

Discussion Paper

Discussion Paper No. 96-25

Delta-Neutral Volatility Trading with Intra-Day Prices: An Application to Options on the DAX

Christian Schmitt
Jürgen Kaehler

7616M

01. JULI 1997

Wirtschaft
Kiel

W 636 (96.25) m. dr. sig. gfa

W 636 (96.25)



ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

International Finance Series

Delta-Neutral Volatility Trading with Intra-Day Prices: An Application to Options on the DAX

by

Christian Schmitt* and Jürgen Kaehler**

** Zentrum für Europäische Wirtschaftsforschung (ZEW)*

*** University of Exeter, UK and ZEW*

December 1996

Abstract

This paper evaluates the profitability of applying four different volatility forecasting models to the trading of straddles on the German stock market index DAX. Special care has been taken to use simultaneous intra-day prices and realistic transaction costs. Furthermore, straddle positions were evaluated on a daily basis to preserve delta neutrality. The four models applied in this paper are: historical volatility, two ARCH models, and an autoregressive model for the volatility index VDAX. The ARCH models perform best in generating profits for market makers. Forecasts based on historical volatility also produce statistically and economically significant profits over the two-year simulation period of 1993 and 1994. In general, a filter rule with a small filter of 0.5 per cent produces the best results for both the ARCH models and historical volatility. However, the VDAX-AR model generates much lower and usually insignificant profits, and for some filter rules this model even has cumulative losses for market makers. For non-market-makers and non-members of exchange, however, larger transaction costs imply that no significant profits can be gained with any model of volatility forecasts.

Acknowledgements

Christian Schmitt gratefully acknowledges the financial support of the German Research Foundation (DFG), grant Ko 219/19-3. We wish to thank Deutsche Börse AG, Frankfurt/Main for providing the data. We are also grateful to Katinka Barysch and Elisabeth Hehn for helpful discussions and comments.

1 Introduction

In the Black-Scholes world of option pricing, volatility is the great unknown. The exercise price and the expiry date of the option are fixed in the contract, the spot price and the interest rate are observable, but the volatility of the underlying is not directly observable; it can only be estimated. Market practice has found two ways around this problem but both have their potential pitfalls.

The first approach to the estimation of volatility is to derive implied volatilities from actual option prices. According to the Black-Scholes model, implied volatility should be independent of the maturity and the moneyness of the option and implied volatility should be the same for call and put options on the same underlying. However, volatility smiles and skews are typically observed in practice and the term structure of volatility is usually not flat.¹ If these patterns are systematic then the Black-Scholes model is not applicable. Moreover, if one derives implied volatilities one relies on the efficient estimation of volatility by other participants and this can lead to circular reasoning.

The second approach is known as the estimation of historical volatility (HV). According to this approach, volatility is simply estimated as the standard deviation from a time series of the underlying's returns. This approach is perfectly consistent with the Black-Scholes model where it is assumed that the returns follow a Gaussian diffusion process. However, market practice often restricts the sample of returns to a relatively short period of a few weeks of daily data, or matches the time to maturity of the option that is to be priced with the sample length for volatility estimation.² This market practice is inconsistent. In the Black-Scholes model, volatility is a constant. It follows from first principles of statistics that a maximum sample size would be optimal. The fact that market participants do not want to use 'stale returns' reflects the view that volatility is not constant. But if volatility is variable and stochastic, then again, the Black-Scholes model cannot be applied. Furthermore, if market participants believe that volatility changes, then efforts should be directed at deriving a model of volatility dynamics and this model should be built into the option-pricing approach.

The Black-Scholes model of option pricing and its variants assume that the returns of the underlying asset follow a simple Wiener process in continuous time. The discrete-time equivalent to this assumption is a Gaussian white noise process. However, there is ample empirical evidence that returns in financial markets consistently violate this assumption. It is a strong and robust stylised fact of financial markets that short-run returns are heteroskedastic and have a leptokurtic distribution, i.e. both the

¹ See e.g. Tompkins (1994), Ch. 5.

² See e.g. Natenberg (1994).

distributional assumption and the assumption of independence are violated in practice. Heteroskedasticity, or volatility clustering, of course, confirms the market view that volatility varies over time in a systematic way.

In recent years, research has been directed at identifying models for financial-markets dynamics which capture these stylised facts. In particular, the class of ARCH³ models, first introduced by Engle (1982), has proved to be a promising and fruitful approach. Over the years, a plethora of ARCH-type models have been introduced but the GARCH model of Bollerslev (1986) and the EGARCH model of Nelson (1991) have established themselves as a kind of 'industry standard'. For surveys on ARCH models see Bollerslev, Chou and Kroner (1992), Bera and Higgins (1993) and Bollerslev, Engle and Nelson (1994).

Empirical applications confirm that these models capture the stylised facts of leptokurtosis and heteroskedasticity. Moreover, ARCH models tend to outperform other models in forecasting volatility. In particular, ARCH models have a more flexible and reasonable forecast function than historical volatilities (HV). The forecasts derived from HV estimates are constant for the whole forecast horizon. Furthermore, the HV estimator and forecast function attaches a uniform weight of 1 over the sample size (T) to all observations in the sample, whereas in an ARCH model, more recent observations would typically receive a larger weight than more distant observations.

However, the comparison of volatility forecasts has its problems. Since volatility is not observable, the forecast error cannot be directly determined. It is common practice to compare volatility forecasts with the squares of returns or with residuals from return equations but the squared (residual) returns are, strictly speaking, only volatility estimates from a sample of size 1. In this sense, these volatility forecast evaluations should perhaps be better described as data fitting exercises. Therefore, the adequate metric to evaluate the performance of ARCH models does not seem to be the simple forecasting competition.

In more recent years, attention has been directed to the implications of GARCH models for the pricing of options. Duan (1995) shows that a GARCH model is consistent with the often observed smile pattern of implied volatilities. Schmitt (1996) extends this analysis to EGARCH models and demonstrates that both smile and skew patterns can be obtained from an EGARCH option pricing model.

In this paper, we explore further the application of ARCH models to the pricing and trading of options. In particular, we implement volatility trading strategies, based on the GARCH model and the EGARCH model, for options on the German stock market

³ The acronym ARCH (autoregressive conditional heteroskedasticity) is used here as a generic name.

index DAX as the underlying. This approach can be viewed as an attempt to evaluate the forecasting performance of ARCH models in an adequate metric. Alternatively, this paper can be regarded as a contribution to the study of market efficiency for options on the German stock market.⁴ This aspect should also be of interest to practitioners who want to know whether excess profits can be obtained from applying ARCH models to option markets.

The rest of this paper is organised as follows. Section 2 provides a review of related literature and explains where this paper goes beyond previous studies. Section 3 introduces the GARCH and EGARCH models, explains their forecasting functions and reports parameter estimates. The alternative volatility trading strategies, based on historical, implied, GARCH and EGARCH volatilities, are explained in section 4. Section 5 provides a detailed description of the data and section 6 reports the empirical results from applying the trading rules. Section 7 summarises the main results and draws conclusions.

2 Literature Review

A few studies have attempted to explore the potential profitability of trading rules derived from ARCH models. Engle, Hong and Kane (1990) examine the profits from applications of GARCH option trading against the following three competitors: A moving average of HV, HV from the residuals of an AR(1) model for returns, and an ARMA(1,1) model for the residuals of an AR(1) model of returns. Volatility forecasts from these four models for the NYSE index are used as input into the Black-Scholes formula. Agents, representing one model each, trade on the basis of hypothetical one-day maturity put and call options. Reassuringly for both academics and practitioners, the GARCH trader is the clear winner. However, this study has a methodological flaw in that, again, the Black-Scholes formula is not applicable once the assumption of constant or deterministic volatility is dropped.

Engle, Kane and Noh (1993) adopt a similar methodology but, following Hull and White (1987), compute average option prices from 1000 simulated sample paths of the GARCH model. They also extend the previous analysis by including options with maturities of up to one year. The benchmark is the 300-day rolling HV. It turns out that trades based on Hull-White adjustments to GARCH prices are the most profitable strategies. Extending this strategy to longer maturities, however, would require frequent rebalancing of the option positions. The volatility trading strategy applied by this study is to buy an at-the-money straddle (i.e. a call and a put with the same exercise price) if option prices appear to be based on an underestimation of volatility and to sell an at-the-money straddle when volatility estimates of a rival

⁴ It should go without saying that the methodology developed in this paper can readily be adapted to other option markets.

model seem to be too high. An at-the-money straddle has approximately a delta value of zero, i.e. this is initially a pure volatility trading position. However, as the price of the underlying moves up (down), a long straddle position will become delta positive (negative) and vice versa for the selling of a straddle. To keep the straddle position (approximately) delta neutral, the initial straddle position needs to be closed and a new at-the-money straddle position has to be established. The fact that this re-balancing is missing in the study of Engle, Kane and Noh implies that, ultimately, the trading strategy is not a pure volatility strategy.

In a further study, Noh, Engle and Kane (1994) examine the profitability of a straddle strategy derived from a GARCH model against forecasts from autoregressive models based on implied volatilities. The underlying is the S&P 500 index and the options have maturities of more than 15 days. Again, the GARCH model turns out to provide the superior strategy.

Berglund, Hedvall and Liljebloom (1992) adopt a methodology which is similar to the one adopted by Engle et al. and apply it to a portfolio of shares of the Helsinki Stock Exchange. The competitor to the GARCH model is HV adjusted for serial correlation in returns. Here, too, the GARCH trader is the winner.

The main result that option trading strategies based on GARCH models produce significant and superior profitability is also confirmed by Kang and Brorsen (1995) for options on wheat futures and by Kroner and Levin (1996) for foreign currency options.

The methodologies adopted in these previous studies can and should be further improved in several respects and this is the starting point of our study. First, some studies use hypothetical options instead of traded options.⁵ The primary motivation for most studies seems to be to find an adequate metric for the forecasting performance of ARCH models rather than studying the efficiency of option markets. Second, for the evaluation of the profitability of trading strategies it is essential that real transaction costs are considered. Several studies neglect transaction costs altogether while others use assumed transaction costs. Third, none of the studies provides any details about the simultaneity and consistency of option prices and prices of the underlying. In general, it would be inadequate to use closing prices. Instead, time-stamped intra-day data should be used to make sure that a trading strategy was actually feasible and executable at the given prices. The precise aim of our study is to bring the simulation of trading strategies as closely as possible to the real world of the trading desk and to make sure that the strategy could actually have been implemented at the given prices.

⁵ It is unclear whether Kroner and Levin (1996) use hypothetical or traded options.

Finally, all previous studies applied the GARCH model to the simulation of option trading rules. However, the EGARCH model is often regarded as superior to the GARCH model since it can capture asymmetric volatility effects which can typically be found in stock returns (see for example Nelson (1991) and Schmitt (1996)). Empirical evidence shows that volatility tends to increase (decrease) with a fall (rise) in stock prices. There are various theoretical explanations for this stylised fact (see Nelson (1991)) which is often called the 'leverage effect' since one of these theories is based on the financial leverage of firms.

3 ARCH Estimates and Volatility Forecasts

The two most popular econometric volatility models are the GARCH model of Bollerslev (1986) and the EGARCH model of Nelson (1991). In the GARCH model it is assumed that the variance σ_t^2 , conditional on information available at time $t-1$, is a linear function of lagged conditional variances and of squared (residual) returns e_{t-i}^2 :

$$\sigma_t^2 = \alpha_0 + \alpha_1 \sigma_{t-1}^2 + \dots + \alpha_p \sigma_{t-p}^2 + \beta_1 e_{t-1}^2 + \dots + \beta_q e_{t-q}^2 \quad (1)$$

In most empirical applications it has proved to be sufficient to use the lag lengths of $p = 1$ and $q = 1$.

The EGARCH model has a different functional form and can be written as

$$\sigma_t^2 = \exp \left\{ a_0 + a_1 \ln \sigma_{t-1}^2 + b_{1a} \varepsilon_{t-1} + b_{1b} (|\varepsilon_{t-1}| - E[|\varepsilon_{t-1}|]) \right\} \quad (2)$$

where $\varepsilon_t = e_t / \sigma_t$ is Gaussian white noise with unit variance⁶. It follows that $E[|\varepsilon_t|]$, the expected value of a half-normal distribution, is equal to $\sqrt{2/\pi}$.⁷ The parameter b_{1a} captures the leverage effect. For "good news" ($\varepsilon_t > 0$), the impact of the innovation ε_{t-1} is $(b_{1b} + b_{1a})\varepsilon_{t-1}$ and for "bad news" ($\varepsilon_t < 0$) it is $(b_{1b} - b_{1a})\varepsilon_{t-1}$. If $b_{1a} = 0$, then $\ln \sigma_t^2$ responds symmetrically to ε_t . To produce a leverage effect, b_{1a} must be negative.

The functional form of the EGARCH model has several advantages compared to that of the simple GARCH model. First, the fact that an exponential form is used for the

⁶ More lags of ε_t and σ_t^2 could be used in the EGARCH model but lags of 1 are usually adequate. We, therefore, use the EGARCH(1,1) model throughout this paper.

⁷ See Johnson and Kotz (1970), p. 49.

conditional variance σ_t^2 guarantees that σ_t^2 is always positive. As a consequence, EGARCH models permit a wide range of variance effects that are not restricted by non-negativity constraints on the parameters. Second, instead of making σ_t^2 a function of squared (residual) returns e_t , as in the GARCH model, the EGARCH has the absolute value of a standard normal variable ε_t as an argument in the conditional variance function. This simplifies the determination of the moments of σ_t^2 and of the stationarity conditions for σ_t^2 . Third, ε_t enters into the conditional variance function in its level and as a deviation of the absolute value of ε_t from its expected value. Both components have, of course, an expected value of zero. The second component captures the size effect of shocks. Since the size effect enters in the form of an absolute value and not as a square, the volatility effects are dampened.⁸

We applied the GARCH model and the EGARCH model to a series of daily data of the German stock market index DAX. The sample has a size of 1000 and covers the period from January 1989 to January 1993. The DAX index includes 30 stocks selected with respect to market capitalisation and turnover. The great advantage of using the DAX index, instead of, say, the S&P500, is the fact that it is a performance index which adjusts not only for stock splits and capital changes but also for dividend payments. The shares are weighted by their share capital and the index is calculated with two decimals at one minute intervals.

Results of the estimation are reported in Table 1. Both the GARCH model and the EGARCH model achieve very satisfactory fits and all parameter estimates are highly significant. The estimate of b_{1a} of the EGARCH model confirms that there is a leverage effect in the DAX series. As expected for a series with volatility clustering, the estimates of α_1 and β_1 for the GARCH model and the estimates of a_1 and b_{1b} for the EGARCH model are positive. Since both $\alpha_1 + \beta_1$ and a_1 are smaller than 1, the stationarity conditions for σ_t^2 are satisfied in both models.

As a tool for model diagnostics we apply the Ljung-Box Statistics LBQ to the 'residuals' of the GARCH and EGARCH model.⁹ More specifically, we compute the LBQ for the absolute values of the 'residuals' at lag 20. The LBQs for both models are not significant at any reasonable level. We may conclude, therefore, that the heteroskedasticity of the DAX series is well captured by both models. Note, however, that the LBQ for the EGARCH model is lower than the LBQ for the GARCH model.

In order to compare the fit of both models we use the Schwarz Information Criterion (SIC). Table 1 shows that the SIC of the EGARCH model is lower than the SIC of the GARCH model. Therefore, the EGARCH model gives a better fit to the data than

⁸ However, the exponential function magnifies the impact effect for large shocks.

⁹ See Granger and Ding (1995).

the GARCH model.¹⁰ This might be due to the fact that the EGARCH model is able to capture the leverage effect whereas the GARCH model has a symmetric volatility function. Although both models appear to be adequate, the EGARCH model seems to have a certain edge over the GARCH model.

Table 1: Estimates of the GARCH and EGARCH model for the DAX

GARCH			EGARCH		
Coefficient	Estimate	t-statistic	Coefficient	Estimate	t-statistic
α_0	2.38×10^{-5}	6.81	a_0	-0.989	-5.31
α_1	0.644	14.71	a_1	0.885	42.91
β_1	0.253	11.52	b_{1a}	-0.080	-4.64
			b_{1b}	0.293	8.21
LBQ(20)	18.7		LBQ(20)	15.3	
SIC	-5972.2		SIC	-5982.7	

For the application in our trading strategies, only the 1-step forecasts of the GARCH(1,1) model

$$\hat{\sigma}_{t+1}^2 = \alpha_0 + \alpha_1 \sigma_t^2 + \beta_1 e_t^2 \quad (3)$$

and of the EGARCH(1,1) model

$$\sigma_t^2 = \exp\left\{a_0 + a_1 \ln \sigma_{t-1}^2 + b_{1a} \varepsilon_{t-1} + b_{1b} \left(|\varepsilon_{t-1}| - \sqrt{2/\pi}\right)\right\} \quad (4)$$

will be used.¹¹

4 Trading Strategies

The aim of this paper is to evaluate the profitability of applying different models to a pure volatility trading strategy. As noted above, a straddle would be an appropriate

¹⁰ This is confirmed in a more detailed analysis by Schmitt (1994).

¹¹ See Heynen, Kemna and Vorst (1994) for the more complicated formula for the multi-step forecasts.

strategy for being bullish (or bearish) on volatility and neutral on the market¹². However, market movements would imply that delta neutrality is lost since only at-the-money straddles are (approximately) delta neutral. It is market practice to tolerate relatively small deltas but, surely, large (cumulative) price changes should lead to a rebalancing of the position. Tompkins (1994) reports that in practice positions with an absolute delta smaller than 0.10 are considered to be delta neutral. We simulate market practice by assuming that a straddle is only 'rolled over' to a new exercise price if the price of the underlying surpasses one of the exercise prices.

The straddle position will also have sensitivities with respect to the interest rate (ρ) and the time to maturity (θ). However, these sensitivities ('Greeks') do not require monitoring and rebalancing. ρ is usually very small and, in fact, within the Black-Scholes framework the interest rate is assumed to be constant. Moreover, the decay of time value, as measured by θ , is inevitable.

The simple trading strategy to be adopted is to buy (sell) a straddle if the volatility forecast for the next day is larger (smaller) than the implied volatility. The position will be closed on the next day unless the strategy requires to establish exactly the same position with the same expiry date and the same exercise price. However, positions will be rolled over to the next expiry month if the time to maturity is less than 15 days. This rule is meant to protect against strong and erratic price movements towards the end of a contract caused by the unwinding of positions.

We use the volatility index VDAX¹³ as an estimate of implied volatilities of DAX options. The VDAX is based on the implied volatilities of options with two maturities which bracket the maturity of 45 calendar days. For each maturity value, call and put options with eight different exercise prices are used.¹⁴ The VDAX is updated every day at 1.30 pm.

¹² LIFFE in London has 12 recognised option strategies for being bullish (or bearish) on volatility and neutral on the market. Tompkins (1994, Ch. 7) recommends a long straddle or a long strangle for someone who expects volatility to rise and the price level to be stable. According to Tompkins (1994, p.240), "the strangle is preferred by those traders who wish to bet on increases of the implied volatility, while the straddle traders are betting on both an increase in the actual volatility (gamma effect) and the absolute impact of the implied volatility (vega effect)."

¹³ The VDAX is similar in its concept and construction to the S&P-100 volatility index (VIX) of the Chicago Board of Option Exchange (CBOE).

¹⁴ The VDAX is derived in two steps. In the first step, two sub-indices of the implied-volatility (σ_1^2 and σ_2^2) for options with a maturity of less than 45 days and more than 45 days are computed through a regression. In the second step a weighted average of the two sub-indices is calculated from

The trading strategy is to buy (sell) a straddle whenever an ARCH volatility forecast for the next day is larger (smaller) than the VDAX. In order to provide benchmarks for the profitability of the ARCH strategies, we enter two further volatility estimators into the contest. The first alternative strategy is based on HV. As explained above, it is common market practice to use the HV estimator with a rather small data window, especially for the pricing of options with short maturities. We mimic this market practice by using the 30-day HV computed from the closing prices of the DAX.¹⁵ The trading strategy is to go long (short) in an at-the-money straddle when the HV is larger (smaller) than the VDAX of the previous day at 1.30 pm.

The second rival strategy uses an autoregressive model of order 1, AR(1) for short, for the VDAX:

$$VDAX_t = \phi_0 + \phi_1 VDAX_{t-1} + u_t \tag{5}$$

where u_t is a random error term. The AR(1) process is stationary, or mean reverting, if $|\phi_1| < 1$. It is a well-established empirical regularity that stock-market volatilities follow a mean-reverting process (see e.g. French, Schwert and Stambaugh (1987) or Harvey and Whaley (1992)). Moreover, in models of stochastic volatility, such as Hull and White (1987), it is typically assumed that volatility is mean reverting. In fact, the AR(1) process is a linear approximation of the Ornstein-Uhlenbeck process often adopted in stochastic volatility models.

The parameters of the AR(1) model are estimated on a rolling basis with a data window of 250 trading days. Estimates of ϕ_1 were always in the range from 0.9 to 1.0 but strictly smaller than 1.0, confirming the property of mean reversion but showing a large degree of persistence which is typical for volatility clustering. The trading strategy for the AR(1) model is similar to that of the other strategies: buy (sell) an at-the-money straddle if the volatility forecast derived from (5) is larger (smaller) than the current value of the VDAX.

5 Description of Options Data

Options on the DAX index have been traded at the German Futures and Options Exchange (Deutsche Terminbörse, DTB) since August 1991. The DAX option has

$$VDAX = \sqrt{\frac{T_2 - 45}{T_2 - T_1} \sigma_1^2 + \frac{45 - T_1}{T_2 - T_1} \sigma_2^2}$$

where T_1 and T_2 are the calendar days to maturity of the two contract series.

¹⁵ The historical volatility computed from daily data is annualised with the factor $\sqrt{252}$.

now the largest trading volume of all options at the DTB. In 1995, the total trading volume was 24.3 million contracts which compares with 4.0 million contracts on the FT-SE 100 traded at LIFFE, 69.6 million contracts on the S&P100 traded at the CBOE, and 5.2 million contracts on the Nikkei 225 traded at the Osaka Stock Exchange.

The value of an option contract is equal to the current index level multiplied by 10 German marks (DEM). The tick size is 0.1 points with a corresponding tick value of DEM 1.00. The exercise prices have fixed increments of 25 index points. There are at least five option series for each expiry date: two are in the money, two are out of the money and one is (approximately) at the money. DAX options are European style and five different expiry months are always available with a maximum time to maturity of 9 months.¹⁶

The data set contains intra-day time-stamped transaction prices of call and put options. When a trading strategy generates a buy or sell signal, the transactions are assumed to have taken place at the first call and put prices after 11.00 am. However, only those prices are considered which are based on trades with at least 10 contracts in order to exclude small trades with, presumably, unfavourable bid-ask spreads. Moreover, only those contracts are selected which have a remaining maturity of between 15 and 45 days and which are nearest to the spot index value. This ensures sufficient liquidity and appropriate delta neutrality, respectively.

According to information provided by the DTB, normal transaction costs are DEM 4.00 per straddle but market makers have to pay only DEM 0.80 (see DTB (1994)). We will distinguish three groups of investors in this study: market makers, other registered traders (or traders for short), and non-members of the exchange. Further discussions with market participants revealed that the major players are usually able to negotiate mid-prices between the bid and ask quotes through 'prearranged' trades. We will, therefore, assume that the transaction costs are 0.80 for market makers and 4.00 for traders. According to market sources, the transaction costs for non-members are about 1 per cent of the straddle price. We assume that non-members have to pay DEM 8.00 in transaction costs since this would correspond approximately to 1 per cent of the average straddle premium of DEM 780 over the two year period from 1993 to 1994. We can, therefore, be confident that our trade simulations are based on realistic assumptions and that the trades could have been executed at the given prices and transaction costs.

Interest rate effects of volatility trading are considered in the following way. Purchases of options and losses from trading have to be financed whereas income from sales and profits from trading can be invested. To a large extent the two effects

¹⁶ In March 1996, the DTB introduced the DAX-XXL options with maturities of up to 24 months.

will cancel but if there are net interest-rate effects then profits and losses will increase. As an approximation, we use the 1-day interest rate but we neglect the bid-ask spread in interest rates because it would have only a minute impact.

6 Empirical Results

For all trading strategies, the VDAX of the previous day, released at 1.30 pm, is the benchmark. ARCH forecasts and historical volatilities are based on daily returns up to 11.00 am of the day of trading whereas the VDAX-AR generates buy (sell) signals if at 1.30 pm on the previous day, the forecast for the next day is larger (smaller) than the current VDAX.


Trading strategies are simulated on a daily basis from January 1993 to December 1994 with a total of 493 trading-day data available. Figure 1 plots the cumulative profits and losses for market makers from the adoption of the four rival strategies. It is apparent from the figure that the strategies tended to suffer net losses in the early part of the two-year period, especially the strategies based on EGARCH and HV. However, in the second part of 1993, these strategies strongly recovered and then continued to produce profits in 1994. At the end of 1994, the EGARCH and the HV strategy were almost level with an impressive cumulative profit of 288 per cent for EGARCH and 307 per cent for HV whereas the VDAX-AR model trailed with a profit of 154 per cent. The highest overall return was achieved by the GARCH model with a net profit of 348 per cent. For most of the two-year period, the HV profits and the GARCH profits are very similar in magnitude and it is only in the last two months that the GARCH model outperforms the HV forecasts. It is also noteworthy that in the first six months or so, the ARCH strategies and the HV strategy quite often generated different trading signals which is reflected in the different developments of cumulative profits for these strategies.

The frequencies of trades also differs between the strategies. The GARCH model generated trading signals on all 493 trading days, whereas the EGARCH model and the HV generated trading signals on all but one day. On the other hand, the VDAX-AR model did not induce a position on 32 days. It should be noted that the VDAX is reported with two decimals and that all forecasts were also rounded to two decimals. Therefore, trades were not initiated every day.

Since volatility forecasts are subject to error, it would seem reasonable to adopt a filter for trading strategies. A filter of 2.0 per cent, say, would imply that a buy (sell) signal is only generated if the volatility forecast for the next day is 2.0 per cent higher (lower) than the current VDAX. This tends to reduce the number of transactions. In fact, with a 2.0 per cent filter, the ARCH and HV forecasts produce buy or sell signals on between 234 and 293 days. An application of a 2.0 per cent filter rule to VDAX-AR forecasts, however, would imply that no transactions would be initiated

since the estimate of ϕ_1 is always close to 1 and, therefore, the forecast of tomorrow's VDAX is close to the current VDAX.


Figure 1: Cumulative profits of trading strategies without filters for market makers



Obviously, overall profits will also be affected by the introduction of filter rules. As Figure 2 indicates, the cumulative profits are substantially reduced for all three trading strategies when a 2.0 filter is introduced for the ARCH and HV forecasts and a 0.2 filter is applied to the VDAX-AR model. A smaller filter is required for the VDAX-AR model since this model is based on the same data, i.e. implied volatilities, as the benchmark. Therefore, the difference between forecasts and benchmark is smaller for this model than for the other two models. At the end of the two-year trading period, the cumulative profits are at 341 per cent for the GARCH model, 157 per cent for the EGARCH model, 136 per cent for the HV model, and 24 per cent for the VDAX-AR model.

The overall pattern of profitabilities in Figure 2 looks similar to the one in Figure 1 but the number of trades is substantially reduced and, apparently, many profit opportunities are lost through the application of these filters. The only model which retains its profit level in spite of the filter is the GARCH model. Compared with the zero-filter application, the cumulative profit falls only by 8 per cent although the number of transactions falls from 493 to 293.

Figure 2: Cumulative profits of trading strategies with filters for market makers



The fact that profits depend on the filter size raises the question of what the best filter would be. Table 2 tries to give an answer. It reports the number of trade signals generated by the four models and the cumulative profits as a function of the filter size. Note that the first number in the filter column applies to the ARCH models and the HV model, whereas the second filter is used for the VDAX-AR model. It is obvious that an increase in filter size will reduce the number of trades. The GARCH forecasts generate more trading signals than any other forecast and the GARCH model tends to dominate the other models in terms of profitability. However, the highest profit is realized with the EGARCH model and a filter of 0.5 per cent where the cumulative return is a very impressive 363 per cent.

The VDAX-AR model is clearly inferior to the other three models. It produces the best return of 154 per cent with a no-filter rule but would produce overall losses with two filter sizes. On the other hand, the ARCH models and the HV model show profits for all filter rules and achieve the best performance with filters of 0.5 or 1.0. Note, too, that even with a much lower filter, the VDAX-AR model induces much fewer trades than the other models.

Table 2: Trading signals and profits with different filters for market makers

Filter*	Number of trade signals**			
	GARCH	EGARCH	HV	VDAX-AR
0.0	493 (255/238)	492 (267/225)	492 (176/316)	461 (191/270)
0.5 / 0.05	448 (234/214)	401 (229/172)	432 (147/285)	217 (89/128)
1.0 / 0.10	402 (209/193)	297 (173/124)	374 (122/252)	123 (47/76)
1.5 / 0.15	350 (183/167)	234 (140/94)	306 (85/221)	73 (37/36)
2.0 / 0.20	293 (157/136)	165 (111/54)	250 (58/192)	15 (7/8)
2.5 / 0.25	222 (125/97)	108 (78/30)	185 (41/144)	4 (3/1)
3.0 / 0.30	154 (98/56)	68 (56/12)	121 (13/108)	1 (0/1)
Filter*	Cumulative profits in per cent			
	GARCH	EGARCH	HV	VDAX-AR
0.0	347.8	287.9	307.1	154.4
0.5 / 0.05	349.4	362.8	336.3	-0.1
1.0 / 0.10	350.7	297.5	217.5	22.7
1.5 / 0.15	336.3	257.8	185.0	58.1
2.0 / 0.20	341.4	157.0	136.4	23.5
2.5 / 0.25	239.9	89.4	70.0	3.3
3.0 / 0.30	129.8	64.4	49.6	-0.1

* Note: The first filter applies to the ARCH models and HV whereas the second filter applies to the VDAX-AR model.

** Note: The numbers in brackets denote the number of long resp. short trade signals.

Although the overall performance of the models looks impressive, it remains to be shown that the returns are statistically and economically significant. The results from a simple *t*-test for the mean of daily returns, against the null hypothesis of zero, are reported in Table 3. It shows the empirical significance levels for a one-tailed test. Again, the VDAX-AR model performs disappointingly. For the daily returns none of the filter rules are significantly different from zero. On the other hand, the ARCH and HV forecasts generate significant profits for low-filter strategies when a conventional significance level of 5 per cent is applied. The strongest performance comes from the GARCH model where the profits are significantly positive for six of the seven filters.

Table 3: Statistical and economic significance of profits for market makers

Filter*	Empirical significance levels in per cent			
	GARCH	EGARCH	HV	VDAX-AR
0.0	2.05	4.58	3.60	17.94
0.5 / 0.05	1.61	1.18	1.92	50.05
1.0 / 0.10	1.41	1.82	6.56	40.97
1.5 / 0.15	1.33	2.80	8.34	24.46
2.0 / 0.20	0.95	8.05	10.63	19.19
2.5 / 0.25	1.81	12.89	22.89	36.78
3.0 / 0.30	6.54	17.56	27.02	84.13
Filter*	Sharpe ratios			
	GARCH	EGARCH	HV	VDAX-AR
0.0	1.41	1.15	1.23	0.60
0.5 / 0.05	1.48	1.56	1.42	-0.08
1.0 / 0.10	1.51	1.43	1.02	0.07
1.5 / 0.15	1.53	1.30	0.92	0.39
2.0 / 0.20	1.62	0.92	0.81	0.29
2.5 / 0.25	1.42	0.70	0.44	-0.67
3.0 / 0.30	0.98	0.54	0.33	-79.20

* Note: The first filter applies to the ARCH models and HV whereas the second filter applies to the VDAX-AR model.

In order to evaluate the economic significance of the trading strategies, we computed Sharpe ratios for the models as reported in the lower panel of Table 3. The Sharpe ratio is defined as the ratio of the excess return of a strategy to its corresponding standard deviation. The excess return is derived by subtracting the average one-day interest rate from the average annualised return. Table 2 shows that returns tend to decrease with an increase of filter size. It should be intuitive that an increase of filter size will also reduce the volatility of returns. The net effect, however, is a general decrease of the Sharpe ratio for the EGARCH model and the HV model with increasing filter size whereas the Sharpe ratio is relatively stable for the GARCH model. Overall, the GARCH model outperforms all other models in terms of Sharpe ratios and, interestingly, scores the highest ratio with a filter of 2.0. The EGARCH model comes second in the ranking of economic significance since it outperforms the HV model for almost all filter sizes except the zero size. In comparison, the VDAX-AR model trails with markedly lower Sharpe ratios and three are even negative. An overall benchmark can be provided by the Sharpe ratio for a buy-and-hold strategy for the spot-market DAX. This ratio is 0.59.

The overall verdict is that the ARCH forecasts and HV forecasts are able to generate statistically and economically significant profits for market makers, especially when low filters are applied. But, in general, GARCH forecasts outperform the other forecasts of volatility. On the other hand, the VDAX-AR model performs rather poorly. It requires much lower filters to induce any trades and it does not seem to be able to produce significant profits. With some filters it even generated cumulative losses.

The previous analysis was based on the assumption that transaction costs of DEM 0.80 have to be paid per straddle. These are the transaction costs for market makers. Other traders and non-members, however, have to incur higher transaction costs of DEM 4.00 and DEM 8.00, respectively. These higher transaction costs will obviously reduce the profits of the volatility strategies. We assume, however, that the number of trades is not affected by the higher transaction costs. In practice, a larger filter would presumably be applied to balance higher transaction costs.

Table 4: Profits with different filters for traders and non-members

Filter*	Cumulative profits in per cent for traders			
	GARCH	EGARCH	HV	VDAX-AR
0.0	94.0	14.6	53.8	-86.2
0.5 / 0.05	117.9	136.0	113.2	-118.4
1.0 / 0.10	138.8	128.7	19.5	-44.3
1.5 / 0.15	149.8	116.5	27.7	22.4
2.0 / 0.20	170.1	50.3	9.1	15.0
2.5 / 0.25	102.5	13.0	-26.6	0.9
3.0 / 0.30	30.1	15.9	-20.3	-0.6
Filter*	Cumulative profits in per cent for non-members			
	GARCH	EGARCH	HV	VDAX-AR
0.0	-223.4	-327.0	-262.9	-386.9
0.5 / 0.05	-171.5	-147.6	-165.7	-266.3
1.0 / 0.10	-126.1	-82.3	-228.1	-128.0
1.5 / 0.15	-83.4	-60.2	-168.8	-22.2
2.0 / 0.20	-43.9	-83.1	-150.0	4.2
2.5 / 0.25	-69.4	-82.4	-147.5	-2.0
3.0 / 0.30	-94.6	-44.7	-107.7	-1.3

* Note: The first filter applies to the ARCH models and HV whereas the second filter applies to the VDAX-AR model.

A comparison between Table 4 and Table 2 shows in fact that for traders and non-members cumulative profits are substantially lower than for market makers. Although

the number of trades is not affected by the transaction costs, the overall pattern of profits and losses for traders and non-members differs from that for market makers in Table 2. Among the four models it is again the GARCH model which achieves the highest returns. The most successful strategy would have been to apply a GARCH model with a filter of 2.0 per cent and this strategy would have made a profit of 170 per cent. For traders using EGARCH and HV, the filter of 0.5 per cent is the best filter but profits are substantially lower at 136 per cent and 113 per cent, respectively. Traders using the VDAX-AR, however, would make losses with four of the seven filter sizes.

Table 4 also highlights how sensitive the overall profits and losses are to transaction costs. Non-members, who pay the highest transaction costs, would fare much worse than traders. In almost all cases they would make overall losses, the only exception being the very modest profit of 4.2 per cent for a filter of 0.2 per cent with the VDAX-AR. This profit does not have any statistical or economic significance. In fact, the return of this strategy is lower than the risk-free interest rate.

Table 5 shows that, according to our simulations, traders would also be unable to reap significant profits. The upper panel of the table reports the statistical significance level of a one-tailed test. No entry is lower than the notional level of 5 per cent. Sharpe ratios, which measure the performance and economic significance of the strategies, are shown in the lower panel of Table 5. A comparison with the Sharpe ratios for market makers in Table 3 shows that traders would have to be content with a much reduced performance of their strategies. The GARCH model with a filter of 2 per cent has the best Sharpe ratio of 0.78. However, for an EGARCH strategy the highest obtainable Sharpe ratio of 0.59 just matches the Sharpe ratio of a buy-and-hold strategy for the spot-market DAX.

The implications of our simulation results are that for traders and non-members the hypothesis of market efficiency cannot be rejected. For these groups it is not possible to earn (risk-adjusted) excessive profits from implementing volatility trading strategies based on one of the four strategies. On the other hand, it appears that market makers are able to make very large and economically significant profits, as reflected in the high Sharpe ratios. It is not quite clear, however, whether these profits are a sign of market inefficiency. After all, market makers are obliged to quote continuously bid and ask prices in order to ensure sufficient market liquidity. Therefore, they are subject to substantial position risks. It has to remain an open question whether the large profits of market makers are just a compensation for this risk or whether the profits reflect market inefficiency. A more detailed analysis of market micro-structure would be required to answer this question which would go beyond the scope of the present study.

Table 5: Statistical and economic significance of profits for traders

Filter*	Empirical significance levels in per cent			
	GARCH	EGARCH	HV	VDAX-AR
0.0	29.0	46.58	37.67	69.60
0.5 / 0.05	23.4	19.65	24.30	85.55
1.0 / 0.10	19.1	18.10	44.64	67.26
1.5 / 0.15	16.0	19.21	41.81	39.40
2.0 / 0.20	11.9	32.58	46.70	28.58
2.5 / 0.25	18.4	43.36	61.08	45.94
3.0 / 0.30	36.2	40.72	59.79	84.13
Filter*	Sharpe ratios			
	GARCH	EGARCH	HV	VDAX-AR
0.0	0.34	0.01	0.17	-0.42
0.5 / 0.05	0.46	0.55	0.44	-0.84
1.0 / 0.10	0.57	0.59	0.03	-0.41
1.5 / 0.15	0.65	0.56	0.08	0.09
2.0 / 0.20	0.78	0.24	-0.02	0.07
2.5 / 0.25	0.56	0.01	-0.29	-0.89
3.0 / 0.30	0.15	0.04	-0.29	-15.05

* *Note:* The first filter applies to the ARCH models and HV whereas the second filter applies to the VDAX-AR model.

7 Summary and Conclusions

This paper uses a data set and applies a methodology that make it possible to simulate volatility-trading strategies which could have been applied in practice. Special care has been taken to use simultaneous intraday prices and realistic transaction costs. Furthermore, straddle positions were evaluated on a daily basis to make sure that positions, which were initially delta neutral, were, if necessary, readjusted to new strike prices to preserve delta neutrality. A further methodological improvement over previous studies is the application of the EGARCH model, alongside the GARCH model, to capture the well-known leverage effect of stock market returns. Volatility trading strategies were applied to options on the German stock market index DAX. These option contracts, traded at the DTB in Frankfurt, are not only among the most popular stock-market derivatives world-wide but also have the advantage of being based on a true performance index. Option strategies are based on three approaches to volatility estimation: the conventional methods of historical volatility and implied volatility (as measured by the VDAX index) and the more recent ARCH models (which subsume the GARCH model and the EGARCH model).

Of the four rival models considered in this paper, the GARCH model, in general, performs best in generating significant profits from buying or selling straddles when

its volatility forecasts indicate that volatility will be higher or lower than currently anticipated by the market. Forecasts based on the EGARCH model and on historical volatility also produce statistically and economically significant profits for market makers over the two-year simulation period of 1993 and 1994 but the performance is often inferior to that of the GARCH model. The fact that the 'simpler' GARCH model outperforms the EGARCH model, which can accommodate asymmetric volatility reactions, is quite surprising. However, for analysts of financial markets this result should not appear to be unparalleled. In fact, one finds again and again that more sophisticated models are superior in in-sample comparisons but these models are often outperformed in out-of-sample forecasts by simpler and more robust models. The prime example is presumably the random-walk model which, despite its simplicity, is often difficult to beat in forecasting contests.

In practice, small differences between current volatility estimates and volatility forecasts would not trigger trades. We, therefore, experimented with different filter rules. It turns out that a filter rule of 2 per cent is usually optimal for the GARCH model whereas a smaller filter of 0.5 per cent produces the best results for the EGARCH model and for HV strategies. The fourth model, an autoregressive model for the volatility index VDAX, generates much lower and usually insignificant profits from volatility trading. Filters have to be much smaller for the VDAX-AR model but for some filter rules this model even has cumulative losses.

These results can be viewed in two ways. Practitioners will certainly welcome a new approach which promises to be profitable, at least for market makers, in a real-world situation. Academics, on the other hand, might have to conclude from the large volatility trading profits obtainable for market makers, especially with ARCH models, that this options market is not efficient in exploiting the available price data, i.e. the hypothesis of weak-form market efficiency might have to be rejected. However, at the current stage, we cannot be sure because the large profits might just compensate for the position risk market makers have to bear when providing continuous bid and ask quotes.

For the two other groups of market participants considered in this paper, traders and non-members, higher transaction costs imply that their profits from volatility trading rules are substantially reduced. In fact, traders would not be able to make profits that are statistically or economically significant whereas non-members would even have to incur cumulative losses in almost all combinations of models and filter rules.

Our approach can be extended in various directions. Most obviously, it can be applied to other markets and other underlying assets. It would also be interesting to consider other trading strategies such as butterflies, condors, and strangles. However, all applications should be implemented under the basic premise of our study: to simulate trading strategies under assumptions and conditions which are as close as possible to actual market practice.

References

- Bera, Anil K. / Higgins, Matthew L. (1993):** *ARCH Models: Properties, Estimation and Testing*, Journal of Economic Surveys, Vol. 7, 305-366.
- Berglund, Tom / Hedvall, Kaj / Liljebloom, Eva (1992):** *Predicting Volatility of Stock Indexes for Option Pricing on a Small Security Market*, Swedish School of Economics and Business Administration, Helsinki, mimeo.
- Bollerslev, Tim (1986):** *Generalized Autoregressive Conditional Heteroscedasticity*, Journal of Econometrics, Vol. 31, 307-327.
- Bollerslev, Tim / Chou, Ray Y. / Kroner, Kenneth F. (1992):** *ARCH Modelling in Finance - A Review of Theory and Empirical Evidence*, Journal of Econometrics, Vol. 52, 5-59.
- Bollerslev, Tim / Engle, Robert F. / Nelson, Daniel B. (1994):** *ARCH Models*, in: Engle, Robert F. / McFadden, Daniel L. (eds.): *Handbook of Econometrics*, Vol. 4, Elsevier, Chapter 49, 2961-3038.4
- DTB Deutsche Terminbörse (1994):** *Regelwerk*, Frankfurt am Main.
- Duan, Jin-Chuan (1995):** *The GARCH Option Pricing Model*, Mathematical Finance, Vol. 5, No. 1, 13-32.
- Engle, Robert F. (1982):** *Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation*, Econometrica, Vol. 50, No. 4, 987-1007.
- Engle, Robert F. / Hong, C.-S. / Kane, Alex (1990):** *Valuation of Variance Forecasts with Simulated Option Markets*, NBER Working Paper No. 3350.
- Engle, Robert F. / Kane, Alex / Noh, Jaesun (1993):** *Index-Option Pricing with Stochastic Volatility and the Value of Accurate Variance Forecasts*, Discussion Paper 93-43, Department of Economics, University of California, San Diego.
- French, Kenneth R. / Schwert, G. William / Stambaugh, Robert F. (1987):** *Expected Stock Returns and Volatility*, Journal of Financial Economics, Vol. 19, 3-29.
- Granger, Clive W.J. / Ding, Zhuanxin (1995):** *Some Properties of Absolute Return - An Alternative Measure of Risk*, Annales d'Economie et de Statistique, Vol. 40, 67-92.
- Harvey, Campbell R. / Whaley, Robert E. (1992):** *Market Volatility Prediction and the Efficiency of the S&P 100 Index Option Market*, Journal of Financial Economics, Vol. 31, 43-73.

- Heynen, Ronald / Kemna, Angelien / Vorst, Ton (1994):** *Analysis of the Term Structure of Implied Volatilities*, Journal of Financial and Quantitative Analysis, Vol. 29, No. 1, 31-56.
- Hull, John / White, Alan (1987):** *The Pricing of Options on Assets with Stochastic Volatilities*, The Journal of Finance, Vol. 42, 281-300.
- Johnson, Norman L. / Kotz, Samuel (1970):** *Distributions in Statistics - Continuous Univariate Distributions I*, Wiley & Sons.
- Kang, Taehoon / Brorsen, B. Wade (1995):** *Conditional Heteroskedasticity, Asymmetry, and Option Pricing*, The Journal of Futures Markets, Vol. 15, No. 8, 901-928.
- Kroner, Kenneth F. / Levin, Asriel (1996):** *Creating and Evaluating Volatility Trading Strategies: GARCH versus Neural Networks*, Proceedings of the Third International Conference "Forecasting Financial Markets", London, 27 - 29 March 1996.
- Natenberg, Sheldon (1994):** *Option Volatility & Pricing: Advanced Trading Strategies and Techniques*, Probus.
- Nelson, Daniel B. (1991):** *Conditional Heteroskedasticity in Asset Returns: A New Approach*, *Econometrica*, Vol. 59, 347-370.
- Noh, Jaesun / Engle, Robert F. / Kane, Alex (1994):** *Forecasting Volatility and Options Prices of the S&P 500 Index*, Journal of Derivatives, Fall, 17-30.
- Schmitt, Christian (1994):** *Volatilitätsprognosen für deutsche Aktienkurse mit ARCH- und Markov-Mischungsmodellen*, Discussion Paper No. 94-07, Zentrum für Europäische Wirtschaftsforschung, Mannheim.
- Schmitt, Christian (1996):** *Option Pricing Using EGARCH Models*, in: **Albrecht, Peter (ed.):** *Aktuarielle Ansätze für Finanz-Risiken*, Vol. 2, VVW, 1311-1336.
- Tompkins, Robert (1994):** *Options Explained*², Macmillan.