,t}$ and the absolute return $|r_{S,t+j}|$. In this specification, absolute returns serve as an alternative measure of realized volatility. It is apparent from Figure 1 that there is no noticeable correlation of absolute returns before t with DAX return in t. The contemporaneous correlation in t is negative, which means that negative returns are typically larger in magnitude than positive returns. After t, absolute returns are negatively correlated with $r_{S,t}$. This relationship gets weaker the larger the lag, but it is recognizable for all $j \in \{1, \ldots, 250\}$. These observations are similar to the US results in the study of Bollerslev et al. (2006). The correlation series for absolute returns shows that negative DAX returns typically increase subsequent return dispersion. Again, this is compatible with a return-driven effect. However, the analysis focuses on total cross-autocorrelations only and leaves out partial cross-autocorrelations. It could be the case that correlations computed for lags $j \geq 2$ are completely due to the correlation at lag $j = 1$. A more detailed study is thus necessary to identify causality and the number of lagged DAX returns which have an impact on contemporaneous volatility returns.

3 Causality analysis

3.1 Granger causality test

We carry out a Granger causality test, i.e. each variable is regressed on a constant and p of its own lags as well as on p lags of the other variable in terms of the following VAR(p) vector autoregression:

$$\mathbf{R}_t = \mathbf{c} + \sum_{i=1}^{p} \mathbf{\Phi}^{(i)} \cdot \mathbf{R}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{2}$$

where $\mathbf{R}_t$ is the (2×1) vector of DAX and volatility returns, $\mathbf{c}$ is the (2×1) vector of constants and $\mathbf{\Phi}^{(i)}$ is the (2×2) matrix of autoregressive slope coefficients for lag i. We call the two equations of system (2) the *index return regression* (*IRR*) and *volatility return regression* (*VRR*). The matrices $\mathbf{\Phi}^{(i)}$, $i = 1, \ldots, p$ are lower triangular if the relationship is return-driven and upper triangular if it is volatility-driven. The two explanations cohabit if some matrices $\mathbf{\Phi}^{(i)}$ are full.

Table 4 summarizes the results of the Wald F-test for the eleven-year sample period and each sampling frequency for up to five lags ($p = 5$).[5] We report the p-value of the F-test (first line) and the number of observations included in the regression (second line). When more lags are considered, the number of complete return series and therefore the number of observations sharply decreases.[6] This is why we focus our attention primarily on the case of $p = 1$. We regard the rejection of an effect as more reliable if the p-values for both option maturities are significant.

The results provide evidence in favour of a return-driven relationship. At the highest sampling frequency, the null hypothesis that past index returns do not contribute to the explanation of current implied volatility return is always rejected at least at the 1% significance level. This

[5] According to the Akaike and Schwartz information criteria, the optimal number of lags varies between 2 and 5 for sampling frequencies of 5 and 15 minutes and is equal to 1 for hourly and daily data.

[6] The volume of hourly data is relatively low, because we calculate hourly returns only if a transaction price is available from the last 60 seconds (see Section 2).

10 Days < Time-to-Maturity ≤ 1 Month										
	Return-driven relationship with $p = \ldots$					Volatility-driven relationship with $p = \ldots$				
	1	2	3	4	5	1	2	3	4	5
5m.	<.0001	<.0001	<.0001	<.0001	<.0001	.1143	.0536	.0128	.0019	.0069
	39137	27434	19930	14858	11302	39137	27434	19930	14858	11302
15m.	<.0001	<.0001	<.0001	<.0001	<.0001	.0735	.8248	.3546	.8687	.9148
	11138	7120	4671	3099	2073	11138	7120	4671	3099	2073
1h.	.0012	.0021	.0319	.4499	.7632	.4472	.2405	.5287	.2835	.6831
	598	281	133	70	39	598	281	133	70	39
1d.	.6335	.3989	.4981	.6869	.5154	.9092	.2530	.3598	.5171	.6044
	1591	1458	1325	1193	1061	1591	1458	1325	1193	1061
1 Month < Time-to-Maturity ≤ 2 Months										
	Return-driven relationship with $p = \ldots$					Volatility-driven relationship with $p = \ldots$				
	1	2	3	4	5	1	2	3	4	5
5m.	<.0001	<.0001	<.0001	<.0001	<.0001	.0887	.7341	.0073	.1023	.4523
	12070	6642	4086	2706	1854	12070	6642	4086	2706	1854
15m.	<.0001	<.0001	<.0001	.0005	.1135	.1026	.7916	.6121	.7598	.5702
	3272	1668	920	539	324	3272	1668	920	539	324
1h.	.0002	<.0001	.0187	.1102	.7569	.7663	.8132	.0989	.0704	.2540
	746	265	116	54	24	746	265	116	54	24
1d.	.0130	.0532	.0219	.0491	.0137	.0594	.1783	.1621	.1446	.1178
	2774	2773	2772	2771	2770	2774	2773	2772	2771	2770

Table 2: Granger causality test over the total sample period from 1995 to 2005. We report p-values of the F-tests for up to 5 lags in the first line and the number of available data (i.e. of successive returns) in the second line. Sampling frequencies are 5 minutes (5m.), 15 minutes (15m.), hourly (1h.) and daily (1d.).

causality is discernible for intervals of up to 60 minutes. The same results hold true in yearly subsamples except for hourly data. The return-driven relationship for hourly data only comes from the subperiod 03/08/2000 to 03/12/2003. The test statistics do not give evidence for a volatility-driven effect. Past volatility returns do not significantly add to the explanatory power of past index returns in explaining current index returns.[7] Results for periodical subsamples are similar to those for the total sample period.[8] The only notable difference is that we find a causality running in both directions in the bearish market from 2000 to 2003 when considering volatility computed from options with the second nearest time-to-maturity. In fact, the significant daily return-driven relationship that we find for the total sample comes entirely from this period. In all, we conclude that the lead-lag relationship of DAX returns and implied volatility returns is compatible with Granger causality running from index return to volatility.

The results achieved so far do not rule out the possibility that a feedback-effect was not detected because it occurs in a more subtle fashion. For instance, one may suspect that only large volatility returns have an impact on index returns. To investigate if such non-linear feedback effects exist, we performed a non-parametric causality test introduced by Baek/Brock (1992) and extended and improved by Hiemstra/Jones (1994) and Diks/Panchenko (2005). This test examines if the probability distribution of future index returns is different if the information set contains either the history of both DAX and volatility returns or the history of DAX returns alone. The test statistics of the test by Diks/Panchenko (2005) (not shown here) do not provide evidence in favour of a non-linear feedback effect.

3.2 Contemporaneous versus lagged relationship

The finding of a return-driven effect in high-frequency data leaves open the question of how important this lead-lag relationship is compared to the strong contemporaneous correlation of index and volatility return. To enable this comparison, we extend the volatility return regression by adding contemporaneous DAX returns ($r_{S,t}$) as explanatory variable for volatility returns ($r_{v,t}$):

$$r_{v,t} = c^* + \sum_{i=1}^{p} \phi^*_{i,1} \cdot r_{S,t-i} + \sum_{i=1}^{p} \phi^*_{i,2} \cdot r_{v,t-i} + \beta^* r_{S,t} + \varepsilon^*_t. \tag{3}$$

We compare the unrestricted model (3) with two restricted versions:

- restricted model 1, characterized by $\beta^* = 0$, and
- restricted model 2, characterized by $\phi^*_{i,1} = 0 \;\; \forall i = 1, \dots, p$.

Restricted model 1 is identical to the volatility return regression of the last section, whereas restricted model 2 replaces lagged index returns by the contemporaneous index return as explanatory variable. Using OLS with Newey/West (1987) standard errors, we estimate the three

[7] There is weak evidence in favour of a volatility-driven effect when considering 5-minute returns. However, this result depends on the number of lags and the time-to-maturity. Therefore, its economic significance seems questionable.

[8] The results for the three subsamples P1 to P3 are available on request.

regressions for all sampling frequencies and the two times-to-maturity. We also decompose the variance V of volatility returns according to:

$$\begin{aligned} V(r_{v,t}) &= V(\sum_{i=1}^{p} \phi^*_{i,1} \cdot r_{S,t-i} + \sum_{i=1}^{p} \phi^*_{i,2} \cdot r_{v,t-i} + \beta^* r_{S,t} + \varepsilon^*_t) \\ &= V(r_{v,t}) \cdot (VL + VCR + COV + VE), \end{aligned} \tag{4}$$

where

$$\begin{aligned} VL &= V(\sum_{i=1}^{p} \phi^*_{i,1} \cdot r_{S,t-i} + \sum_{i=1}^{p} \phi^*_{i,2} \cdot r_{v,t-i})/V(r_{v,t}), \\ VCR &= V(\beta^* r_{S,t})/V(r_{v,t}), \\ COV &= 2Cov(\sum_{i=1}^{p} \phi^*_{i,1} \cdot r_{S,t-i} + \sum_{i=1}^{p} \phi^*_{i,2} \cdot r_{v,t-i}, \beta^* r_{S,t})/V(r_{v,t}), \\ VE &= V(\varepsilon^*_t)/V(r_{v,t}). \end{aligned}$$

VL, VCR, COV and VE measure the percentage of the overall variance of $r_{v,t}$ explained by lagged DAX and volatility returns (VL), contemporaneous DAX returns (VCR), covariance between lagged DAX and volatility returns and contemporaneous DAX returns (COV) and variance of the residuals (VE).

In the first three columns of Table 5, we report the sampling frequency, the number of lags employed[9] and the number of valid observations. The p-values 1 and 2 refer to a test of the hypothesis that the MSE of a forecast of $r_{v,t}$ based on the unrestricted model is the same as the MSE based on restricted models 1 and 2, respectively. In the case of restricted model 1, this hypothesis is always rejected at the 99% confidence level. Thus, adding $r_{S,t}$ to the set of regressors improves the explanatory power of the model. This finding confirms that part of the relationship occurs contemporaneously. In the second comparison we test whether the model with lagged and contemporaneous returns (unrestricted model) has additional explanatory power above restricted model 2 which only uses contemporaneous index returns. Again, with one exception, all p-values are below 1%. Thus, even after controlling for contemporaneous returns, a significant part of volatility returns can be traced back to leading index returns. For high-frequency data, a substantial part of the variance of $r_{v,t}$ can be attributed to leading returns (VL). At lower frequencies, the lead-lag-relationship is negligible, and the variation of $r_{v,t}$ is primarily attributed to contemporaneous index returns (VCR).

3.3 Impulse-response functions

As a natural extension of the Granger causality analysis, we use impulse-response functions (IRFs) to illustrate the dynamic relations between DAX and implied volatility returns. They

[9] The number of lags is taken to be alternatively $p = 1$ or the optimal choice indicated by the Akaike and Schwartz criterion. For hourly and daily data, the latter choice is equal to $p = 1$.

10 Days < Time-to-Maturity ≤ 1 Month								
	p	*#data*	p-value 1	p-value 2	*VL*	*VCR*	*COV*	*VE*
5m.	1	39137	<0.0001	<0.0001	2.17%	4.22%	0.38%	93.22%
	5	11302	<0.0001	<0.0001	6.55%	4.53%	0.67%	88.25%
15m.	1	11138	<0.0001	<0.0001	1.20%	14.36%	0.60%	83.84%
	3	4671	<0.0001	<0.0001	9.55%	7.75%	0.47%	82.23%
1h.	1	598	<0.0001	0.0009	0.19%	41.90%	-0.02%	57.92%
1d.	1	1591	<0.0001	0.5923	0.24%	40.09%	-0.03%	59.70%
1 Month < Time-to-Maturity ≤ 2 Months								
	p	*#data*	p-value 1	p-value 2	*VL*	*VCR*	*COV*	*VE*
5m.	1	12070	<0.0001	<0.0001	1.74%	2.91%	0.09%	95.27%
	4	2706	<0.0001	<0.0001	9.58%	2.81%	-0.18%	87.79%
15m.	1	3272	<0.0001	<0.0001	1.91%	10.06%	0.10%	87.93%
	2	1668	<0.0001	<0.0001	4.84%	7.40%	0.58%	87.17%
1h.	1	746	<0.0001	0.0013	1.20%	14.36%	0.60%	83.84%
1d.	1	2774	<0.0001	0.0185	0.19%	41.90%	-0.02%	57.92%

Table 3: Analysis of variance for the total sample period from 1995 to 2005. The first two columns display the sampling frequency and the number of lags employed in the regressions. The third and fourth columns report the p-value of testing whether the MSE of the unrestricted model is the same as the MSE of the restricted model 1 or 2 (p-value 1 and p-value 2). The last four columns report the portion of volatility returns' variance explained by the various subsets of regressors.

allow us to assess how important the impact of a one-time impulse in one variable on future values of the other variable is and how long the impact lasts. Let subscripts i and j refer to DAX and volatility returns. We denote by s a forecast period starting from date t (forecast horizon $t + s$) and assume that the state of the system as of date t is known. Then, IRF is a function of s whose values correspond to the revision in the forecast of $r_{i,t+s}$ induced by the information that the value of $r_{j,t+s}$ is higher than expected ($\varepsilon_{j,t} > 0$). We orthogonalize the impulses to ensure that direct and indirect effects of an impulse are considered (for details, see Hamilton (1994), pp. 318-323).

Figures 2 and 3 illustrate the orthogonalized IRFs for DAX and volatility return innovations. The magnitude of the shock is fixed at one standard deviation of the uncorrelated (orthogonalized) innovation. We add two standard error bands from Monte Carlo simulations with 100,000 paths. The horizon s varies from 4 to 12 intervals depending on the sampling frequency. This corresponds to a range of one hour (5-minute data) to 5 days (daily data). The units on the vertical axis are in DAX or volatility return standard deviations.

The IRFs for responses of volatility return to an impulse of DAX return are typically negative (see Figure 2). For 5- and 15-minute data, the responses directly after a shock have a magnitude of about -0.1 to -0.2 standard deviations. The IRFs then remain significantly negative for about 15 to 45 minutes. As expected, the responses are less important for lower frequencies. Figure 3 shows that the impact of a volatility shock on DAX returns is very limited. This observation is compatible with our findings of Section 3.1.

3.4 Effect of liquidity and net buying pressure

Bollen and Whaley (2004) find that implied volatility returns of S&P500 options are directly related to net buying pressure for index puts. The results suggest that this buying pressure, which is typically decreasing with moneyness, drives the downward sloping shape of the implied volatility function. In a related paper, Chan et al. (2005) argue that net buying pressure and liquidity affect the response of implied volatility to returns. The response seems to be distorted by net buying pressure for low moneyness options and thin trading of medium moneyness options. Chan et al. (2003) report similar results for options on the Hang Seng Index (HSI).

In order to avoid distortions due to varying degrees of net buying pressure in this study, we only include ATM options. In our sample, call option trades are more numerous and voluminous than put option trades, which indicates that buying pressure is presumably not substantial in this moneyness range. In addition, ATM DAX options used in our study represent the highest transaction number (about 40% of all options) and trading volume (about 45% of all options) among different moneyness classes. Liquidity in out-of-the-money options is not sufficiently high to carry out our high frequency analysis.

To ensure that the remaining effect of net buying pressure and thin trading is small, we subdivide all trading days of our sample period in 3 quantiles of low, medium and high liquidity days and separately in 3 quantiles of days with low, medium and high buying pressure. Combining both criteria, we obtain nine groups of trading days. We use trading volume (in Euro per day) as liquidity measure and the ratio of trading volume of puts to the trading volume of calls (per

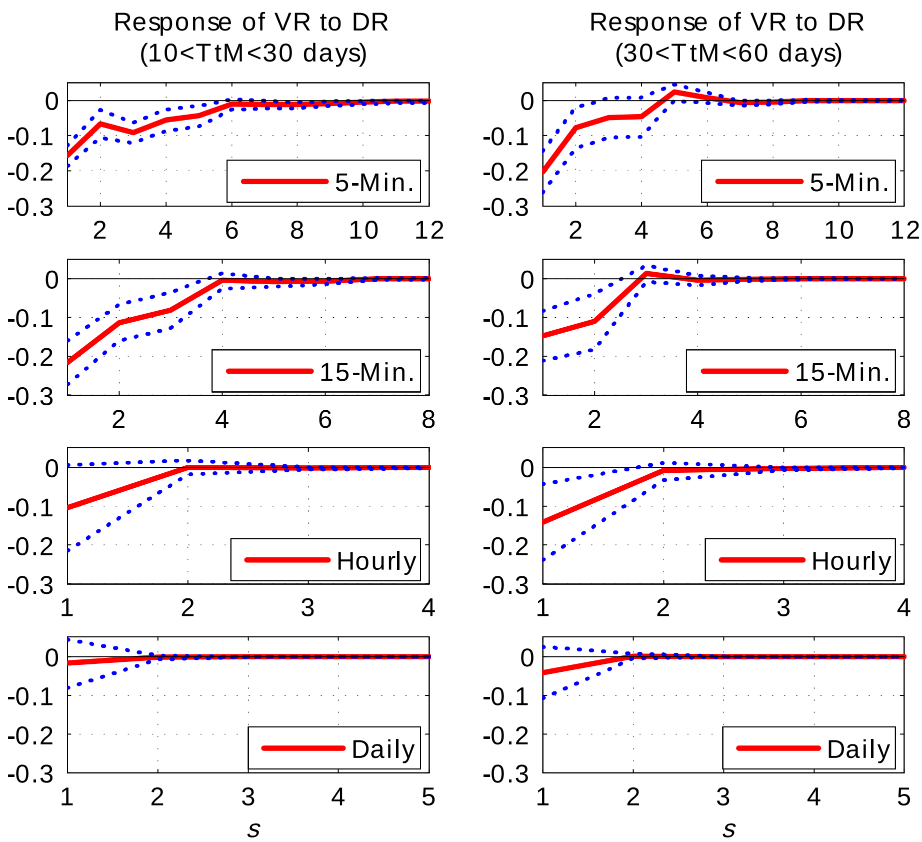

Figure 2: Orthogonalized impulse responses of volatility returns (VR) to DAX returns (DR). The panels on the left (respectively on the right) display the IRFs for the volatility computed from options with the time-to-maturity (TtM) ranging from 10 to 30 days (30 to 60 days).

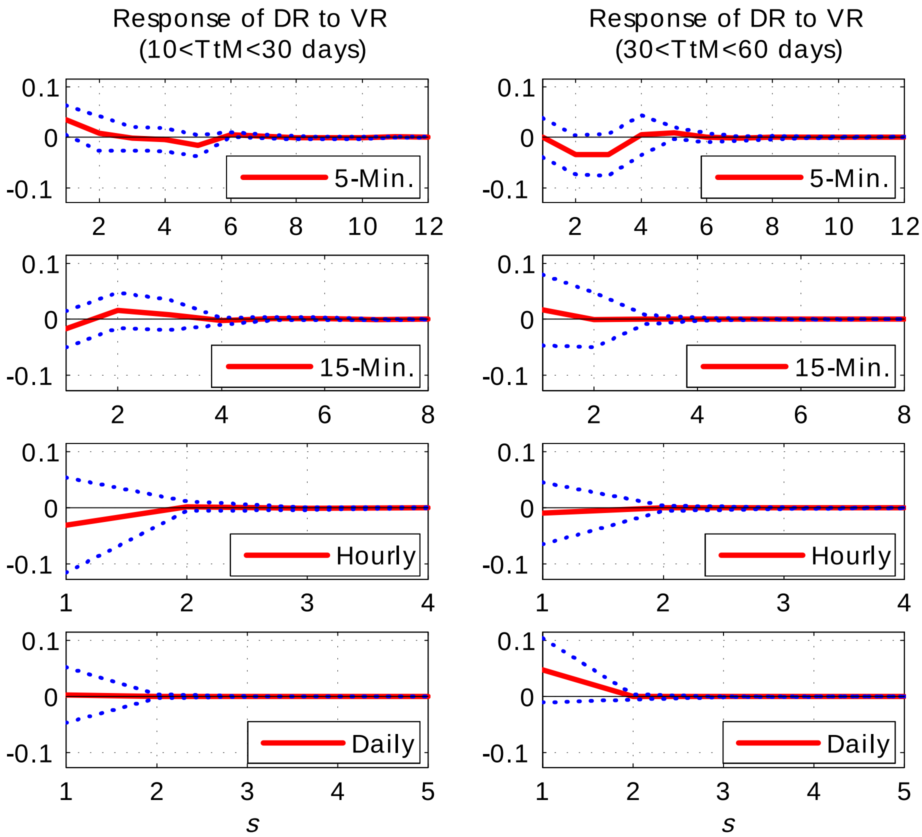

Figure 3: Orthogonalized impulse responses of DAX returns (DR) to volatility returns (VR). The panels on the left (respectively on the right) display the IRFs for the volatility computed from options with the time-to-maturity (TtM) ranging from 10 to 30 days (30 to 60 days).

day) as a measure of net buying pressure. We repeat our empirical analyses for each of the nine subgroups. While significance levels of the Granger causality tests are partly lower due to the smaller number of observations, we do not find any substantial differences across groups. Graphs of the lead-lag relationships within subgroups are almost identical to the structure in Figure 1.[10] Thus, it seems safe to conclude that our results are not driven by net buying pressure or liquidity effects.

4 Conclusion

It is well known that index returns are inversely related to volatility returns. The relationship is so strong that it constitutes an important stylized fact in finance. Nevertheless, the origin and the causes of the effect are not yet well understood, which is particularly true for financial markets in Europe. In this paper, we analyze the return-volatility relationship at the German market. We calculate return series for 5-minute intervals from tick-by-tick DAX option and futures data over the time period from 1995 to 2005. We also consider lower return frequencies for the purpose of comparison. Our volatility measure is the implied volatility of at-the-money options. This allows us to more accurately detect changes in volatility than previous studies which use measures of realized volatility. In addition, as ATM implied volatilities can be determined independently of the underlying asset return, index and volatility returns can be modelled jointly in a VAR model. This provides a flexible framework for running Granger causality tests and computing impulse-response functions.

We find that a lead-lag relationship exists only in high-frequency data. The relationship is return-driven in the sense that index returns Granger cause volatility returns. This causal relationship is statistically and economically significant and can be clearly separated from the contemporaneous correlation. A volatility feedback effect does not show up. Either it does not exist, or the market promptly incorporates all direct and indirect impulses into market prices so that the feedback effect fully evolves within the 5-minute intervals.

Our paper does not account for jumps either in the index level or the (implied) volatility process. The relationship around such discontinuities could offer further insight into the nature of the effect. The behaviour of the relationship through time could also be of interest, because it is closely related to the pricing of volatility and the dynamics of the volatility risk premium. But the most important topic for further research still seems to be the question why this strong effect exists and which economic fundamentals are effectively driving it.

References

Andersen, T./Bollerslev, T./Das, A. (2001), Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns, *Journal of Finance* **56**(1), pp. 305–327.

[10] These results are available on request.

Andersen, T./Bollerslev, T./Lange, S. (1999), Forecasting financial market volatility: Sample frequency vis-à-vis forecast horizon, *Journal of Empirical Finance* **6**(5), pp. 457–477.

Aydemir, A./Gallmeyer, M./Hollifield, B. (2006), *Financial Leverage Does Not Cause the Leverage Effect*, Working Paper, Mays Business School, Texas University.

Bae, J./Kim, C.-J./Nelson, C. (2007), Why are stock returns and volatility negatively correlated?, *Journal of Empirical Finance* **14**(1), pp. 41–58.

Baek, E./Brock, W. (1992), *A General Test for Nonlinear Granger Causality: Bivariate Model*, Working Paper, Iowa State University and University of Wisconsin, Madison.

Bekaert, G./Wu, G. (2000), Asymmetric volatility and risk in equity markets, *Review of Financial Studies* **13**(1), pp. 1–42.

Black, F. (1976), *Studies of Stock Price Volatility Changes*, Proceedings of the meetings of the American Statistical Association, Business and Economics Section, Chicago, 177-181.

Black, F./Scholes, M. (1973), The pricing of options and corporate liablilities, *Journal of Political Economy* **81**(3), pp. 637–654.

Bollen, N./Whaley, R. (2004), Does net buying pressure affect the shape of implied volatility functions?, *Journal of Finance* **59**(2), pp. 711–753.

Bollerslev, T./Litvinova, J./Tauchen, G. (2006), Leverage and volatility feedback effects in high-frequency data, *Journal of Financial Econometrics* **4**(3), pp. 353–384.

Bollerslev, T./Zhou, H. (2006), Volatility puzzles: A unified framework for gauging return-volatility regression, *Journal of Econometrics* **131**(1-2), pp. 123–150.

Bouchaud, J./Matacz, A./Potters, M. (2001), The leverage effect in financial markets: Retarded volatility and market panic, *Physical Review Letters* **87**(22), pp. (228701) 1–4.

Campbell, J. (1987), Stock returns and the term structure, *Journal of Financial Economics* **18**(2), pp. 373–399.

Campbell, J./Hentschel, L. (1992), No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* **31**(3), pp. 281–331.

Chan, K./Cheng, L./Lung, P. (2003), Moneyness and the response of the implied volatilities to price changes: The empirical evidence from HSI options, *Pacific-Basin Finance Journal* **11**(4), pp. 527–553.

Chan, K./Cheng, L./Lung, P. (2005), Asymmetric volatility and trading activity in index futures options, *The Financial Review* **40**(3), pp. 381–407.

Chelley-Steeley, P./Steeley, J. (2005), The leverage effect in the UK stock market, *Applied Financial Economics* **15**(6), pp. 409–423.

Chen, X./Ghysels, E. (2007), *News - Good or Bad - and its Impact over Multiple Horizons*, Working Paper, University of North Carolina at Chapel Hill.

Cheung, Y.-W./Ng, L. K. (1992), Stock price dynamics and firm size: Am empirical investigation, *The Journal of Finance* **47**(5), pp. 1985–1997.

Christie, A. (1982), The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* **10**(4), pp. 407–432.

Corsi, F./Zumbach, G./Müller, U./Dacorogna, M. (2001), Consistent high-precision volatility from high-frequency data, *Economic Notes* **30**(2), pp. 183–204.

Dennis, P./Mayhew, S./Stivers, C. (2006), Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon, *Journal of Financial and Quantitative Analysis* **41**(2), pp. 381–406.

Diks, C./Panchenko, V. (2005), A note on the hiemstra-jones test for granger non-causality, *Studies in Nonlinear Dynamics and Econometrics* **9**(2), pp. 1234–1234.

Duffee, G. (1995), Stock returns and volatility: A firm level analysis, *Journal of Financial Economics* **37**(3), pp. 399–420.

Dufour, J.-M./Garcia, R./Taamouti, A. (2006), *Measuring Causality Between Volatility and Returns with High-Frequency Data*, Working Paper, Université de Montréal.

Easley, D./O'Hara, M./Srinivas, P. (1998), Option volume and stock prices: Evidence on where informed traders trade, *Journal of Finance* **53**(2), pp. 431–465.

Engle, R./Ng, V. (1993), Measuring and testing the impact of news on volatility, *Journal of Finance* **48**(5), pp. 1749–1778.

Ericsson, J./Huang, X./Mazzotta, S. (2007), *Leverage and Asymmetric Volatility: The Firm Level Evidence*, Working Paper, McGill University.

Figlewski, S./Wang, X. (2000), *Is the Leverage Effect a Leverage Effect?*, Working Paper Series 00-37, New York University.

French, K./Schwert, G./Stambaugh, R. (1987), Expected stock returns and volatility, *Journal of Financial Economics* **19**(1), pp. 3–29.

Ghysels, E./Santa-Clara, P./Valkanov, R. (2005), There is a risk-return tradeoff after all, *Journal of Financial Economics* **76**(3), pp. 509–548.

Giot, P. (2005), On the relationships between implied volatility indices and stock index returns, *The Journal of Portfolio Management* **26**, pp. 12–17.

Glosten, L./Jaganathan, R./Runkle, D. (1993), On the relation between the expected value and the volatility of the nominal excess return on the stocks, *Journal of Finance* **48**(5), pp. 1779–1801.

Gwilym, O./Buckle, M. (2001), The lead-lag relationship between the FTSE100 stock index and its derivative contracts, *Applied Financial Economics* **11**(4), pp. 385–393.

Hafner, R./Wallmeier, M. (2001), The dynamics of DAX implied volatilities, *International Quarterly Journal of Finance* **1**(1), pp. 1–27.

Hafner, R./Wallmeier, M. (2007), Volatility as an asset class: European evidence, *European Journal of Finance* **13**(7), pp. 621–644.

Hamilton, J. (1994), *Time Series Analysis*, Princeton.

Harvey, C. (2001), The specifications of conditional expectations, *Journal of Empirical Finance* **8**(5), pp. 573–638.

Hiemstra, C./Jones, J. (1994), Testing for linear and nonlinear granger causality in the stock price-volume relation, *Journal of Finance* **49**(5), pp. 1639–1664.

Jackwerth, J. (2000), Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* **13**(2), pp. 433–451.

Jackwerth, J. C./Rubinstein, M. (1996), Recovering probability distributions from option prices, *Journal of Finance* **51**(5), pp. 1611–1631.

Kempf, A./Bühler, W. (1995), DAX index futures: Mispricing and arbitrage in German markets, *Journal of Futures Markets* **15**(7), pp. 833–859.

Lundblad, C. (2007), The risk return tradeoff in the long-run: 1836-2003, *Journal of Financial Economics* **85**(1), pp. 123–150.

Nelson, D. (1991), Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* **59**(2), pp. 347–370.

Newey, W. K./West, K. D. (1987), A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**(3), pp. 703–708.

Pindyck, R. (1984), Risk, inflation, and the stock market, *American Economic Review* **74**(3), pp. 335–351.

Roll, R. (1984), A simple impicit measre of the effective bid-ask sread in an efficient market, *Journal of Finance* **39**(4), pp. 1127–1139.

Scruggs, J. (1998), Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* **53**(2), pp. 575–603.

Stephan, J./Whaley, R. (1990), Intraday price change and trading volume relations in the stock and stock option markets, *Journal of Finance* **45**(1), pp. 191–220.

Taylor, S. (2006), *Asset Price Dynamics, and Prediction*, Princeton University Press, Princeton.

Turner, C./Startz, R./Nelson, C. (1989), A markov model of heteroskedasticity, risk, and learning in the stock market, *Journal of Financial Economics* **25**(1), pp. 3–22.