The performance of implied volatility in forecasting future volatility: an analysis of three major equity indices from 2004 to 2010

by

Vladimir M. Ionesco

Submitted to the MIT Sloan School of Management on May 6, 2011 in partial fulfillment of the requirements for the degree of Master of Science in Management Studies.

Abstract

In this thesis, we investigate whether implied volatility is an efficient estimator of future one-month volatility from an informational perspective and whether it outperforms historical volatility in this regard.

We first compare the predictive powers of implied volatility, simple historical volatility, and exponential historical volatility, using monthly observations of the S&P 500, FTSE 100, and DAX equity and option markets from 2004 to 2010.

Then, we introduce a GARCH(1,1) model and compare in-sample GARCH-fitted volatility and implied volatility from 2004 to 2010, as well as out-of-sample GARCH-forecasted volatility and implied volatility from 2005 to 2010, using data on the S&P 500.

We find that implied volatility is not only an efficient estimator of future volatility, but also that its information content is at least as good, if not much better, than that of historical volatility. Our results also suggest that implied volatility systematically subsumes the information included in historical volatility, even when a GJR-GARCH model is utilized.

Thesis Advisor: Eric Jacquier
Title: Visiting Associate Professor of Finance

Contents

List of Tables

List of Figures

1 Introduction

Volatility is a most critical concept in both the theory and practice of finance. It is a measure of the uncertainty about a financial instrument's probability distribution of returns and is commonly defined as the standard deviation of the returns of the said instrument within a set time span. Volatility is thus intimately linked to the fundamental concept of risk, "*the central element that influences financial behavior*" as Nobel Laureate Robert C. Merton once put it. Consequently, many investors seek to forecast volatility, for risk management, portfolio selection, or valuation purposes, or for designing trading strategies, such as volatility arbitrage. To do so, they mainly rely on two estimators: historical volatility, also known as past realized volatility, and implied volatility, a concept formally born in 1973 with the publication of the Black-Scholes option pricing model and the creation of the Chicago Board Options Exchange (CBOE), the first market of its kind in the United States. While historical volatility is directly computable from past market data, implied volatility is extracted from option prices, using models such as Black-Scholes.

The rationale behind the use of historical volatility to forecast future volatility lies in the idea that the past tends to repeat itself, which leads to common financial principles such as mean-reversion. By definition, simple historical volatility is an unconditional predictor that ignores the most recent publicly available data. Moreover, when computed using standard statistical techniques, historical volatility fails to reflect the possible predictability of "true" volatility. In contrast, implied volatility is widely seen as the market's estimate of future volatility, and if markets are efficient, it should thus reflect all the information available at a given time, including that contained in historical volatility. Bodie and Merton (1995) use the period that precedes the Persian Gulf War of 1991 to illustrate the superiority of implied volatility over historical volatility when estimating future volatility.

Nevertheless, despite the widely-shared belief among finance practitioners that implied volatility is, indeed, a much better estimator than historical volatility, research has produced intriguingly divergent results about whether implied volatility actually estimates future volatility or whether it does so efficiently.

Early academic articles support the superiority of the forecasting characteristics of implied volatility over those of historical volatility. In the first study on this kind, Latane and Rendleman (1976) test the relationship between an average of implied volatilities and subsequent realized volatility, using the closing call option and stock prices of twenty four companies traded on the Chicago Board Options Exchange. They conclude that weighted implied volatility is generally a better predictor of future volatility than historical volatility. Chiras and Manaster (1978) reach the same conclusion, after analyzing twenty three monthly observations between 1973 and 1975. Although the authors discover no significant difference between the forecasting characteristics of historical volatility and implied volatility in the first nine months covered by their study, they find that implied volatility becomes a much better predictor of future volatility in the following fourteen months. The findings of Beckers (1981) also suggest that implied volatility incorporates and outperforms the predictive information of historical volatility. However, in the absence of large time-series data, the above-mentioned research resorts to a static cross-sectional regression approach.

Using time-series data in a dynamic context, several later works reach opposite conclusions after studying the actively-traded OEX options on the S&P 100 index. Day and Lewis (1992) analyze these options between 1983 and 1989 and conclude that, although implied volatility may contain some information about subsequent volatility, it is still outperformed by time-series models of conditional volatility, such as GARCH and EGARCH. In their study of individual stock options from 1982 to 1984, Lamoureux and Lastrapes (1993) also find that the information contained in historical volatility is superior to that contained in implied volatility. Canina and Figlewski (1993) reach a more radical conclusion, arguing that "*implied volatility has virtually no correlation with future volatility*" and that "*it does not incorporate the information contained in recent observed volatility*".

However, recent research, notably since Christensen and Prabhala (1998), tends to support the idea that implied volatility does not only contain substantial information about future volatility, but that it is also more predictive than historical volatility. Christensen and Prabhala (1998) study OEX options with longer time series and non-overlapping data covering the period between November 1983 and May 1995. They show that implied volatility is an unbiased and efficient estimator of subsequent volatility and that, in some cases, implied volatility includes the information contained in past realized volatility. Hansen (1999) analyzes the Danish option and equity markets and also concludes that implied volatility is a good forecaster of subsequent realized volatility, that its bias is negligible, and that it incorporates the information of historical volatility. Christensen and Hansen (2002) confirm the results of Christensen and Prabhala (1998) in their study of a more recent period with checks of robustness. In addition, they extend their analysis to put options and find that put implied volatility is also predictive, though not as much as call implied volatility. A subsequent study of the S&P 500 index and its options by Shu and Zhang (2003) also support the idea that implied volatility is a superior predictor of future volatility. In a different context, Szakmary et al. (2003) study thirty five options markets and find that, for a large majority of commodities, implied volatility outperforms historical volatility in forecasting the volatility of the underlying prices. More recently, after analyzing the S&P/ASX 200 index options traded on the Australian Stock Exchange, Li and Yang (2008) also conclude that the implied volatilities of both calls and puts are better than historical volatility at predicting subsequent volatility. Moreover, they find that the volatility implied in call options is a nearly unbiased estimator of future volatility.

Contrasting with past findings, recent research thus gives credit to the widely-shared belief that implied volatility does contain some information about future volatility, and that it is superior in this respect to historical volatility.

By analyzing the time-series data of three major equity indices - S&P500, FTSE 100, and DAX - and their options from 2004 to 2010, we aim to investigate whether the forecasting power of implied volatility for future volatility, as demonstrated by Christensen and Prabhla (1998) and confirmed by subsequent research, is still verified in the most recent years and in both American and foreign major option and stock markets. In this respect, the aftermath of the financial crisis of 2008 deserves a close examination.

2 Methodology

2.1 Data specifications

Our empirical analysis focuses on the S&P 500 (SPX), FTSE 100 (ESX), and DAX equity indices and their options. The data covers the period from January 2004 to December 2010 for the S&P 500 and the DAX, and the period from November 2004 to December 2010 for the FTSE 100. The option data used in the present study has been provided by IVolatility.com.

S&P 500 options are European-style options traded on the Chicago Board Option Exchange (CBOE). The options are quoted in index points, with a multiplier of US$100. The average daily volume in January-November 2010 was 711,231. The last trading day is the last business day before the third Friday of the expiration month. The expiration date is the Saturday following the third Friday of the expiration month. The nearest twelve calendar months are available for trading. SPX options are cash-settled.

FTSE 100 options are European-style options traded on the NYSE Liffe London exchange. The options are quoted in index points, with a multiplier of 10. The last trading day is the third Friday of the expiration month. Trading stops after 10:15 am, London time. In the event of the third Friday not being a business day, the last trading day is the last business day prior such Friday. The settlement day is the business day following the last trading day. The expiration months are the nearest eight of March, June, September, and December, as well as the nearest four calendar months. FTSE100 options are cash-settled.

DAX options are European-style options traded on the Frankfurt Stock Exchange. The options are quoted in index points, with a multiplier of EUR25. The last trading day is the third Friday of the expiration month. In the event of the third Friday not being a business day, the last trading day is the last business day prior such Friday. The expiration day is the business day following the last trading day. The next, the second, and the third quarter-end months (March, June, September, December) are available for trading. DAX options are cash-settled.

2.2 Sampling procedure

The sampling procedure partly follows that used by Christensen and Prabhala (1998). Our data consists of monthly observations with no overlaps, so as to avoid drawbacks such as high residual auto-correlation, as well as imprecise or biased regression estimates. In order to limit any excess volatility resulting from the opening of new contracts and obtain consistent and accurate data, we record the closing characteristics of both call and put options on the S&P 500, FTSE 100, and DAX indices between 2004 and 2010 that are:

1. the nearest-to-the-money, the least deep out-of-the-money, and the second least deep out-of-the-money;
2. traded on the Wednesday following the last trading day of a given month, or on the following business day when such Wednesday happens to not be a trading day;
3. traded with significant volume;
4. expiring the coming month.

This procedure gives 82 monthly observations for the S&P 500, 73 monthly observations for the FTSE 100, and 82 monthly observations for the DAX. The corresponding historical volatility and ex-post future volatility over the remaining life of the options are computed according to the procedures detailed below. The data is subsequently divided into two distinct samples - 2004-2007 and 2008-2010 - in order to examine the relationship between implied volatility and future volatility with and without the potentially-distorting impact of the financial crisis of 2008. We thus end with 47 observations in 2004-2007 and 35 in 2008-2010 for both the S&P 500 and the DAX indices, and with 38 in 2004-2007 and 35 in 2008-2010 for the FTSE 100 index.

2.3 Variable definitions

2.3.1 Implied volatility

For any call option price C_t observed at time t, the implied volatility $\sigma_{ic,t}$ is computed by numerically solving the Black-Scholes call option pricing formula, assuming no dividends:

$$C_t = S_t N(d_{1,t}) - K_t e^{-r_{f,t}\tau_t} N(d_{2,t}) \tag{1}$$

where

$$d_{1,t} = \frac{ln(S/K) + (r + \frac{1}{2}\sigma_{ic,t}^2)\tau_t}{\sigma_{ic,t}\sqrt{\tau_t}} \tag{2}$$

$$d_{2,t} = d_{1,t} - \sigma_{ic,t}\sqrt{\tau_t} \tag{3}$$

For any close put option price P_t observed at time t, the implied volatility

$\sigma_{ip,t}$ is computed by numerically solving the Black-Scholes put option pricing formula, assuming no dividends:

$$P_t = K_t e^{-r_{f,t}\tau_t} N(-d_{2,t}) - S_t N(-d_{1,t}) \tag{4}$$

where

$$d_{1,t} = \frac{ln(S/K) + (r_{f,t} + \frac{1}{2}\sigma_{ip,t}^2)\tau_t}{\sigma_{ip,t}\sqrt{\tau_t}} \tag{5}$$

$$d_{2,t} = d_{1,t} - \sigma_{ip,t}\sqrt{\tau_t} \tag{6}$$

In the above equations, K_t denotes the strike price of the option, S_t the spot price of the underlying index, τ_t the time to expiration, $r_{f,t}$ the risk-free interest rate, and $N(.)$ stands for the standard normal cumulative distribution function. The risk-free rate is the London inter-bank borrowing rate in the Eurodollar market (LIBOR). The time to expiration is generally around 24 calendar days. The data has been provided by IVolatility.

On each observation date t, the final implied volatility $\sigma_{i,t}$ is computed as a weighted average of the three call and three put option implied volatilities recorded:

$$\sigma_{i,t} = 0.2(\sigma_{ic1,t} + \sigma_{ip1,t}) + 0.2(\sigma_{ic2,t} + \sigma_{ip2,t}) + 0.1(\sigma_{ic3,t} + \sigma_{ip3,t}) \tag{7}$$

where $\sigma_{ic1,t}$, $\sigma_{ip1,t}$, $\sigma_{ic2,t}$, $\sigma_{ip2,t}$, $\sigma_{ic3,t}$, and $\sigma_{ip3,t}$ denote, respectively, the implied volatility at time t of the nearest-to-the-money call, the nearest-to-the-money put, the least deep out-of-the-money call, the least deep out-of-the-money put, the second least deep out-of-the-money call, and the second least deep out-of-the-money put. Such a weighted average is both simple to implement and more likely than a single implied volatility to attenuate the noise that might impact an observation.

2.3.2 Historical volatility

Merton (1980) has shown that the accuracy of an estimate of volatility using past volatility increases with the sampling frequency within a given overall observation period. We thus choose to use daily data. Furthermore, we decide to compute both a simple average and an exponentially-weighted (EW) average of historical volatility. Assuming that volatility varies with time, the EW version compensates to some extent for one of the shortcomings of simple historical volatility, namely the inevitable lag in information incorporation, by giving greater weight to the most recent daily observations, and thus to more recent information known to the market. Usually, EW historical volatility is used to forecast one or a few days ahead, using high frequency data. However, we still would like to assess how it behaves as an estimator of one-month volatility.

Simple average

For each option price observation at time t, we measure the historical volatility over the past τ_t trading days, including the day of the observation, where τ_t is the number of days until the expiration of the option, by the sample standard deviation of the daily index returns during that period. We decide to omit the usual estimator of the mean to avoid excessive noise, and use the following formula:

$$\sigma_{h,t} = \sqrt{\frac{252}{\tau_t} \sum_{i=t-\tau_t}^{t} (R_i)^2} \tag{8}$$

where R_i denotes the log-return on day i. Let S_i be the index level on the same day i, we have:

$$R_i = ln(S_i/S_{i-1}) \tag{9}$$

where ln denotes the natural logarithm.

Exponentially-weighted average

For each option price observation at time t, we also measure the historical volatility using an exponentially-weigthed average of past daily volatility, including the day of the observation, with a decay factor of 0.94. This method of computation is inspired by that developed by Riskmetrics. We use the formula:

$$\sigma_{he,t} = \sqrt{252 * 0.06 * \sum_{i=0}^{t} 0.94^n (R_i)^2} \tag{10}$$

where 0.94 is the decay factor, 0.06 the sum of the weights, and where R_i denotes the log-return on day i according to (9).

2.3.3 Ex-post future volatility

For each time-t option price observation, we measure the ex-post future volatility by the sample standard deviation of the daily index returns over the remaining life τ_t of the option. Again, we deliberately omit the estimator of the mean, which would have been too noise-sensitive. We use the following formula:

$$\sigma_{f,t} = \sqrt{\frac{252}{\tau_t} \sum_{i=t+1}^{t+\tau_t} (R_i)^2} \tag{11}$$

where R_i denotes the log-return on day i according to (9).

Finally, to improve the normality of our variables, we apply a logarithmic transformation to our observations of implied volatility, historical volatility, and ex-post future volatility. Namely, for a given observation of volatility with value σ, we use the value $ln(\sigma)$ instead, where ln denotes the natural logarithm. The ordinary least squares (OLS) regressions are run with these logvalues.

2.4 Regressions

To examine the information content of implied volatility and historical volatility before and after the financial financial crisis of 2008, we run the following regressions for each index, first in 2004-2007 and then in 2008-2010:

$$\sigma_{f,t} = \alpha + \beta_h \sigma_{h,t} + \epsilon_t \tag{12}$$

$$\sigma_{f,t} = \alpha + \beta_{he} \sigma_{he,t} + \epsilon_t \tag{13}$$

$$\sigma_{f,t} = \alpha + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{14}$$

$$\sigma_{f,t} = \alpha + \beta_h \sigma_{h,t} + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{15}$$

$$\sigma_{f,t} = \alpha + \beta_{he} \sigma_{he,t} + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{16}$$

These regressions allow us to test three relevant hypotheses (Christensen and Prabhala (1998), Christensen and Hansen (2002)). First, if implied volatility contains some information about future realized volatility, then the coefficient of implied volatility β_{iv} should be nonzero in equations (14), (15), and (16). Second, if implied volatility is an unbiased estimator, β_{iv} should be equal to 1 and the intercept α to 0 in equations (14),(15), and (16).

Finally, if implied volatility is an informationally-efficient predictor of future realized, it should subsume any information contained in historical volatility, and thus the coefficients β_h in equation (15) and β_{he} (16) should both be equal to 0 and the adjusted R-squared of regressions (15) and (16) should not be higher than that of regression (14). Based on a number of results and previous research, we expect historical volatility to contain some information about future volatility, and thus β_h and β_{he} to be nonzero and statistically significant in equations (12) and (13). We will also check whether exponentially-weighted estimates of standard deviation are better than simple estimates.

3 Descriptive statistics

Tables 1, 2, and 3 display relevant statistics describing the data on ex-post future volatility, simple historical volatility, exponentially-weighted volatility, and implied volatility.

First of all, we observe that the magnitude of the means and standard deviations of ex-post future, simple, exponentially-weighted, and implied volatilities across the three markets are somewhat equivalent, and since volatility is not rising or decreasing uniformly throughout the sample, we would not expect otherwise. For instance, in 2004-2007, the means of ex-post future volatility are 0.127 for the S&P 500, 0.126 for the FTSE 100, and slightly higher for the DAX, with a value of 0.147, while the standard deviations are 0.044, 0.065, and 0.053 respectively. Besides, as demonstrated by figure 2, we can see that simple and EW historical volatility differ only slightly.

Second, we notice that between 2004-2007 and 2008-2010, both the means and standard deviations of the four measured volatilities rise sharply. For example, the mean of the simple historical volatility of the S&P 500 increases from 0.112 to 0.259, that of the FTSE 100 from 0.116 to 0.241, and that of the DAX from 0.142 to 0.268, while the standard deviations rise by an even greater factor: from 0.037 to 0.259, from 0.053 to 0.128, and from 0.042 to 0.145 respectively.

Third, in the cases of the S&P 500 and the DAX, the cross-correlations between historical and implied volatility increase substantially in 2008-2010, virtually reaching a perfect correlation of 1. For instance, the correlation between simple historical and implied volatility rises from 0.78 to 0.95 in the case of the SP, and from 0.70 to 0.92 in the case of the DAX, while it only rises from 0.86 to 0.88 in the case of the FTSE.

Fourth, while implied volatility is slightly greater than the three other volatilities across the three markets in 2004-2007, it stops being so in 2008-2010. However, in both periods, its standard deviations are always lower than those of ex-post future volatility, simple historical volatility, and exponential historical volatility.

Finally, the Jarque-Bera statistics demonstrate the effectiveness of applying a logarithmic scale to normalize a set of data. As an example, the Jarque-Bera statistic of the DAX ex-post future volatility in 2004-2007 falls from a huge 140.33 to an acceptable level of 5.38.

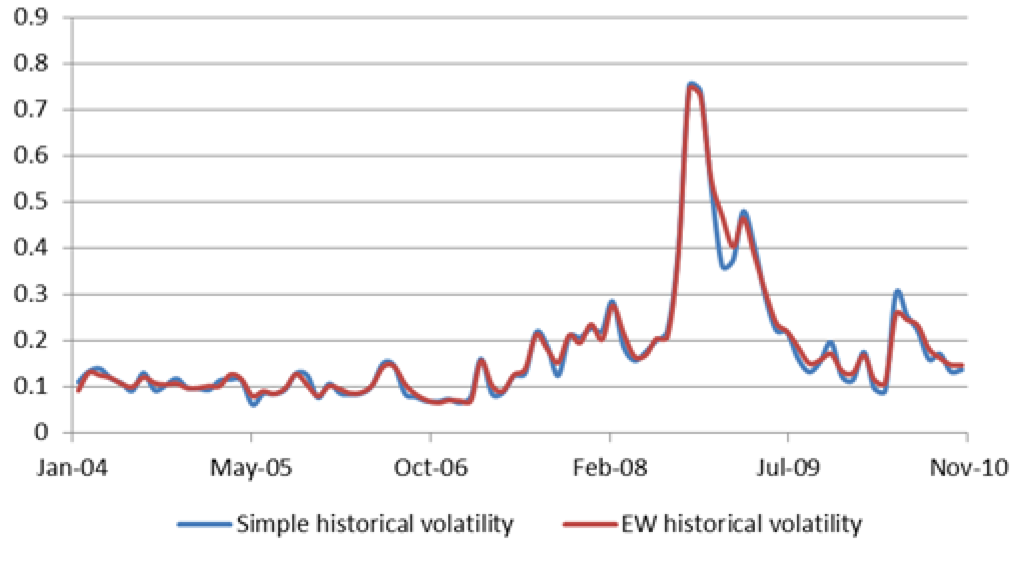

Figure 1: S&P 500: simple and exponential historical volatility from 2004 to 2010

Table 1: Data based on 47 observations of the S&P 500 index and its options from January 2004 to December 2007 and 35 observations from January 2008 to December 2010. Values have been annualized.

2004-2007

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.127	0.112	0.113	0.125
StDev	0.044	0.037	0.035	0.035
Skewness	1.50	1.29	1.33	1.62
Kurtosis	4.79	4.52	4.69	6.44
Jarque-Bera	23.85	17.58	19.50	43.64

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-2.207	-2.239	-2.225	-2.111
StDev	0.332	0.305	0.284	0.248
Skewness	0.85	0.55	0.59	0.86
Kurtosis	3.03	2.98	3.19	3.60
Jarque-Bera	5.61	2.27	2.78	6.52

2008-2010

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.256	0.259	0.265	0.256
StDev	0.164	0.162	0.160	0.112
Skewness	1.94	1.79	1.76	1.63
Kurtosis	6.93	6.11	5.76	5.76
Jarque-Bera	44.46	32.78	29.23	26.63

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-1.514	-1.496	-1.463	-1.437
StDev	0.536	0.527	0.499	0.377
Skewness	0.59	0.61	0.83	0.80
Kurtosis	3.23	2.97	3.02	3.07
Jarque-Bera	2.09	2.18	3.98	3.78

Correlation matrix (log of σ)

	2004-2007			2008-2010		
	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$
$\sigma_{h,t}$	1	0.95	0.78	1	0.98	0.95
$\sigma_{he,t}$	0.95	1	0.82	0.98	1	0.97
$\sigma_{iv,t}$	0.78	0.82	1	0.95	0.97	1

Table 2: Data based on 38 observations of the FTSE 100 index and its options from November 2004 to December 2007 and 35 observations from January 2008 to December 2010. Values have been annualized.

2004-2007

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.126	0.116	0.120	0.125
StDev	0.065	0.053	0.054	0.047
Skewness	1.76	1.68	1.60	1.85
Kurtosis	5.72	5.27	5.04	6.63
Jarque-Bera	31.41	26.07	22.80	42.46

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-2.173	-2.233	-2.201	-2.130
StDev	0.423	0.382	0.380	0.314
Skewness	1.06	1.06	1.04	1.14
Kurtosis	3.18	3.28	3.00	3.74
Jarque-Bera	7.18	7.22	6.83	9.07

2008-2010

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.240	0.241	0.249	0.249
StDev	0.132	0.128	0.130	0.101
Skewness	2.25	2.06	2.02	1.68
Kurtosis	9.34	8.10	7.68	5.60
Jarque-Bera	88.23	62.81	55.63	26.43

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-1.532	-1.527	-1.492	-1.454
StDev	0.441	0.440	0.433	0.344
Skewness	0.87	0.77	0.81	0.99
Kurtosis	3.73	3.53	3.47	3.34
Jarque-Bera	5.23	3.84	4.20	5.93

Correlation matrix (log of σ)

	2004-2007			2008-2010		
	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$
$\sigma_{h,t}$	1	0.97	0.86	1	0.98	0.88
$\sigma_{he,t}$	0.97	1	0.92	0.98	1	0.93
$\sigma_{iv,t}$	0.86	0.92	1	0.88	0.93	1

Table 3: Data based on 47 observations of the DAX index and its options from January 2004 to December 2007 and 35 observations from 2008 to 2010. Values have been annualized.

2004-2007

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.147	0.142	0.143	0.157
StDev	0.053	0.042	0.036	0.036
Skewness	2.10	0.81	0.67	0.83
Kurtosis	10.34	3.07	2.65	3.30
Jarque-Bera	140.33	5.18	3.80	5.62

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-1.967	-1.995	-1.978	-1.872
StDev	0.315	0.288	0.249	0.218
Skewness	0.68	0.23	0.22	0.37
Kurtosis	3.98	2.47	2.33	2.55
Jarque-Bera	5.38	0.98	1.28	1.45

2008-2010

Statistics	Realized volatility	Historical volatility (SA)	Historical volatility (EWA)	Implied volatility
Mean	0.262	0.268	0.273	0.264
StDev	0.140	0.145	0.142	0.098
Skewness	2.07	1.64	1.76	1.44
Kurtosis	7.77	5.59	6.14	4.41
Jarque-Bera	58.15	25.54	32.38	14.97

Statistics	Log realized volatility	Log historical volatility (SA)	Log historical volatility (EWA)	Log implied volatility
Mean	-1.442	-1.429	-1.402	-1.388
StDev	0.433	0.465	0.439	0.325
Skewness	0.92	0.66	0.76	0.88
Kurtosis	3.60	2.80	3.05	2.98
Jarque-Bera	5.49	2.59	3.38	4.53

Correlation matrix (log of σ)

	2004-2007			2008-2010		
	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$
$\sigma_{h,t}$	1	0.95	0.70	1	0.99	0.92
$\sigma_{he,t}$	0.95	1	0.80	0.99	1	0.94
$\sigma_{iv,t}$	0.70	0.80	1	0.92	0.94	1

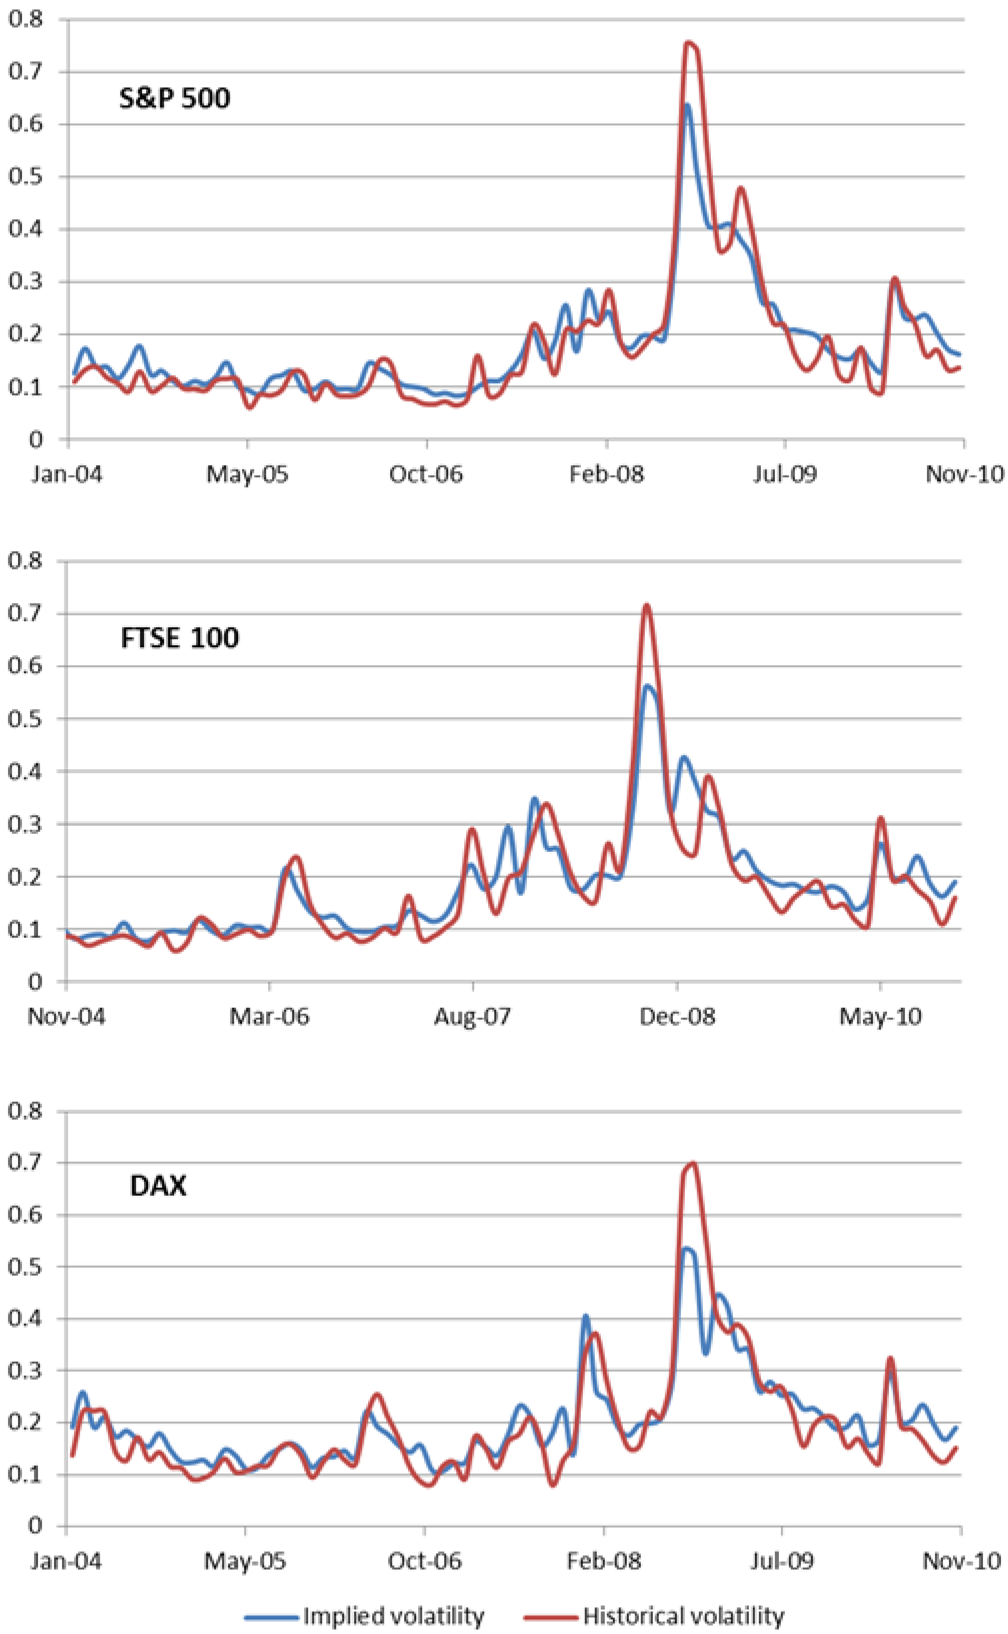

Figure 2: S&P, FTSE, and DAX: simple historical volatility and implied volatility from 2004 to 2010

4 Empirical Results

4.1 S&P 500

The ordinary least squares (OLS) coefficient estimates for the S&P 500 are displayed in Table 4.

Let us first focus on 2004-2007. In our first regression, we observe that the coefficient estimate for simple historical volatility has a value of 0.567, which, with a t-stat of 4.10, is statistically significant. The adjusted R-squared indicates that, when considered alone, simple volatility accounts for 25.57% of the variance of future volatility. Our second regression shows that exponentially-weighted historical volatility performs negligibly better. If the equally-significant coefficient estimate has a higher value of 0.633, we obtain an R-squared of 0.2765, which is virtually equal to that of our previous regression. However, we observe in our third regression that implied volatility outperforms both simple and exponentially-weighted volatility. Not only is the coefficient estimate greater, with a value of 0.943, and more significant statistically with at t-stat of 6.67, but, above all, the R-squared indicates that implied volatility explains nearly 50% of the variance of future volatility. What is more, in contrast with our two first regressions, we get an intercept that is not statistically significant. These results suggest that while the information content of simple and EW historical volatility are relatively equivalent, that of implied volatility is substantially greater.

Equations (15) and (16) enable us to assess whether the information content of implied volatility subsumes those of our two other estimators. The first striking element is that the R-squared of both regressions do not exceed that of regression (14). Secondly, the coefficient estimates of both simple and EW historical volatility are both not statistically significant at the commonly-accepted level. On the other hand, the coefficient estimates of implied volatility are very close to one and statistically significant, with t-stats greater than 4.

Our analysis of the period 2004-2007 suggests that implied volatility does contain some significant amount information about future volatility, outperforms both simple and EW historical volatility in that regard, and subsumes the information contained in the other two estimators. Besides, with coefficient estimates very close to 1, it appears to be a nearly unbiased estimator of future volatility.

If we move on to the period 2008-2010, we obtain very different results. First of all, the R-Squared of all six regressions, with values within the 0.51-0.56 range, are virtually equal to each other and are also higher than any single R-squared obtained in our analysis of 2004-2007. Of the three estimators, it seems that none outperforms the others. In regressions (12), (13), and (14), the coefficient estimates of simple volatility, EW volatility, and implied volatility are all statistically significant, with values of 0.739, 0.796, and 1.070 respectively, consistent with the hierarchy found in 2004-2007. However, in our last two regressions, none of the coefficient estimates manages to have a t-stat greater than 1.96. A plausible explanation is that the estimators "neutralized" each other because of their higher correlation, as seen in Table 1.

It appears that in 2008-2010, the three estimators show equivalent performance in predicting future volatility, that this performance is greater than in 2004-2007, and that there is no added value in combining these estimators in terms of forecasting power.

Table 4: S&P500: conventional OLS coefficient estimates

2004-2007

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.938	0.567	-	-	0.2557	2.36	2.89	0.05
(-3.00)	(4.10)						
-0.799	-	0.633	-	0.2765	2.64	0.74	0.06
(-2.43)		(4.31)					
-0.216	-	-	0.943	0.4859	14.14	0.58	0.07
(-0.72)			(6.67)				
-0.233	-0.080	-	1.020	0.4764	15.19	1.24	0.10
(-0.76)	(-0.43)		(4.47)				
-0.251	-	-0.145	1.079	0.4794	13.94	2.44	0.11
(-0.81)		(-0.66)	(4.30)				

2008-2010

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.408	0.739	-	-	0.5149	5.94	1.08	0.17
(-2.12)	(6.09)						
-0.350	-	0.796	-	0.5356	7.78	0.94	0.23
(-1.81)		(6.34)					
0.024	-	-	1.070	0.5547	21.20	3.01	0.26
(0.10)			(6.58)				
-0.026	0.123	-	0.908	0.5424	19.35	5.49	0.24
(-0.09)	(0.33)		(1.73)				
-0.054	-	0.210	0.802	0.5435	18.55	5.27	0.25
(-0.18)		(0.43)	(1.25)				

4.2 FTSE 100

Table 5 displays the OLS coefficient estimates for the FTSE 100.

We start our analysis by studying 2004-2007. The coefficient estimates of simple and exponentially-weighted historical volatility are both statistically significant, with values of 0.688 (t-stat of 4.75) and 0.691 (t-stat of 4.73) respectively. The R-squared of the first regression is 0.3684 and that of the second is 0.3661. It thus appears that there is no palpable difference in the predictive power of simple and EW volatility. On the other hand, in our third regression, we obtain a statistically-significant (t-stat of 5.96) coefficient estimate of 0.950 for implied volatility, an insignificant intercept, and an R-squared of 0.4828. Our results tend again to support the idea that implied volatility is a better estimator than historical volatility.

Let us now consider regressions (15) and (16). First and foremost, we observe that their R-squareds (0.4689 and 0.4751 respectively) are equivalent to that of regression (14), suggesting that combining historical and implied volatility does not produce better results in estimating future volatility than implied volatility alone. Secondly, the only coefficient estimates that have statistical significance are those of implied volatility, with a value of 0.883 and a t-stat of 2.79 in regression (15), and a value of 1.215 and a t-stat of 2.91 in regression (16).

Our analysis of 2004-2007 shows that implied volatility is a better estimator than historical volatility, and that it incorporates the information content of the other two estimators. Furthermore, combining implied volatility with historical volatility does not outperform the predictive power of implied volatility taken alone.

Moving on to 2008-2010, we observe that all five regressions account for the variance of future volatility to an equivalent extent, with R-squared of around 0.465. Only simple historical volatility slightly lags behind, with an R-squared of 0.4150 for regression (12). When considering each estimator separately, we get statistically-significant coefficient estimates of 0.658 for simple volatility, 0.707 for EW volatility, and 0.892 for implied volatility, with t-stats at the 5% significance level. Out of these three regressions, only regression (14) gives an insignificant intercept. In regressions (15) and (16), no coefficient estimates, except (barely) for that of simple historical volatility (t-stat of 1.97) in regression (15), is statistically significant.

Our results for 2008-2010 suggest that the performances of the three estimators are relatively equivalent, and that combining historical and implied volatility does not improve the overall predictive power.

4.3 DAX

The OLS coefficient estimates for the DAX are displayed in Table 6.

Again, let us first consider 2004-2007. Surprisingly, and contrary to the results of our studies of the S&P 500 and the FTSE 100, we find not only that the R-squared of all five regressions are equivalent, but that they are also of relatively low level (around 10%). For instance, while implied volatility alone accounts for as much as 48.59% of the variance of future volatility in the case of the S&P 500 and for 48.28% in that of the FTSE 100, here, it only accounts for 10.15%. In our first three regressions, the coefficient estimates of simple, exponentially-weighted, and implied volatility, with values of 0.391, 0.444, and 0.502 respectively, are also much lower. Moreover, regressions (15)

Table 5: FTSE 100: conventional OLS coefficient estimates

2004-2007

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.637	0.688	-	-	0.3684	7.48	6.19	0.07
(-1.94)	(4.75)						
-0.653	-	0.691	-	0.3661	10.44	3.27	0.08
(-2.00)		(4.73)					
-0.150	-	-	0.950	0.4828	17.88	1.50	-0.06
(-0.44)			(5.96)				
-0.150	0.064	-	0.883	0.4689	17.09	2.51	-0.06
(-0.43)	(0.25)		(2.79)				
-0.109	-	-0.237	1.215	0.4751	17.88	1.54	-0.08
(-0.31)		(-0.69)	(2.91)				

2008-2010

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.527	0.658	-	-	0.4150	4.49	0.17	0.13
(-2.52)	(5.01)						
-0.476	-	0.707	-	0.4670	7.24	0.01	0.18
(-2.40)		(5.55)					
-0.235	-	-	0.892	0.4690	13.73	0.33	0.30
(-0.98)			(5.57)				
-0.256	0.201	-	0.666	0.4621	11.23	3.97	0.22
(-1.06)	(1.97)		(0.76)				
-0.309	-	0.357	0.474	0.4702	11.10	2.91	0.23
(-1.24)		(1.09)	(1.04)				

and (16) give a statistically-significant coefficient estimate for none of the three estimators. We also note that the intercept is significant in all five regressions.

Intriguingly enough, our analysis of 2004-2007 indicates that no estimator outperforms the other two, that a combination of historical and implied volatility does not generate greater performance, and that the overall predictive power is comparatively poor.

Now, if we move to 2008-2010, our results are consistent with what we find for the other markets. First of all, the adjusted R-squareds of all five regressions are equivalent, with a level slightly higher than 0.50. In regressions (12), (13), and (14), the coefficient estimates of simple, exponential, and implied volatility are statistically significant, with a t-stat of approximately 6, and their magnitude, respectively 0.674, 0.729, and 0.966, are in line with our findings in the other two markets for 2008-2010. Again, our two last regressions do not result in any significant coefficient estimate for either historical or implied volatility.

In 2008-2010, we find that the three estimators show equivalently-high predictive power, and that combining historical and implied volatility does not entail better performance in that regard.

4.4 Residual diagnostics

Tables 4, 5 and 6 display the Jarque-Bera statistic (JB), the White statistic (White), and the first-order residual autocorrelation (Autocorr.) for each regression.

The Jarque-Bera test is a test of normality, based on kurtosis and skewness, that we conduct on the residuals of each regression. The test statistic has an asymptotic chi-square distribution with two degrees of freedom. This means that with a level of significance of 5%, we can reject the null hypothesis that the residuals are normally distributed if the JB statistic is above 5.99. As we can see, in most regressions, the hypothesis is rejected by a substantial margin.

The White test is a test of homoskedasticity, i.e. whether the variance of the residuals is constant. The statistic has a chi-square distribution with $k-1$ degrees of freedom, where k is equal to the number of regressors excluding the intercept. In the case of the number of regressors being equal to one, the statistic is computed so as to still have a chi-square distribution with one degree of freedom. At a level of significance set a 5%, we can the reject the null hypothesis of homoskedasticity when the White statistic is above 3.84. We can see that the hypothesis of homoskedasticity of the residuals is rejected in only a few instances: in regressions (15) and (16) in 2008-2010 in the case of the S&P 500, and in regressions (12) in 2004-2007 and (15) in 2008-2010 in the case of the FTSE 100.

We also computed the first order autocorrelation of the residuals of each re-

Table 6: DAX: conventional OLS coefficient estimates

2004-2007

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-1.187	0.391	-	-	0.1088	7.20	1.41	0.08
(-3.87)	(2.57)						
-1.090	-	0.444	-	0.1038	9.09	1.38	0.09
(-3.10)		(2.51)					
-1.028	-	-	0.502	0.1015	18.19	2.40	0.10
(-2.71)			(2.49)				
-0.963	0.245	-	0.275	0.1079	13.06	3.53	0.04
(-2.51)	(0.97)		(1.15)				
-0.960	-	0.255	0.268	0.0964	14.10	3.62	0.06
(-2.47)		(0.86)	(0.79)				

2008-2010

Intercept	$\sigma_{h,t}$	$\sigma_{he,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.479	0.674	-	-	0.5092	19.93	0.10	0.20
(-2.85)	(6.02)						
-0.420	-	0.729	-	0.5330	23.77	0.39	0.27
(-2.48)		(6.31)					
-0.101	-	-	0.966	0.5133	26.66	0.68	0.23
(-0.44)			(6.07)				
-0.231	0.337	-	0.525	0.5205	29.52	2.00	0.20
(-0.92)	(1.22)		(1.33)				
-0.231	-	0.337	0.525	0.5205	28.63	0.93	0.24
(-0.93)		(1.22)	(1.33)				

gression. Across the three markets, we observe that autocorrelation tends to be very low in 2004-2007, with a range of [-0.08;0.11], but rises sharply in 2008-2010, with a range of [0.13;0.30], while remaining at a tolerable level.

4.5 Intermediary conclusion

As regards which of simple, exponentially-weighted and implied volatility constitutes the best estimator of future volatility, the results of our analysis seem quite compelling: implied volatility is always at least as good as the other two, and often better. In 2004-2007, the superior informational content of implied volatility is unequivocal. Indeed, not only does implied volatility partly explain the variance of future volatility, but it does so better than past historical volatility and incorporates the information content of the latter. Concerning 2008-2010, it is reasonable to hypothesize that our results were impacted by the violent turbulences of the financial crisis. In that scenario, the three estimators appear somewhat redundant. Nevertheless, these results also demonstrate that the superiority, in a large sense, of implied volatility is robust, for implied was at least as good as the other estimators. Our results suggest that implied volatility is a robust and powerful estimator of future volatility, notably when compared with other usual measures such as simple and exponentially-weighted historical volatility.

5 Introducing GARCH

5.1 The GJR-GARCH(1,1) model

The Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model extends the ARCH model developed by Nobel Prize laureate R. Engle in 1982 as a method of analyzing time-varying volatility. The model considers that future variances depend upon past variances and squared returns; variance is both conditional and autoregressive. It is widely used among practitioners and academic alike to model and forecast volatility. Developed by Glosten, Jagannathan, and Runkle in 1993, GJR-GARCH is an extension that takes into account asymmetries in the response of the conditional variance to an innovation. Focusing on the S&P 500, we decide to use a GJR-GARCH(1,1) model and compare its predictive power to that of implied volatility using two different methods. The model is the following:

$$\sigma_{g,t}^2 = K + \delta\sigma_{t-1}^2 + \alpha\epsilon_{t-1}^2 + \phi\epsilon_{t-1}^2 I_{t-1} \tag{17}$$

where

$$I_{t-1} = 0 \text{ if } \epsilon_t \geq 0 \ , \tag{18}$$

$$I_{t-1} = 1 \text{ if } \epsilon_t \leq 0 \ , \tag{19}$$

$$\epsilon_t = z_t \sigma_t \ , \tag{20}$$

z_t is a sequence of independent and identically distributed variables, and $\sigma_{g,t}$ denotes the GJR-GARCH conditional volatility at time t.

In our first method, we use the data on the S&P 500 from 2004 to 2010 to estimate the parameters of our model by maximum likelihood, before doing an in-sample comparison of GARCH-fitted volatility and implied volatility. In our second method, on each date t when we observe an option from 2005 to 2010, we use the daily volatility data from 2004 up to t to calibrate the GJR-GARCH(1,1) model and forecast the market volatility over the remaining life of the option. We then compare the predictive power of these forecasts with that of implied volatility in 2005-2007 and 2008-2010.

5.2 In-sample GARCH-fitted values v. implied volatility

5.2.1 Methodology

We collect monthly observations of the return of the S&P 500 from 2004 to 2010 on the same dates t we observed historical and implied volatility (see 2.2). At time t, the return is:

$$R_t = ln(S_t/S_{t-\tau_t}) \tag{21}$$

where S_t denotes the level of the S&P 500 at time t and τ_t the days remaining before the expiration of the option whose implied volatility was observed at time t.

Then, we use our collection of 82 returns to estimate the parameters of our GJR-GARCH(1,1) model through maximum likelihood and generate 82 fitted values of conditional volatility. The parameter estimates are presented in Table 7, while the GARCH-fitted values of volatility are plotted in in Figure 3.

Table 7: GJR-GARCH(1,1) parameter estimates

Parameter	Value	T-stat
K	0.0012972	3.7158
δ	0	0
α	0.23906	0.7259
ϕ	1	2.1967

We run the following regressions over 2004-2010 to assess the respective information contents of GJR-GARCH modeled volatility and implied volatility:

$$\sigma_{f,t} = \alpha + \beta_g \sigma_{g,t} + \epsilon_t \tag{22}$$

$$\sigma_{f,t} = \alpha + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{23}$$

$$\sigma_{f,t} = \alpha + \beta_g \sigma_{g,t} + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{24}$$

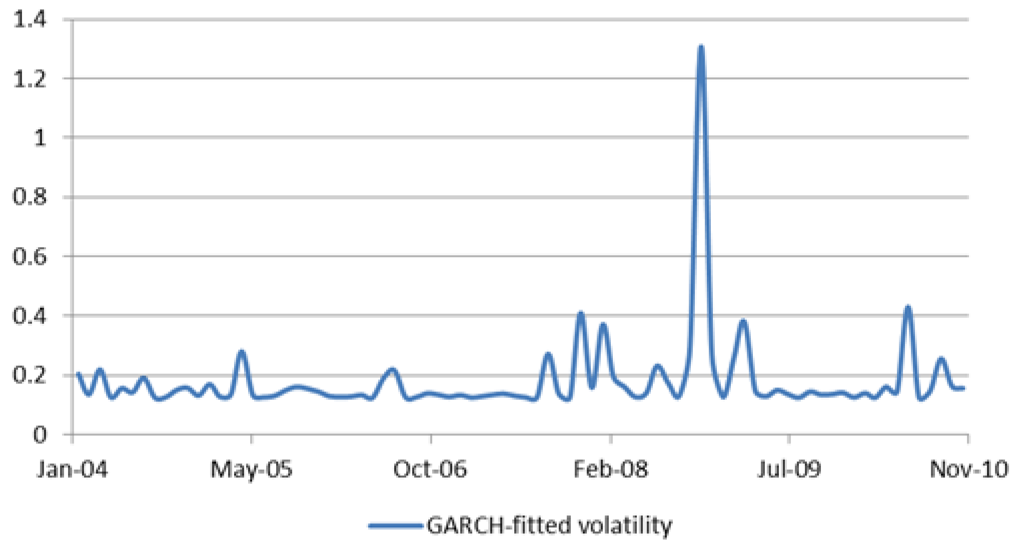

Figure 3: GARCH-fitted volatility

Since the parameters of our GJR-GARCH(1,1) model are estimated using the very same sample data, there is definitely a look-ahead bias in favor of GJR-GARCH volatility. However, considering our results, we deem the shortcoming acceptable.

5.2.2 Empirical results

The results of the regressions are displayed in Table 7. Despite the look-ahead bias entailed by the the in-sample comparison, we observe that the does not perform very well, certainly because of the too-small amount of data used for calibration. In regression 21, we obtain a statistically-significant (t-stat of 4.11) coefficient of 0.599 for fitted volatility, as well as a statistically-significant intercept. The R-squared indicates that the volatility inferred from our GARCH model only accounts for 16.38% of the variance of future volatility, while implied volatility alone accounts for over 71.86%. Moreover, when considering fitted and implied volatility together in regression 23, the R-squared of 71.60% shows that fitted volatility does not add any information content to that already included in implied volatility. We also observe that the coefficient estimate of implied volatility is virtually equal to 1, with a t-stat of 12.51, while both that of fitted volatility and the intercept are very close to 0 and not statistically significant.

Our results suggest that, using this method, volatility fitted with a GJR-GARCH(1,1) model constitutes a relatively poor estimator, even with a look-ahead bias, and that its information content is subsumed by implied volatility, which appears to be a nearly unbiased estimator.

Table 8: S&P500 2004-2010: OLS coefficient estimates of GARCH-fitted volatility and implied volatility

Intercept	$\sigma_{g,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.823	0.599	-	16.38%	17.89	1.12	0.60
(-3.04)	(4.11)					
-0.041	-	1.025	71.86%	50.71	0.93	0.17
(-0.31)		(14.42)				
0.014	0.049	1.006	71.60%	54.80	0.91	0.18
(0.08)	(0.52)	(12.51)				

5.3 Out-of-sample GARCH forecasts v. implied volatility

5.3.1 Methodology

Generally, GARCH models rely on large amounts of high-frequency historical data to forecast a few periods ahead. With our second method, we endeavour to exploit the predictive power of GJR-GARCH(1,1) in a much more efficient way.

From 2006 to 2010, on each date t we collect an observation of implied volatility, use the daily return data from January 2004 up to t to estimate the parameters of our GJR-GARCH(1,1) model, and then forecast the daily volatility of the next τ_t days, where τ_t denotes the number of days before the expiration of the option whose implied volatility we collected on day t. We then compute and annualize the average daily volatility over these τ_t days:

$$\sigma_{g,t} = \frac{\sqrt{252}}{\tau_t} \sum_{i=t+1}^{t+\tau_t} \sigma_{gd,i} \tag{25}$$

By doing so, we are able to avoid the flaws that come with a GJR-GARCH model fed with only a small amount of data.

We subsequently run the following regressions, first on 2005-2007 and then on 2008-2010, to compare the predictive powers of simple historical volatility, GARCH forecasted volatility, and implied volatility:

$$\sigma_{f,t} = \alpha + \beta_g \sigma_{g,t} + \epsilon_t \tag{26}$$

$$\sigma_{f,t} = \alpha + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{27}$$

$$\sigma_{f,t} = \alpha + \beta_h \sigma_{h,t} + \beta_g \sigma_{g,t} + \epsilon_t \tag{28}$$

$$\sigma_{f,t} = \alpha + \beta_g \sigma_{g,t} + \beta_{iv} \sigma_{iv,t} + \epsilon_t \tag{29}$$

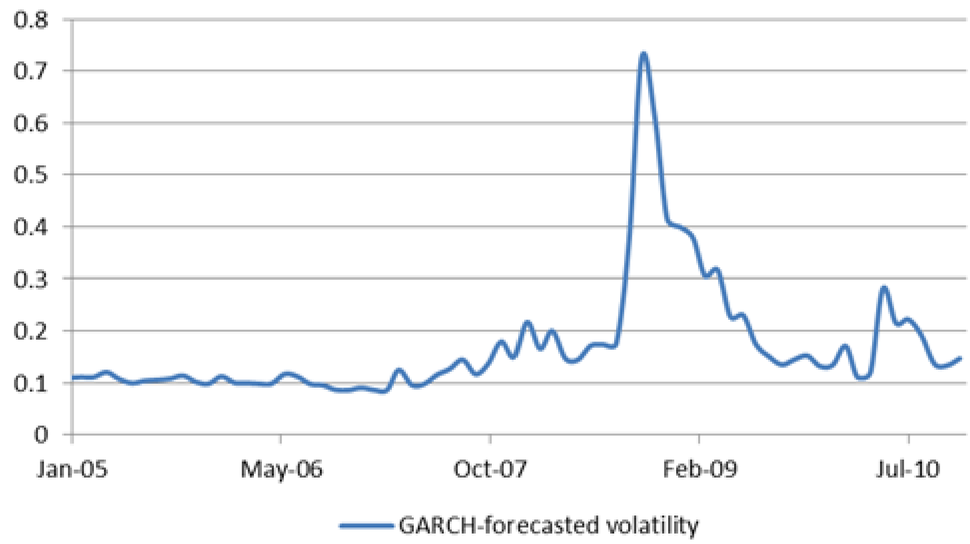

Figure 4: GARCH-forecasted volatility

5.3.2 Empirical results

The results, displayed in Table 8, confirm to some extent what we found before introducing GARCH. We can see that, with a large amount of daily data, the model perform quite well when considered alone. The coefficient estimate of GARCH-forecasted volatility has a statistically-significant value of 1.486 in 2005-2007 and of 0.893 in 2008-2010, while the t-stats of the intercepts fail to reach 1.96. With an R-squared of 40.60 in the first period and of 58.99% in the second, GARCH explains a substantial portion of the variance of future volatility. Furthermore, when combining GARCH-forecasted volatility and simple historical volatility, not only does the R-squared fail to increase above those when GARCH is taken alone, but also the coefficient estimates of historical volatility are insignificant from a statistical perspective. On the other, regression (26) gives us coefficient estimates of 1.607, with a t-stat of 2.57, in 2005-2007 and of 1.044, with a t-stat of 2.45, in 2008-2010 for GARCH-forecasted volatility.

However, in 2005-2007, we observe that regression (26) gives us an R-squared which is higher than that from regression (25), indicating that implied volatility accounts for a greater part of future volatility than GARCH-forecasted volatility. In addition, when combining GARCH and implied volatility, we observe that the R-squared is equivalent to that of regression (26), and that only implied volatility has a coefficient estimate that is statistically significant. In 2008-2010, as we found the fourth part of the present thesis, all measure volatilities seem to perform equally well and to be redundant from an informational perspective.

The results of our second method suggest that, although GJR-GARCH outperforms historical volatility, it is still inferior to implied volatility in terms of forecasting power and does not seem to contain supplementary information.

Table 9: S&P 500: OLS estimates of GARCH forecasted volatility and implied volatility for 2005-2007 and 2008-2010

2005-2007

Intercept	$\sigma_{h,t}$	$\sigma_{g,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
1.108	-	1.486	-	40.60%	4.29	0.31	-0.02
(1.67)		(4.99)					
0.142	-	-	1.093	58.77%	15.96	0.55	-0.07
(0.43)			(7.13)				
1.228	-0.066	1.607	-	38.88%	4.76	7.02	0.00
(1.42)	(-0.22)	(2.57)					
-0.185	-	-0.308	1.260	57.95%	13.64	3.72	-0.04
(-0.28)		(-0.59)	(3.88)				

2008-2010

Intercept	$\sigma_{h,t}$	$\sigma_{g,t}$	$\sigma_{iv,t}$	Adj. R^2	JB	White	Autocorr.
-0.102	-	0.893	-	58.99%	9.11	0.48	0.19
(-0.49)		(7.06)					
0.024	-	-	1.070	55.48%	21.20	0.91	0.26
(0.10)			(6.59)				
-0.0727	-0.139	1.044		57.89%	9.12	1.01	0.20
(-0.32)	(-0.37)	(2.45)					
-0.096	-	0.869	0.031	57.71%	9.47	17.73	0.19
(-0.39)		(1.66)	(0.05)				

6 General conclusion

Our analysis indicates that implied volatility is an informationally-efficient and largely-unbiased estimator of one-month-ahead future volatility. It also shows that implied volatility is superior to simple historical volatility, exponential historical volatility, and GJR-GARCH(1,1)-fitted or -forecasted volatility. In 2004-2007, implied volatility unequivocally outperformed its competitors in forecasting the one-month volatility of the S&P 500 and the DAX, while in 2004-2007 for the DAX and in 2008-2010 across all three markets, it performed equally well. Above all, in every single instance, combining implied volatility with any of the other candidate estimators did not yield a higher R-squared than when implied volatility was used as a lone forecaster, suggesting that implied volatility subsumes the information content of all the other estimators. As in Christensen and Prabahala (1998), we can only conclude that there is undeniable evidence of the superior efficiency of implied volatility as a predictor of future volatility one month ahead. We have also shown that the result has been robust across different markets since 2004.

What is more, in light of our results for the 2008-2010 period, we dare generalize our findings, by formulating the hypothesis that, in times of turbulence and exceptional market volatility, all major estimators share the same information content and forecasting performance regarding one-month future volatility. We leave this statement open for further investigation.

Of course, if the present study confirms the results of recent research, such as Christensen and Prahbala (1998), it also paves the way for more studies. First of all, it would be interesting to check the robustness of the conclusion in a wide range of other markets, particularly with different levels of liquidity. Secondly, it would be interesting to decrease the length of the forecast and see if implied volatility keeps its advantage, notably over GARCH models, when estimating volatility over shorter periods of time. Finally, comparing implied volatility with more sophisticated historical volatility models would perhaps yield different results.

References

[1] Beckers, S.: *Standard deviation implied in option prices as predictors of future stock price variability*, Journal of Banking and Finance (1981)

[2] Black, F. & Scholes M.: *The valuation of options and corporate liabilities*, Journal of Political Economy (1973)

[3] Bodie, Z. & Merton, R. C.: *The informational role of asset prices*, The Global Financial System (1995)

[4] Bollerslev, T.: *Generalized Autoregressive Conditional Heteroskedasticity*, Journal of Econometrics (1986)

[5] Canina, L. & Figlewski, S.: *The informational content of implied volatility*, Review of Financial Studies (1993)

[6] Chiras, D. P. & Manaster, S.: *The information content of prices and a test of market efficiency*, Journal of Financial Economics (1978)

[7] Christensen, B. J. & Prabhala, N. R.: *The relation between implied and realized volatility*, Journal of Financial Economics (1998)

[8] Christensen, B. J. & Hansen, C. S.: *New evidence on the implied-realized volatility relation*, The European Journal of Finance (2002)

[9] Day, T. E. & Lewis, C. M.: *Stock market volatility and the information content of stock index options*, Journal of Financial Economics (1992)

[10] Engle, R.: *Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation*, Econometrica (1982)

[11] Glosten, L. R., Jagannathan, R. & Runkle, D. E.: *On the relation between the expected value and the volatility of the nominal excess return on stocks*, Journal of Finance (1993)

[12] Hansen, C. S.: *The Relationship between implied and realized volatility in the Danish option and equity markets* (1999)

[13] Lamoureux, C. G. & Lastrapes, W.: *Forecasting stock return variance: understanding stochastic implied volatility*, Review of Financial Studies (1993)

[14] Latane, H. & Rendleman, R.: *Standard deviation of stock price rations implied in option prices*, Journal of Finance (1976)

[15] Li, S. & Yang, Q.: *The Relationship between implied and realized volatility: evidence from the Australian stock index option market*, Review of Quantitative Finance and Accounting (2009)

[16] McDonald, R. L.: *Derivatives markets* (2006)

[17] Merton, R. C.: *Theory of rational option pricing*, The Bell Journal of Economics and Management Science (1973)

[18] Pan, J.: *Time-series patterns in equity, Part 1 & 2*, Teaching slides for 15.433 at MIT Sloan School of Management (2011)

[19] Szakmary, A. et al.: *The predictive power of implied volatility: evidence from 35 futures markets*, Journal of Banking and Finance (2003)