The Skew Pattern of Implied Volatility in the DAX Index Options Market

Silvia Muzzioli*

Abstract

The aim of this paper is twofold: to investigate how the information content of implied volatility varies according to moneyness and option type and to compare option-based forecasts with historical volatility. The different information content of implied volatility is examined for the most liquid at-the-money and out-of-the-money options: put (call) options for strikes below (above) the current underlying asset price. Two hypotheses are tested: unbiasedness and efficiency of the different volatility forecasts. The investigation is pursued in the Dax index options market, by using synchronous prices matched in a one-minute interval. It was found that the information content of implied volatility has a humped shape, with out-of-the-money options being less informative than at-the-money ones. Overall, the best forecast is at-the-money put implied volatility: it is unbiased (after a constant adjustment) and efficient, in that it subsumes all the information contained in historical volatility.

Keywords: Implied Volatility, Volatility Smile, Volatility Forecasting, Option Type.
JEL classification: G13, G14.

* Department of Economics and CEFIN, University of Modena and Reggio Emilia, Viale Berengario 51, 41100 Modena (I), e-mail: silvia.muzzioli@unimore.it. The author wishes to thank the Editor and the anonymous referee for helpful comments and suggestions. The author gratefully acknowledges financial support from MIUR and Fondazione Cassa di Risparmio di Modena. Usual disclaimer applies.

1 – Introduction

Black-Scholes implied volatility is a forward-looking measure of the expected volatility between a given time and the expiration of the option. Even if theoretically the Black-Scholes model postulates a constant volatility, in empirical terms implied volatility varies according to the option's strike price, describing a smile or skew, depending on the shape of the relation. As it is often necessary to have implied volatilities that correspond to strike prices that are not traded in the market, implied volatilities are usually interpolated (e.g. by cubic splines) in order to obtain a smile or skew function. The skew function is fundamental both for the construction of option-implied trees (see e.g. Derman and Kani (1994)) that are used to price and hedge exotic options, and for the computation of a number of market volatility indexes (see e.g. CBOE VIX, for the Chicago Board Options Exchange, or the VDAX-New for the German equity market).

The recent turmoil in the financial markets following the sub-prime crisis has clearly highlighted the important role of market volatility indexes in the detection and anticipation of market stress. These indexes are highly correlated with future market volatility and with risk factors embedded in credit spreads of sovereign debt: as such they are deemed to capture what is known as market "fear".

Numerous papers have investigated the forecasting power of Black-Scholes volatility versus a time-series volatility forecast (we refer the interested reader to Poon (2005), that examines 93 studies on the issue of volatility forecasting, and concludes that predictions based on implied volatility are on average superior to time-series volatility models). However, as far as we know, there is little evidence about the different information content of implied volatilities extracted from options with different strike price and type (call or put), that are used in the computation of the smile function. As for the strike price dimension, Ederington and Guan (2005), in the S&P 500 options market, show that the information content of implied volatilities varies roughly in a mirror image of the implied volatility smile. As for the option type dimension, Fleming (1998) and Christensen and Hansen (2002), in the S&P 100 options market, find that at-the-money call implied volatility has slightly more predictive power than put implied volatility. These two studies use American-type options on a dividend paying index: the early exercise feature and the dividend yield estimation tend to influence call and put option prices differently, and may have altered the comparison if not properly addressed.

Even if call and put implied volatilities extracted from an option with the very same strike price and time to maturity should theoretically be the same due to no-arbitrage considerations, there are empirically many reasons that may cause call and put implied volatilities to differ. These reasons are amplified if call and put options are compared in a different strike price dimension. First of all, converting option prices into implied volatilities leads to measurement errors (stemming from finite quote precision, bid-ask spreads, non-synchronous observations and other measurement errors): small errors in any of the input may produce large errors in the implied volatility (see e.g. Hentshle (2003)). This is also documented by Fleming (1999) who observes that deviations of call and put option prices from no-arbitrage values do not necessarily signal market inefficiency, but are rather due to transaction costs and other market imperfections. Along the same line of reasoning, the no-arbitrage replication of a put or a call through put-call parity implies going short (long) on the underlying asset. Unlike the long side, the short side usually requires an initial margin, and is exposed to margin calls if the underlying asset price begins to rise.

Second, the demand for put options is inherently different from demand for call options. Put options are used for portfolio insurance purposes, in particular by institutional investors. Rubinstein (1994) finds that out-of-the-money put implied volatilities are usually higher than both in-the-money put and out-of-the-money call implied volatilities due to the crash phobia that became widespread after October 1987. This hedging pressure has been documented both in relation to different moneyness classes, and also in the moneyness category and may lead the implied volatilities of options whose price is impacted by hedging pressure to be less informative about future market volatility.

Last, call and put option volumes are very different. Put options are usually traded for a wider strike price interval, and they are also more frequently traded than call options if compared in the same moneyness class (see e.g. Buraschi and Jackwerth (2001). Bollen and Whaley (2004) provide evidence that the demand for at-the-money put options is much higher than for at-the-money call options. As implied volatility is a forward-looking estimate of future realised volatility, we expect actively traded options to be more informative than less traded options. This has been documented in various papers that have analysed index options markets. Trading volume can be considered as an indicator of investors' information: as pointed out in Donaldson and Kamstra (2005) when trading volumes are high, the forecasting power of implied volatility is also high. Sarwar (2005) finds a positive relation between trading volume and implied volatility, since implied volatility is determined by the activity of informed traders who may prefer options markets rather than stock markets in order to benefit from lower transaction costs and higher leverage.

The aim of this paper is twofold: to investigate how the information content of implied volatility varies according to moneyness and option type, and to compare option-based forecasts with historical volatility in order to see if they subsume all the information contained in historical volatility. The different information content of implied volatility is examined for the most liquid at-the-money and out-of-the-money options: put (call) options for strikes below (above) the current underlying asset price, i.e. the ones that are generally used as inputs for the computation of the smile function. This investigation is important for the understanding of the role of the different ingredients of the smile function and can be seen as a preliminary exercise in order to choose different weights for each volatility input in a volatility index. In particular, for at-the-money volatilities, that are widely used by market participants and are usually inserted in the smile function by computing an average of both call and put implied ones, we investigate whether one option class provides a better forecast of future realised volatility and whether a combination of the two adds substantial benefit. Two hypotheses are tested: unbiasedness and efficiency of the different volatility forecasts with respect to historical volatility. Historical volatility is measured by both lagged realised volatility and a GARCH (1,1) forecast. The investigation is pursued in the DAX-index options market, chosen for two main reasons: the options are European, therefore the estimation of an early exercise premium is not needed and cannot influence the results; the DAX-index is a capital-weighted performance index composed of 30 major German stocks and is adjusted for dividends, stock splits and changes in capital: dividends are assumed to be reinvested in the shares and they do not affect the index value.

This paper makes at least three contributions to the ongoing debate on the performance of option-based volatility forecasts and their efficiency with respect to historical volatility. First, unlike previous studies, that use settlement prices, it uses a rich data set consisting of the more informative synchronous prices between the options and the underlying asset. This is important to stress, since our implied volatilities are real “prices”, as determined by synchronous no-arbitrage relations. Moreover, the choice of the DAX-index option market avoids the estimation of the early exercise premium and of the dividend yield and makes our data set less prone to measurement error.

Second, since the forecasting performance of implied volatilities extracted from options with different strike price and type has not been extensively tested, our results can contribute to an understanding of how the information content of implied volatility varies with moneyness and type. This feature has potential implications for all the fields in which volatility measures can be used: portfolio selection models, derivative pricing models, hedging and risk management techniques in general.

Third, the results of the present paper are important for an assessment of which option class is best to include in a market volatility index. In fact, many market volatility indexes, such as the VIX or the V-DAX New, have changed their computational methodology from a formula based on only at-the-money implied volatilities to a formula that uses the information content of the whole smile function: what is known as model-free implied volatility (see e.g. Britten Jones and Neuberger (2000)). Criticisms of the use of the whole smile function can be found in various papers (see e.g. Taylor et al. (2006), Becker at al. (2007)). Among others, Andersen and Bondarenko (2007) find that implied volatility measures obtained from a narrow corridor of strikes, closely related to the concept of at-the-money implied volatility, perform better in predicting future realised volatility than broader corridor-implied volatility measures, such as model-free implied volatility. If the most informative option class is still the at-the-money one, it follows that this class should have a greater weight than other option classes in the volatility index.

The plan of the paper is as follows. Section 2 illustrates the data set, the sampling procedure and the definition of the variables. Section 3 describes the methodology used in order to address the unbiasedness and efficiency of the different volatility forecasts. Sections 4 and 5 report the results of univariate and augmented regressions respectively, and assess the relative performance of the different volatility forecasts (at-the-money and out-of-the-money call and put implied volatilities, lagged realised volatility and GARCH (1,1)). Section 6 investigates the forecasting performance of a combination of at-the-money call and put implied volatilities. The last section concludes.

2 - The Data set and the definition of the variables

The data set[1] consists of intra-daily data on DAX-index options, recorded from 19 July 1999 to 31 December 2005. Each record reports the strike price, expiration month, transaction price, contract size, hour, minute, second and centisecond. As for the underlying asset we use intra-daily prices

[1] The data source for DAX-index options and the DAX index is the Institute of Finance, Banking, and Insurance of the University of Karlsruhe (TH). The risk-free rate is available on Data-Stream.

of the DAX-index recorded in the same time period. As a proxy for the risk-free rate we use the one-month Euribor rate.

DAX-options started trading on the German Options and Futures Exchange (EUREX) in August 1991. They are European options on the DAX-index, which is a capital-weighted performance index composed of 30 major German stocks and is adjusted for dividends, stocks splits and changes in capital. Dividends are assumed to be reinvested in the shares and they do not affect the index value. Moreover, the fact that the options are European avoids the estimation of the early exercise premium. This feature is important since our data set is by construction less prone to estimation error if compared to the majority of previous studies that use American-style options.

Several filters are applied to the option data set. First, we eliminate option prices that are less than one euro, since the closeness to the tick size may affect the true option value. Second, in order not to use stale quotes, we eliminate options with trading volumes of less than one contract. Third, as it is standard practice in the literature to estimate the smile by using only the more liquid at-the-money and out-of-the-money options, following Jiang and Tian (2005) we eliminate in-the-money options (call options with moneyness[2] $(X/S) < 0{,}97$ and put options with moneyness $(X/S) > 1{,}03$). Fourth, we eliminate option prices violating standard no-arbitrage bounds. Finally, in order to reduce computational burden, we only retain options that are traded between 3.00 and 4.00 p.m, (the choice is motivated by the level of trading activity in this interval).

As for the sampling procedure, in order to avoid the telescoping problem described in Christensen, Hansen and Prabhala (2001), we use monthly non-overlapping samples. In particular, we collect the prices recorded on the Wednesday following the expiry of the option (third Saturday of the expiry month) because the week immediately following the expiry date is one of the most active. These options have a fixed maturity of almost one month (from 17 to 22 days to expiry). If the Wednesday is not a trading day, we move to the trading day immediately after.

Implied volatility is computed separately for out-of-the-money and at-the-money call and put prices. We start from the cleaned data set of option prices that is composed of at-the-money and out-of-the-money call and put prices recorded from 3.00 to 4.00 p.m. We compute call and put implied volatilities by using synchronous prices, matched in a one-minute interval, by inverting the Black-Scholes formula. Implied volatilities are grouped into four sets depending on the option's moneyness and type and averaged in order to

[2] Moneyness is defined as *X/S*, where *X* is the strike price and *S* is the underlying asset.

obtain four implied volatility estimates: at-the-money call (ATMC) implied volatility (σ_{ATMC}), at-the-money put (ATMP) implied volatility (σ_{ATMP}), out-of-the-money call (OTMC) implied volatility (σ_{OTMC}), out-of-the-money (OTMP) implied volatility (σ_{OTMP}) (OTMC if $(X/S) > 1,03$, ATMC e ATMP if $0,97 \leq (X/S) \leq 1,03$, OTMP if $(X/S) < 0,97$).

Unlike Ederington and Guan (2005), who use settlement prices, we are using the more informative synchronous prices, matched in a one-minute interval. This is important to stress, since our implied volatilities are real "prices", as determined by no-arbitrage relations. As a result it is unlikely to have observations for each day for all the 12 categories of moneyness that Ederington and Guan (2005) use, since most of the trading concentrates on at-the-money and close-to-the-money options. Therefore, in order to avoid the case in which one option class is empty (that is found in Ederington and Guan (2005)) and to have a simple clear-cut comparison between at-the-money and out-of-the-money options, we chose to examine much broader classes than Ederington and Guan (2005).

Implied volatility is an ex-ante forecast of future realised volatility on the time period until the option expiry. Therefore we compute realised volatility (σ_R) in month *t,* as the sample standard deviation of the daily index returns over the option's remaining life:

$$\sigma_R = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(R_i - \overline{R})^2}$$

R_i is the return of the DAX-index on day i and $\overline{R}$ is the mean return of the Dax-INDEX in month t. We annualize the standard deviation by multiplying it by $\sqrt{252}$.

In order to examine the predictive power of implied volatility versus historical volatility, following Christensen and Prabhala (1998) and Jorion (1995) we chose to use two different time series volatility forecasts: lagged realized (LR), i.e. one month before, volatility (σ_{LR}) and a GARCH (1,1) (GAR) forecast (σ_{GAR}). Using daily data on the DAX index, the GARCH (1,1) variance equation is defined as:

$$\sigma_{t+1}^2 = a_0 + a_1 R_t^2 + b_1 \sigma_t^2 ,$$

where R_t is the de-meaned DAX-index return on day t (for more details see Bollerslev (1986)). Given the length of our sample period (similar to that of Jorion (1995) and Fleming (1998)) and our sampling procedure, that avoids telescoping samples, it is not possible to split the sample period into two sub-periods (one for estimating the model parameters and one for computing the out-of-sample forecasts) without reducing too much the number of observations for the other implied volatility forecasts. Therefore, as in Jorion (1995) and Fleming (1998), the GARCH model was estimated via maximum likelihood over the entire data set. Following Fleming (1998) the GARCH forecast (σ_{GAR}) of the average volatility over the life of the option is defined as:

$$\sigma_{GAR} = \sqrt{\frac{1}{T-t-1}\sum_{j=1}^{T-t}\tilde{\sigma}^2_{t+j|t}} \ ,$$

where $\tilde{\sigma}^2_{t+j|t}$ is the forecast at time t of the variance j days into the future, and T is the maturity of the option. We annualize the standard deviation by multiplying it by $\sqrt{252}$. The GARCH forecast, estimated over the entire sample period, benefits from information that is not available to other forecasts.

Descriptive statistics for volatility and log volatility series are reported in Table 1. It may be seen that on average realized volatility is lower than the implied volatility estimates (except for out-of-the-money call implied volatility), with on average put implied volatility that is higher than call implied volatility. This skew pattern, depicted in Figure 1[3], is typical for index options and is consistent with the crash-phobia explanation, since the demand for out-of-the-money put options to hedge against downside risk pushes implied volatility to rise at low strikes. As for the standard deviation, realised volatility is slightly more volatile than the implied volatility estimates (except for out-of-the-money put implied volatility). The volatility series are highly skewed (long right tail) and leptokurtic. In line with the literature (see e.g. Jiang and Tian (2005)) we use the natural logarithm of the volatility series instead of the volatility itself in the empirical analysis for the following reasons. First log-volatility series conform more closely to normality than pure volatility series: this is documented in various papers and it is the case in our sample (see Table 1). Second, natural logarithms are less likely to be affected by outliers in the regression analysis.

[3] In the graph ATMP implied volatility is to the left of ATMC implied volatility because at-the-money call (put) implied volatility is mainly obtained from options with $1 < X/S \leq 1{,}03$ ($0{,}97 \leq X/S < 1$), since these are the most traded strike price intervals.

Figure 1. The skew pattern of implied volatility.

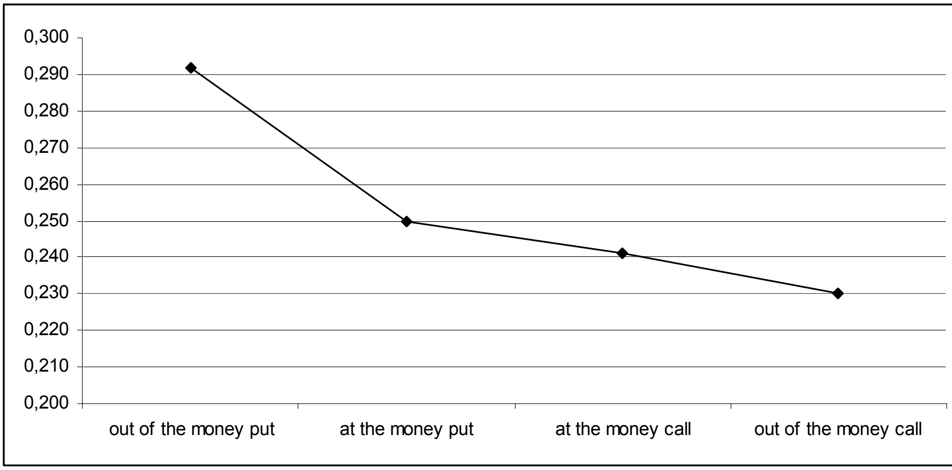

Table 1. Descriptive statistics for volatility and log-volatility series.

Statistic	σ_{ATMC}	σ_{OTMC}	σ_{ATMP}	σ_{OTMP}	σ_{R}	σ_{LR}	σ_{GAR}
Mean	0,241	0,230	0,250	0,292	0,238	0,239	0,240
std dev	0,111	0,100	0,109	0,134	0,127	0,125	0,110
Skewness	1,748	1,658	1,560	1,873	1,245	1,255	1,520
Kurtosis	6,137	5,590	5,560	6,440	3,976	4,003	4,740
Jarque Bera	70,770	56,800	52,300	83,050	23,250	23,450	39,190
p-value	0,000	0,000	0,000	0,000	0,000	0,000	0,000
	ln(σ_{ATMC})	**ln(σ_{OTMC})**	**ln(σ_{ATMP})**	**ln(σ_{OTMP})**	**ln(σ_{R})**	**ln(σ_{LR})**	**ln(σ_{GAR})**
Mean	-1,506	-1,565	-1,465	-1,311	-1,558	-1,550	-1,507
std dev	0,395	0,383	0,386	0,381	0,486	0,482	0,395
Skewness	0,644	0,692	0,543	0,873	0,376	0,357	0,693
Kurtosis	3,277	3,197	2,972	3,512	2,340	2,374	2,888
Jarque Bera	5,576	6,263	3,790	10,622	3,220	2,899	6,218
p-value	0,062	0,044	0,150	0,005	0,199	0,235	0,045

3 - The methodology

The information content of implied volatility is examined both in univariate and in augmented regressions. In univariate regressions, realized volatility is regressed against one of the six volatility forecasts in order to examine the different predictive power of each forecast. The univariate regressions are the following:

$$\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i) \qquad (1)$$

where σ_R = realized volatility and σ_i= volatility forecast, *i=ATMC, OTMC, ATMP, OTMP, LR, GAR*.

In augmented regressions, realized volatility is regressed against two or more volatility forecasts in order to distinguish which one has the highest explanatory power. We chose to compare first pair-wise one volatility forecast with a time series volatility forecast in order to see if implied volatility subsumes all the information contained in historical volatility. The augmented regressions used are the following:

$$\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i) + \gamma \ln(\sigma_j) \qquad (2)$$

where σ_R = realized volatility, σ_i= implied volatility, *i= ATMC, OTMC, ATMP, OTMP* and σ_j = *LR, GAR*.

Moreover, we compare pair-wise the four implied volatility forecasts in order to understand whether the information carried by one option class is more valuable than the information carried by the other:

$$\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i) + \gamma \ln(\sigma_j) \qquad (3)$$

where σ_R = realized volatility, σ_i= *ATMC, OTMC, ATMP, OTMP* and σ_j= *ATMC, OTMC, ATMP, OTMP, i≠j*.

We also compare the two time-series volatility forecasts, in order to see which one has the highest forecasting power on future realised volatility:

$$\ln(\sigma_R) = \alpha + \beta \ln(\sigma_{LR}) + \gamma \ln(\sigma_{GAR}) \qquad (4)$$

Following Christensen and Prabhala (1998) there are three hypotheses tested in univariate regressions (1). The first hypothesis concerns the amount of information about future realized volatility contained in the volatility forecast. If the volatility forecast contains some information, then the slope coefficient should be significantly different from zero. Therefore we test whether $\beta = 0$ and we see whether it can be rejected. The second hypothesis is about the unbiasedness of the volatility forecast. If the volatility forecast is an unbiased estimator of future realised volatility, then the intercept should be zero and the slope coefficient should be one (H_0: $\alpha = 0$ and $\beta = 1$). In case this hypothesis is rejected, we see if at least the slope coefficient is equal to one (H_0: $\beta = 1$) and, if not rejected, following Jiang and Tian (2005) we interpret the volatility forecast as unbiased after a constant adjustment. Finally if implied volatility is efficient, then the error term should be white noise and uncorrelated with the information set.

In augmented regressions (2) there are two hypotheses to be tested. The first is about the efficiency of the volatility forecast: we test whether the implied volatility (ATMC, OTMC, ATMP, OTMP) forecast subsumes all the information contained in historical volatility. If the forecast is efficient, then the slope coefficient of historical volatility should not be significant, (H_0: $\gamma = 0$). Moreover, as a joint test of information content and efficiency we test in equations (2) if the slope coefficients of historical volatility and implied volatility (ATMC, OTMC, ATMP, OTMP) are equal to zero and one respectively (H_0: $\gamma = 0$ and $\beta = 1$). Following Jiang and Tian (2005), we ignore the intercept in the null hypothesis, and if our null hypothesis is not rejected, we interpret the volatility forecast as unbiased after a constant adjustment.

Moreover, we investigate the different information content of each option class with respect to the others. To this end we test, in augmented regressions (3), whether $\gamma = 0$ and $\beta = 1$, or $\gamma = 1$ and $\beta = 0$, in order to see if the implied volatility of one option class subsumes all the information contained in the other. Finally we test, in augmented regression (4), whether $\gamma = 0$ and $\beta = 1$, or $\gamma = 1$ and $\beta = 0$, in order to see if one time series volatility forecast subsumes all the information contained in the other.

Unlike other papers (see e.g. Christensen and Prabhala 1998, Christensen and Hansen (2002)) that use American options on dividend paying indexes, our data set of European-style options on a non-dividend paying index avoids measurement errors that may arise in the estimation of the dividend yield and the early exercise premium. Moreover, we carefully cleaned the data set and we use synchronous prices. Nonetheless, as we are averaging different implied volatilities in a single class, some measurement errors may still affect our estimates. Therefore we adopt an instrumental variable procedure (IV), we regress implied volatility in each class on an instrument (in univariate regressions) and on an instrument and any other exogenous variable (in augmented regressions) and replace fitted values in the original univariate and augmented regressions. As instrument for implied volatility in each class, we use both LR volatility, GAR, and past implied volatility in the same class, as they are possibly correlated to the true implied volatility, but unrelated to the measurement error associated with implied volatility one month later. As an indicator of the presence of errors in variables, we use the Hausman (1978) specification test statistic[4].

4 - The results of univariate regressions

The results of the OLS univariate regressions (equation (1)) are reported in Table 2 (p-values in parentheses). In all the regressions the residuals are normal, homoscedastic and not autocorrelated (the Durbin Watson statistic is not significantly different from two and the Breusch-Godfrey LM test confirms non-autocorrelation up to lag 12[5]).

First of all, in all the univariate regressions all the beta coefficients are significant at the 1% level: this means that all the six volatility forecasts contain some information about future realised volatility. Among the two time series volatility forecasts, GAR performs much better than LR volatility: this is not surprising, since GAR has been estimated on the entire data set and therefore benefits from information that is not available for other forecasts. Overall put implied volatility obtains a better performance than call implied volatility.

The adjusted R^2 is the highest for ATMP implied volatility, followed by OTMP implied, and by ATMC and GAR, that obtain a similar performance. LR volatility and OTMC implied volatility have the lowest adjusted R^2. If we plot the R^2 against the option's moneyness (bearing in mind that ATMP (ATMC) implied volatility is mainly obtained from options with $1 < X/S \leq 1{,}03$ ($0{,}97 \leq X/S < 1$), since these are the most traded strike prices, we find the pattern depicted in Figure 2.

---

[4] The Hausman specification test is defined as: $m = \dfrac{\left(\hat{\beta}_{IV} - \hat{\beta}_{OLS}\right)^2}{Var(\hat{\beta}_{IV}) - Var(\hat{\beta}_{OLS})}$ where: $\hat{\beta}_{IV}$ is the beta obtained through the TSLS procedure, $\hat{\beta}_{OLS}$ is the beta obtained through the OLS procedure and *Var(x)* is the variance of the coefficient *x*. The Hausman specification test is distributed as a $\chi^2(1)$.

[5] In the regression that includes as an explanatory variable lagged realised volatility, the Durbin's alternative was computed. The results confirmed the non-autocorrelation of the residuals. The results of the Durbin's alternative and of the Breusch-Godfrey LM test are available on request.

Table 2. OLS univariate regressions.

Dependent variable: log realized volatility

Independent variables

Inter.	$\ln(\sigma_{ATMC})$	$\ln(\sigma_{OTMC})$	$\ln(\sigma_{ATMP})$	$\ln(\sigma_{OTMP})$	$\ln(\sigma_{LR})$	$\ln(\sigma_{GAR})$	Adj. R^2	DW	X^2	H test
-0,002	$1{,}03^{+++}$						0,70	1,97	3,16	8,50
(0,99)	(0,00)								(0,21)	
0,083		$1{,}05^{+++}$					0,68	1,95	0,40	10,59
(0,53)		(0,00)							(0,81)	
0,057			$1{,}10^{+++}$				0,76	1,94	13,95	2,11
(0,60)			(0,00)						(0,00)	
-0,123				$1{,}09^{+++}$			0,73	1,74	75,84	5,78
(0,24)				(0,00)					(0,00)	
-0,302					$0{,}81^{+++}$***		0,64	2,19	7,50	
(0,01)					(0,00)				(0,02)	
-0,002						$1{,}03^{+++}$	0,70	2,17	2,92	
(0,98)						(0,00)			(0,23)	

Note: The numbers in brackets are the p-values. The χ^2 report the statistic of a χ^2 test for the joint null hypothesis $\alpha = 0$ and $\beta = 1$ (p-values in parentheses) in the following univariate regressions: $\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i)$, where σ_R = realized volatility and σ_i= volatility forecast i=ATMC, OTMC, ATMP, OTMP, LR, GAR. The superscripts ***, **, * indicate that the slope coefficient is significantly different from one at the 1%, 5%, and 10% critical levels respectively. The superscripts $^{+++}$, $^{++}$, $^{+}$ indicate that the slope coefficient is significantly different from zero at the 1%, 5%, and 10% critical levels respectively. The last column shows the Hausman (1978) specification test statistic (one degree of freedom) 5% critical level = 3,841.

The results highlight the fact that the information content of implied volatility has a humped shape, with out-of-the-money options being less informative than at-the-money ones. This is consistent with the hedging pressure argument documented in Bollen and Whaley (2004): out-of-the-money options are less informative than at-the-money ones. Unlike the results in Ederington and Guan (2005) the forecasting power of implied volatility does not vary in a mirror image of the implied volatility smile: rather it exactly follows the volatility skew pattern, the only exception being OTMP implied volatility that has a smaller forecasting power than it should have by looking at the skew. The difference can be attributed to the fact that, with respect to Ederington and Guan (2005), our option classes are broader[6] and our results are based on synchronous prices, matched in a one-minute interval.

The null hypothesis that the volatility forecast is an unbiased estimate of future realized volatility is not rejected for both call implied volatility forecasts (ATMC and OTMC) and for GAR. However, it is rejected for both put implied volatility forecasts (ATMP and OTMP). This is probably due to the fact that, in our sample, realized volatility is on average much lower than both ATMP and OTMP implied volatility forecasts. However, the null hypothesis β=1 cannot be rejected even at the 10% critical level for both put implied volatility forecasts. Therefore also ATMP and OTMP implied volatilities can be considered as unbiased after a constant adjustment given by the intercept of the regression. LR volatility obtains the worst performance: it is not unbiased even after a constant adjustment.

Figure 2. The adjusted R^2 for different moneyness classes.

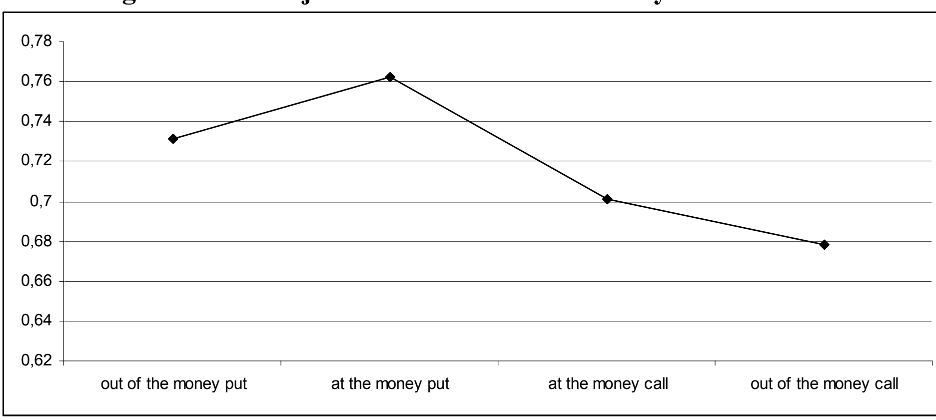

Finally, in order to test our results for robustness, and see whether implied volatility was measured with errors, we adopt an instrumental variable procedure and run a two-stage least squares. The Hausman (1978) specification test reported in the last column of Table 2 indicates that the error in variables problem is not significant only for ATMP. Therefore we report in Table 3 the TSLS regressions. As expected, the TSLS estimates of the beta coefficients are higher than the OLS estimates. Nonetheless, the results are virtually the same as the OLS case, with ATMC and OTMC being unbiased

[6] The choice was made in order to avoid having samples of different length, caused by missing observations for some dates.

and ATMP and OTMP being unbiased after a constant adjustment. Therefore, the forecasting power of each volatility forecast is not substantially changed with respect to the OLS case.

Table 3. TSLS univariate regressions.

Dependent variable: log realized volatility							
Independent variables							
Inter.	**ln(σ_{ATMC})**	**ln(σ_{OTMC})**	**ln(σ_{ATMP})**	**ln(σ_{OTMP})**	**Adj. R^2**	**DW**	**X^2**
0,185	$1{,}16^{+++}$*				0,69	2,11	6,08
(0,18)	(0,00)						(0,05)
0,328		$1{,}21^{+++}$**			0,66	2,14	4,62
(0,04)		(0,00)					(0,10)
0,118			$1{,}15^{+++}$*		0,76	1,97	15,36
(0,31)			(0,00)				(0,00)
-0,01				$1{,}18^{+++}$**	0,73	1,81	77,25
(0,93)				(0,00)			(0,00)

Note: The numbers in brackets are the p-values. The χ^2 report the statistic of a χ^2 test for the joint null hypothesis $\alpha = 0$ and $\beta = 1$ (p-values in parentheses) in the following univariate regressions: $\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i)$, where σ_R = realized volatility and σ_i= volatility forecast, i=ATMC, OTMC, ATMP, OTMP. The superscripts ***, **, * indicate that the slope coefficient is significantly different from one at the 1%, 5%, and 10% critical levels respectively. The superscripts $^{+++}$, $^{++}$, $^{+}$ indicate that the slope coefficient is significantly different from zero at the 1%, 5%, and 10% critical levels respectively.

5 - The results of augmented regressions

The results of the OLS augmented regressions (equations (2), (3) and (4)) are reported in Table 4. In all the regressions the residuals are normal, homoscedastic and not autocorrelated (the Durbin Watson statistic is not significantly different from two and the Breusch-Godfrey LM test confirms non-autocorrelation up to lag 12[7]).

[7]In the regressions that include as explanatory variable lagged realised volatility, the Durbin's h-statistic should be computed. The Durbin's h-statistic is defined as:

$$h = \left(1 - \frac{1}{2}d\right)\sqrt{\frac{T}{1 - TVar(\hat{\beta})}},$$

where d is the Durbin-Watson statistic, T is the sample size and $Var(\hat{\beta})$ is the estimated variance of the regression coefficient of the lagged dependent variable. However, as in all the

In augmented regressions (2), we compare each implied volatility forecast with historical volatility in order to see whether any of the implied volatility forecasts is efficient, i.e. if it subsumes all the information contained in historical volatility. For historical volatility we use both LR volatility and GAR. As the results are similar, in the following we use the term historical volatility, without specifying the forecasting method. The results differ somehow across option type: overall put implied volatilities are more efficient than call implied ones. Only ATMP implied volatility is clearly efficient. In fact, we cannot reject the null hypothesis that the slope coefficient of historical volatility is equal to zero even at the 10% level for both LR volatility and GAR, indicating that ATMP implied volatility subsumes all the information contained in historical volatility. Moreover, from the comparison of univariate and augmented regressions, the inclusion of historical volatility does not improve the goodness of fit according to the adjusted R^2.

We cannot reject the null hypothesis that the slope coefficient of ATMP implied volatility is equal to one even at the 10% level and the joint test of information content and efficiency $\gamma = 0$ and $\beta = 1$ does not reject the null hypothesis, indicating that ATMP implied volatility is efficient and unbiased after a constant adjustment. OTMP implied volatility is marginally inefficient, since the coefficient of historical volatility is significant only at the 5% level, and the joint test of information content and efficiency $\gamma = 0$ and $\beta = 1$ does not reject the null hypothesis only at the 1% level. For ATMC and OTMC implied volatilities the results are quite similar, with ATMC performing slightly better. In both cases, from the comparison of univariate and augmented regressions, the inclusion of historical volatility improves the goodness of fit according to the adjusted R^2. In fact, the slope coefficient of historical volatility is highly significant and the joint test of information content and efficiency $\gamma = 0$ and $\beta = 1$ rejects the null hypothesis (the only exception being ATMC implied volatility with respect to LR volatility at the 1% level).

In order to see if any one of the implied volatilities subsumes all the information contained in the others, we test in augmented regressions (3) if $\gamma = 0$ and $\beta = 1$ or $\gamma = 1$ and $\beta = 0$.

---

regressions $TVar(\hat{\beta}) > 1$, the test is not applicable. The results of the Breusch-Godfrey LM test are available on request.

Table 4. Augmented regressions.

Dependent variable: log realized volatility

Independent variables

Inter.	**ln (σ_{ATMC})**	**ln (σ_{OTMC})**	**ln (σ_{ATMP})**	**ln (σ_{OTMP})**	**ln (σ_{LR})**	**ln (σ_{GAR})**	**Adj. R^2**	**DW**	**$X^{2\,a}$**	**$X^{2\,b}$**	**H test**
-0,009	$0,70^{+++}$				$0,32^{+++}$		0,73	2,26	7,39		2,907
(0,12)	(0,14)				(0,12)				(0,02)		
0,045		$0,65^{+++}$			$0,37^{+++}$		0,72	2,28	11,21		3,019
(0,13)		(0,14)			(0,11)				(0,00)		
0,048			$0,95^{+++}$		0,14		0,76	2,08	3,38		0,931
(0,11)			(0,15)		(0,12)				(0,18)		
-0,930				$0,82^{+++}$	$0,25^{++}$		0,74	2,04	6,34		1,455
(0,10)				(0,15)	(0,12)				(0,04)		
0,088	$0,56^{+++}$					$0,53^{+++}$	0,74	2,23	11,42		0,016
(0,12)	(0,16)					(0,16)			(0,00)		
0,134		$0,50^{+++}$				$0,60^{+++}$	0,73	2,22	11,51		0,122
(0,12)		(0,16)				(0,16)			(0,00)		
0,09			$0,86^{+++}$			0,26	0,77	2,11	4,40		0,003
(0,11)			(0,18)			(0,17)			(0,11)		
0,001				$0,69^{+++}$		$0,43^{+++}$	0,75	2,07	9,01		0,015
(0,11)				(0,16)		(0,16)			(0,01)		
-0,003					0,004	$1,027^{+++}$	0,69	2,17	22,68	0,167	
(0,13)					(0,23)	(0,27)			(0,00)	(0,92)	
-0.00	$1,02^{++}$	0,01					0,69	1,97	0,19	6,07	3,512
(0,13)	(0,43)	(0,44)							(0,91)	(0,04)	
0,031	$-0,91^{++}$		$2,02^{+++}$				0,78	1,90	26,12	7,77	0,949
(0,10)	(0,39)		(0,40)						(0,00)	(0,02)	
-0,014	$0,42^{++}$			$0,70^{+++}$			0,75	1,83	14,77	7,29	2,288
(0,11)	(0,18)			(0,18)					(0,00)	(0,03)	
-0,001		-0,36	$1,45^{+++}$				0,76	1,89	28,94	3,89	1,671
(0,11)		(0,27)	(0,27)						(0,00)	(0,14)	
0,035		$0,40^{++}$		$0,74^{+++}$			0,75	1,84	22,80	8,13	2,101
(0,11)		(0,16)		(0,16)					(0,00)	(0,02)	
0,04			$0,85^{+++}$	0,26			0,77	1,89	3,12	12,47	1,766
(0,11)			(0,26)	(0,26)					(0,21)	(0,00)	

Note: The numbers in brackets are the standard errors. The χ^{2a}, χ^{2b} report the statistic of a χ^2 test for the joint null hypothesis $\gamma = 0$ and $\beta = 1$ or $\gamma = 1$ and $\beta = 0$ (p-values in parentheses) in the following regressions: $\ln(\sigma_R) = \alpha + \beta \ln(\sigma_i) + \gamma \ln(\sigma_j)$, where σ_R = realized volatility, σ_i= volatility forecast i= ATMC, OTMC, ATMP, OTMP, LR, GAR and σ_j= volatility forecast j, j= ATMC, OTMC, ATMP, OTMP, LR, GAR, *i≠j*. The superscripts $^{+++}$, $^{++}$, $^{+}$ indicate that the slope coefficient is significantly different from zero at the 1%, 5%, and 10% critical levels respectively. The last column shows the Hausman (1978) specification test statistic (one degree of freedom) 5% critical level = 3,841.

By looking at the significance of the coefficients and at the results of the χ^2 test, we can see that ATMP implied volatility subsumes all the information contained in both OTMP and OTMC implied volatilities. ATMC implied volatility subsumes all the information contained only in OTMC implied volatility. The comparison of ATMP and ATMC implied volatilities is not straightforward since the coefficient of ATMC is significant only at the 5% level and the χ^2 test marginally rejects the null hypothesis for ATMP at the 5% level. In order to better understand the performance of the two at-the-money implied volatility forecasts, we compute the Diebold and Mariano test statistic (for more details see Diebold and Mariano (1995)). The loss function chosen is the absolute error loss. The Diebold and Mariano test statistic under the null of equal predictive accuracy is distributed as a *N(0,1)*, in our case the test statistic is -2,35, therefore we can reject the null of equal predictive accuracy at the 5% level. Based on these results we can say that ATMP implied volatility has a slightly better predictive power than ATMC implied volatility. Therefore, in our sample, at-the-money put options are priced more efficiently than at-the-money call options, probably due to the larger trading volume, determined by a higher demand.

Finally, in order to identify the best time-series forecast, we test in augmented regression (4) if $\gamma = 0$ and $\beta = 1$ or $\gamma = 1$ and $\beta = 0$. The results highlight that GAR subsumes all the information contained in LR volatility.

Since the volatility forecasts compared in augmented regressions (2), (3), (4) are highly correlated with future realized volatility, the question naturally arises about how the results are affected by multicollinearity. It is standard in the literature on volatility forecasting (see e.g. Poon and Granger (2003) to use regression-based methods for examining the information content of different measures of volatility, either obtained from option prices or from time-series data. It is also expected that, as the different volatility measures are designed to forecast the same unobservable quantity, they will track one another fairly closely.

The encompassing tests illustrated in Section 5 are based on a procedure derived from Fair and Schiller (1990) that is aimed at determining whether one forecast contains different information from the other. There are basically four cases: if the two forecasts contain independent information, then the two slope coefficients will be different from zero. If both forecasts contain information, but the information of the first forecast is contained in the information of the second forecast, and the second forecast contains further relevant information, then only the slope coefficient of the second forecast will be significantly different from zero and we say that the second forecast subsumes all the information contained in the first forecast. If both forecasts contain the same information, then they are perfectly correlated and the two slope coefficients can not be determined. In this case we face the problem of perfect multicollinearity.

High correlation between the explanatory variables need not necessarily cause a problem. Whether or not it is a problem depends on how big the standard errors are and if the t-ratios are significant. If the standard errors are not too inflated and the t-ratios are significant, then the near-multicollinearity problem might not be serious. Moreover, it is important to stress that even in the presence of near multicollinearity, the OLS estimators remain BLUE (still consistent, unbiased and efficient).

In Table 4, overall the standard errors are not too big (in relation to the magnitude of the slope coefficients). As a result, in every augmented regression, either both explanatory variables are statistically significant, or at least one of the two variables is significant. Therefore, either the two forecasts contain independent information, or one forecast subsumes the information contained in the other. Based on these results, we can conclude that in our case the near-multicollinearity problem is not serious.

As a last step, in order to test our results for robustness and see if implied volatility was measured with errors, we adopt an instrumental variable procedure and run a two-stage least squares. The Hausman (1978) specification test reported in the last column of Table 4 indicates that the errors-in-variables problem is significant neither in augmented regressions (2) nor in augmented regressions (3)[8]. Therefore the OLS regression results are reliable.

---

[8] In augmented regressions (3) the instrumental variables procedure is used for the most significant variable in each regression.

6 - A combination of call and put at-the-money volatilities

At-the-money volatilities are widely used by market participants. Call and put at-the-money volatilities are usually inserted in the smile function by computing an average of both option classes. Given that prices are observed with measurement errors (stemming from finite quote precision, bid-ask spreads, non-synchronous observations and other measurement errors) small errors in any of the input may produce large errors in the implied volatility. Quoting Hentshle (2003): "Unfortunately many authors preclude the cancellation of errors across puts and calls by using only the more liquid out-of-the-money options. Unless underlying asset prices and dividend rates are observed with high precision, this practice can result in a substantial loss of efficiency". Moreover, as noted in Moriggia, Muzzioli and Torricelli (2007) the use of both call and put options in the volatility estimation greatly improves the pricing performance of option pricing models based on implied binomial trees.

Therefore, in this section we investigate how to combine at-the-money call and put implied volatilities in a single estimate, in order to convey the information from both call and put prices and cancel possible errors across option type. In the logarithmic specification, the natural candidates for the weights that we may assign to call and put implied volatilities would be the estimated coefficients of augmented regression (3). However, as the beta coefficient of call implied volatility is not significantly different from zero, it is not possible to find an optimal combination of the two with constant weights over time.

In line with the approach by Christensen and Hansen (2002) who propose favouring the most actively traded options, we construct a weighted average of ATMC and ATMP implied volatilities (σ_M), where the weights are the relative trading volume of each option class on the total trading volume:

$$\sigma_M = \frac{\sigma_{ATMC} V_c + \sigma_{ATMP} V_p}{V_c + V_p}$$

where V_i is the trading volume of option in class *i*, *i=c,p*. The weighting rule favours the most actively traded options, that in our sample are the put ones.

Descriptive statistics of average implied volatility and log average implied volatility are reported in Table 5. Average implied volatility is slightly higher than realised volatility. Similarly to the results in Table 1, we can see that the natural logarithm of the volatility series conforms more to normality than the volatility series. Therefore it will be used as an explanatory variable in univariate and augmented regressions.

Table 5. Descriptive statistics for average implied volatility.

Statistic	σ_M	$\ln(\sigma_M)$
mean	0,25	-1,48
std dev	0,11	0,39
skewness	1,67	0,59
kurtosis	5,97	3,14
Jarque Bera	64,15	4,52
p-value	0,00	0,10

In order to analyse the performance of average implied volatility, we run both univariate and augmented regressions (1), (2) and (3)[9] with $\sigma_i=\sigma_M$. Furthermore, in order to test for robustness our results, we look for possible errors in variables. The results are reported in Table 6. In all the regressions the residuals are normal, homoscedastic and not autocorrelated (the Durbin Watson statistic is not significantly different from two and the Breusch-Godfrey LM test confirms non autocorrelation up to lag 12[10]).

In univariate regression (1), the beta coefficient of average implied volatility is highly significant, but the null hypothesis that average implied volatility is an unbiased estimate of future realized volatility is rejected at the 5% level. The null hypothesis that β is equal to one cannot be rejected even at the 10% critical level: therefore we can consider average implied volatility as unbiased after a constant adjustment given by the intercept of the regression.

In augmented regressions (2) we compare average implied volatility with historical volatility in order to ascertain whether average implied volatility subsumes all the information contained in historical volatility. The results provide evidence for both the unbiasedness and efficiency of average implied volatility forecast with respect to LR volatility. The evidence is less clear-cut if we measure historical volatility with GAR, since the joint test of information content and efficiency $\gamma=0$ and $\beta=1$ marginally rejects the null hypothesis. Moreover, if we compare the performance of average implied

[9] In augmented regression 3 we compare average implied only with ATMP implied, since we are looking for an improvement over the best forecast.

[10] In the regression that includes as explanatory variable lagged realised volatility, the Durbin's h-statistic should be computed. However, in this case, the test is not applicable. The results of the Breusch-Godfrey LM test are available on request.

volatility with ATMP, we see that the adjusted R^2 is lower for average implied volatility.

Table 6. OLS and TSLS regressions of realised volatility on average implied volatility.

PANEL A: OLS REGRESSIONS

Dependent variable: log realized volatility

Independent variables

Inter.	ln (σ_M)	ln (σ_{ATMP})	ln (σ_{LR})	ln (σ_{GAR})	Adj. R^2	DW	X^2	$X^{2\,a}$	$X^{2\,b}$	H. test
0,040	1,08				0,74	1,97	7,87			5,280
(0,71)	(0,00)						(0,02)			
0,029	0,84***		$0,22^{+}$		0,74	2,18		4,47		2,319
(0,79)	(0,00)		(0,07)					(0,11)		
0,096	0,71***			$0,40^{++}$	0,75	2,19		6,72		0,433
(0,39)	(0,00)			(0,02)				(0,04)		
0,025	$-1,75^{++}$	$2,85^{+++}$			0,78	1,89		15,27	7,52	0,019
(0,81)	(0,02)	(0,00)						(0,00)	(0,02)	

PANEL B: TSLS REGRESSION

Dependent variable: log realized volatility

Independent variables

Inter.	ln (σ_M)	Adj. R^2	DW	X^2
0,16	1,16***	0,73	2,05	10,37
(0,20)	(0,00)			(0,01)

Note: The numbers in brackets are the p-values. The χ^2 report the statistic of a χ^2 test for the joint null hypothesis $\alpha = 0$ and $\beta = 1$ (p-values in parentheses) in the following univariate regression: $\ln(\sigma_R) = \alpha + \beta \ln(\sigma_M)$, where σ_R = realized volatility and σ_M= average implied volatility. The χ^{2a} , χ^{2b} report the statistic of a χ^2 test for the joint null hypothesis $\gamma = 0$ and $\beta = 1$ or $\gamma = 1$ and $\beta = 0$ (p-values in parentheses) in the following regressions: $\ln(\sigma_R) = \alpha + \beta \ln(\sigma_M) + \gamma \ln(\sigma_j)$, where σ_R = realized volatility, σ_M= average implied volatility σ_j= volatility forecast j, j= ATMP, LR, GAR. The superscripts ***, **, * indicate that the slope coefficient is significantly different from one at the 1%, 5%, and 10% critical levels respectively. The superscripts $^{+++}$, $^{++}$, $^{+}$ indicate that the slope coefficient is significantly different from zero at the 1%, 5%, and 10% critical levels respectively. The last column shows the Hausman (1978) specification test statistic (one degree of freedom): 5% critical level = 3,841.

Moreover from the results in augmented regression (3) we see that average implied volatility does not subsume all the information of ATMP.

Therefore we conclude that the attempt of combining at-the-money call and put implied volatilities in a single estimate does not improve the forecasting power over the simple use of ATMP.

Finally, we test our results for robustness by adopting an instrumental variable procedure. The Hausman (1978) specification test reported in the last column of Table 6 indicates that the errors-in-variables problem is significant only in univariate regression (1). We report in Panel B the TSLS regression output, but the results do not alter the conclusions based on the OLS regression.

7 – Conclusion

In this paper we investigated how the information content of implied volatility varies according to moneyness and option type and we compared the latter option-based forecasts with historical volatility. The information content of implied volatility has been examined for the most liquid at-the-money and out-of-the-money call and put options i.e. the ones that are usually used as inputs for the computation of the smile function. Unlike previous studies, that use settlement prices, we used synchronous prices, matched in a one-minute interval.

The results show that the information content of implied volatility has a humped shape, with out-of-the-money options being less informative than at-the-money ones. This is consistent with the hedging pressure argument documented in Bollen and Whaley (2004), that causes out-of-the-money options to be less informative than at-the-money options. Overall, implied volatility forecasts contain more information about future realised volatility than LR volatility. The GAR forecast obtains roughly the same performance as ATMC implied volatility and is superior to both OTMC implied volatility and LR volatility.

Two hypotheses were tested: unbiasedness and efficiency of the different volatility forecasts. Overall, call implied volatilities forecasts are unbiased, while put implied volatilities are unbiased only after a constant adjustment given by the intercept of the regression. Efficiency was evaluated by assessing whether the implied volatility forecast subsumes all the information contained in historical volatility. Only ATMP implied volatility turns out to be efficient. Among the remaining three volatility forecasts, OTMP is marginally inefficient, while ATMC and OTMC are strongly inefficient.

By comparing pair-wise the four implied volatility forecasts, it is clear that ATMC subsumes all the information contained in OTMC, and ATMP subsumes all the information contained in both OTMP and OTMC. The comparison of ATMC and ATMP is less clear-cut, but we can conclude that ATMP yields a slightly better performance than ATMC. Therefore, in our sample, at-the-money put options are priced more efficiently than at-the-money call options. ATMP options, being more heavily traded than ATMC options, are more informative of future realised volatility. This is an interesting result, unlike previous research (see e.g. Christensen and Hansen (2002)), and is a warning against the a-priori choice of using call implied volatility. The attempt to combine ATMC and ATMP in a single forecast in order to cancel possible errors across option-type does not lead to an improvement over the simple use of ATMP implied volatility.

The results of this paper are significant for the understanding of the information content of option based measures of volatility and have potential implications for all the fields in which these measures can be used: portfolio selection models, derivative pricing models, hedging, risk measurement and risk management techniques in general. Moreover, these results also have a potential influence on the way market volatility indexes are computed. In particular, the VDAX-New, the new volatility index of the German equity market, is now based on the so-called "model-free" implied volatility that uses the information content of the whole smile function. The VDAX-New has replaced the old VDAX, that was computed by using only at-the-money options (pairs of calls and puts with the four strikes below and above the at-the-money point). The present investigation provides some indications to improve the information content of the VDAX-New: overall put options are more informative than call options, ATMP are preferred to ATMC, OTMP predict future realised volatility better than both ATMC and OTMC. How these rules can be embedded in the index and the empirical comparison between the suggested modifications and the existing VDAX-New is left for future research.

References

Andersen, T.G. and O. Bondarenko, 2007. Construction and Interpretation of Model-Free Implied Volatility. In Israel Nelken (Ed.): *Volatility as an Asset Class*, Risk Books, London, 141-181.

Becker, R., A. E. Clements and S. I. White, 2007. Does implied volatility provide any information beyond that captured in model-based volatility forecasts? Journal of Banking and Finance, 31 (8), 2535-2549.

Black, F. and M. Scholes, 1973. The pricing of options and corporate liabilities. Journal of Political Economy, 81, 637-654.

Bollen, N.P.B. and R.E. Whaley, 2004. Does net buying pressure affect the shape of implied volatility functions? The Journal of Finance, 59, 711-753.

Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroscedasticity. Journal of Econometrics, 31, 307-327.

Britten-Jones, M. and A. Neuberger, 2000. Option prices, implied price processes and stochastic volatility. Journal of Finance, 55, 839-866.

Buraschi, A. and J. Jackwerth, 2001. The price of a smile: hedging and spanning in option markets. The Review of Financial Studies, 14, 495-527.

Christensen, B.J. and N.R. Prabhala, 1998. The relation between implied and realized volatility. Journal of Financial Economics, 50, 125-150.

Christensen, B.J., C. S. Hansen and N.R. Prabhala, 2001. The telescoping overlap problem in options data. Working paper, University of Aarhus and University of Maryland.

Christensen, B.J. and C. S. Hansen, 2002. New evidence on the implied-realized volatility relation. The European Journal of Finance, 8 (2), 187-205.

Derman, E. and I. Kani, 1994. Riding on a smile. Risk, 7 (2), 32-39.

Diebold, F.X. and R.S. Mariano, 1995. Comparing predictive accuracy. Journal of Business & Economic Statistics, 13 (3), 134-144.

Donaldson, R.G. and M. Kamstra, 2005. Volatility forecasts, trading volume and the ARCH vs. option-implied tradeoff. Journal of Financial Research, 28 (4), 519-538.

Ederington, L. H. and W. Guan, 2005. The information frown in option prices. Journal of Banking and Finance, 29 (6), 1429-1457.

Fair, R.C. and R.J. Schiller, 1990. Comparing information in forecasts from econometric models. American Economic Review, 80 (3), 375-380.

Fleming, J., 1998. The quality of market volatility forecasts implied by S&P100 index option prices. Journal of Empirical Finance, 5 (4), 317-345.

Fleming, J., 1999. The economic significance of the forecast bias of S&P100 Index option implied volatility. Advances in Futures and Option Research, 10, 219-251.

Godbey, J.M. and J.W. Mahar, 2007. Forecasting power of implied volatility: evidence from individual equities. B>Quest, University of West Georgia.

Hausman, J., 1978. Specification tests in econometrics. Econometrica, 46, 1251-1271.

Hentschel, L., 2003. Errors in implied volatility estimation. Journal of Financial and Quantitative analysis, 38 (4), 779-810.

Jiang, G. J. and Y. S. Tian, 2005. Model-free implied volatility and its information content. The Review of Financial Studies, 18 (4), 1305-1342.

Jorion, P., 1995. Predicting volatility in the foreign exchange market. Journal of Finance, 50 (2), 507-528.

Moriggia, V., S. Muzzioli and C. Torricelli, 2007. Call and put implied volatilities and the derivation of option implied trees. Frontiers in Finance and Economics, 4 (1), 35-64.

Poon, S. and C.W. Granger, 2003. Forecasting volatility in financial markets: a review. Journal of Economic Literature, 41, 478-539.

Poon, S., 2005. A practical guide to forecasting financial market volatility. (Whiley Finance, Chichester).

Rubinstein, M., 1994. Implied binomial trees. The Journal of Finance, 49, 771-818.

Sarwar, G., 2005. The informational role of option trading volume in equity index options markets. Review of Quantitative Finance and Accounting, 24, 159-176.

Taylor, S. J., Y. Zhang and P. K. Yadav, 2006. The Information Content of Implied Volatilities and Model-Free Volatility Expectations: Evidence from Options Written on Individual Stocks. Available at SSRN: http://ssrn.com/abstract=890522.