

What drives option prices?

Frédéric Abergel¹ and Riadh Zaatour^{2*}

Chair of Quantitative Finance

Ecole Centrale Paris, Grande Voie des Vignes,

92290 Châtenay-Malabry, France.

Abstract

We rely on high frequency data to explore the joint dynamics of underlying and option markets. In particular, high frequency data make observable the realized variance process of the underlying, so its effects on option price dynamics are tested. Empirical results are confronted with the predictions of stochastic volatility models. The study reveals that while the modeling of stochastic volatility gives more robust models, the market does not process information on the realized variance to update option prices.

Contents

Introduction	2
1 The data	3
2 Theoretical framework	4
3 Intraday joint dynamics of option and underlying prices	5
3.1 Epps effect	6
3.2 Regression analysis	7
3.2.1 R^2 of option / future increments regression	7
3.2.2 Multiple linear regression on the future and its realized variance	8
4 Stochastic volatility model predictions	10
4.1 implied R^2	10
4.2 Regression slope versus model delta	13
4.2.1 Hedging performance of the regression slope	13
4.2.2 model free vs model deltas	14
5 Several day analysis : implied volatility perspective	16
5.1 Spot dependence	17
5.2 Pure volatility effect on V_0	18
6 Conclusion	20

*We want to thank BNP Paribas equity research team and especially Boris Leblanc and Casimir Bernard-Mantel for the help they provided to make this work possible. We are also indebted to participants of the conference 'Market Microstructure : Confronting Many Viewpoints' held in Paris from December 6th to December 10th 2010 in which a first part of this work was presented. We also thank participants of the Sixth Annual Conference on Asia-Pacific Financial Markets (CAFPM) of the Korean Securities Association (KSA), held in Seoul, Korea on December 3rd 2011, where a preliminary version of this article was presented. The second author is indebted to Jose Da Fonseca for fruitful discussions.

¹Ecole Centrale Paris, email : frederic.abergel@ecp.fr

²Ecole Centrale Paris, email : zaatour_riadh@yahoo.fr (corresponding author)

Introduction

A huge effort has been made to adapt the Black and Scholes model (Black and Scholes, 1973) in order to explain the observed implied volatility smile phenomenon. Alternative pricing models have been proposed that correct some biases of Black and Scholes model and relax its restrictive assumptions, leading for instance to heavy tails and skewness in log returns. Thanks to nice model properties (such as affinity, as, for example, in Heston Model (Heston, 1993)), and efficient numerical methods (for instance the Fast Fourier Transform algorithm (Carr and Madan, 1999)), or closed formula for local volatility model (Dupire, 1994), many models can be perfectly calibrated to a smooth implied volatility surface. Therefore, they can equally well capture the marginal distributions of the spot, while predicting different dynamical behavior for the underlying as well as for the smile itself.

Many empirical studies have established the need for multifactor models to capture the smile dynamics. Stochastic volatility models or jump diffusion models are from this perspective judged to be superior to the local volatility one *a la* Dupire. One of the major drawbacks of this model, is that it predicts a flattening of the forward smile whereas the formers generate a persisting smile, as observed in option markets.

Moreover, among the features that one can add to the Black and Scholes framework (jumps, stochastic volatility, stochastic interest rates ..), stochastic volatility modeling yields the best dynamical model performance as was established in (Cao et al., 1997).

Consequently, stochastic volatility modeling is the corner stone of option pricing theory. This modifies the replication paradigm, as, under stochastic volatility models, perfect replication by dynamic trading in the stock is no longer possible. Another option has to be introduced to complete the market and the replication now relies on dynamic trading in both the stock and this option, to the extent that researchers now speak of a *volatility delta*. A nice exposition of these ideas can be found in (Wilmott, 2006) and (Rebonato, 2004).

As usual when one deals with option pricing, these stochastic volatility models take an 'implied route' : calibrate model parameters in order to replicate the observed market vanilla option prices, then use the model with these calibrated parameters to price more exotic options.

What interests us in this article is that a stochastic volatility model, besides being an interpolation tool between liquid vanilla prices and non quoted exotic option prices, presents also a scenario of option price formation and evolution. According to stochastic volatility models, option prices are the result of the processing, by the market, of the information on spot and variance stochastic processes. We want to know if the data tells the same story.

Indeed, the availability of high frequency data, and especially of the hardware and software capabilities allowing practitioners and researchers to deal with them, enhanced the accuracy of realized variance estimation. One is now able to build very accurate time series of daily realized variance and even estimate it intra daily, making it a market observable if not a tradable asset. Econometric literature dealing with this subject is huge. See for example (Andersen and Benzoni, 2009) and references therein.

A legitimate question is then to reconsider the scenario of option price formation in the light of this newly available data. Does the scenario presented by stochastic volatility models actually occur in real markets? As the spot process is revealed, option prices adjust to be consistent with it, but does the same occur with the variance process that high frequency data allow to reveal? In other words, are option prices adjusted to be consistent with the realized variance process?

This paper is an attempt to answer these questions. We explore the relationship between realized variance time series and option prices, or equivalently, implied volatility. The followed strategy is to explore the joint dynamics of option and underlying markets. The quadratic variation process of the underlying being observable, we seek, ex-post, its effects on option prices.

The paper is organised as a back and forth between theoretical predictions of stochastic volatility models and

data observation.

In section 1, we describe the TRTH (Thomson Reuters Tick History) data we use, as well as its processing in BNP Paribas. We also present the three stock indices that the study deals with, and some characteristics of their respective markets.

Section 2 presents the stochastic volatility framework and its implications on the dynamics of option prices.

In section 3, we conduct an empirical study of the intraday evolution of option prices. Inspired by the framework presented in the previous section, we isolate the spot component in option price dynamics and try to explain the residual by the realized variance. We conclude that the changing intraday underlying variance does not impact the market view of implied volatility, which can be almost fully explained by the spot dynamics.

In section 4, we compare empirical observations to stochastic volatility predictions, quantified thanks to the calibration of Heston stochastic volatility model to market data. This reveals a comparable weight given to the spot component in the option dynamics by the models and data, but a divergence in the projection of the residual : while the model attributes it to stochastic variance, the data does not confirm that. Moreover, the intraday study gives a measure of option price dependence on the spot, which we call the effective delta. This turns out to be a good hedging parameter, well approximated by the Heston total delta, confirming our method of assessing the spot part in the option price dynamics, and confirming other studies' results on the hedging performance of stochastic volatility models, such as (Cao et al., 1997).

In section 5, armed with the confidence gained through the two last sections that a stochastic volatility model describes well the cinematics we are observing, even if it does not identify the same causes, we generalize the study to a daily time scale. Even if an option changes from day to day (as its maturity is diminished by one day), stochastic volatility models give us some more tools to identify a maturity and strike independent quantity characterizing the option price. We study the influence of daily realized variance on the evolution of that quantity. The result remains unchanged: while the stochastic volatility framework gives a satisfactory phenomenological description of option markets, its structural interpretation is not observable in the data. The market does not process information revealed by realized spot variance to fix option prices. Observation of the realized variance does not improve our predictive power on option price or implied volatility dynamics, as compared to the observation of the spot process alone.

Section 6 summarizes the study.

1 The data

We rely on tick by tick data from TRTH (Thomson Reuters Tick History). The data consists of trade and quote files timestamped in milliseconds. These files are processed by BNP-Paribas Equities & Derivatives Quantitative R&D Histo team. Among other things, this processing consists in storing the data in a special format that saves the necessary disk space, and especially makes data requests very efficient.

For this study, we consider three world indices : Eurostoxx 50, Dax 30 and Kospi 200. These indices have in common a very liquid derivatives market, and are, by other aspects, very different, as discussed in the sequel.

Kospi 200 options are the most liquid exchange listed derivatives in the world as far as trading volume is concerned. In 2010, Kospi 200 options traded in Korea accounted for 70% of the index options volume traded worldwide with more than 3.5 billions contracts¹. This is certainly due to the fact that these contracts are much smaller than those traded on other exchanges. Most traded options on Kospi 200 are those of next month expiry. Almost all the trading activity takes place on the exchange. The most actively traded options can register more than 100000 trade events in one day.

¹source : World Federation of Exchanges, 2010 Derivatives Markets Survey Report, <http://www.world-exchanges.org>

Options on Dax and Eurostoxx are also very actively traded. Their turnover is greater than that of Kospi 200 options, but due to a bigger contract value for Dax and Eurostoxx options on the one hand, and to a greater number of traded maturities on the other hand, the number of trades for any particular option is much smaller than in the Kospi 200 case. In fact, next month and quarterly maturities up to 18 months are actively traded. Even if a non negligible proportion of the trades takes place on the OTC markets, the option order books are very active and quotes are regularly updated (the time scale here is of the order of the second).

By other aspects, these three markets are very different. For example, Dax index is a total return index, which means that dividends are reinvested, contrary to both other indices. Additionnaly, while trading activity on Eurostoxx and Dax derivatives is essentially done by finance professionals, Korean market presents the curious property that approximately the third of trading activity on Kospi 200 options is done by individual householders, a fact unique to this market. Moreover, in the KRX market, options are more liquid than futures. In fact, in 2010, while the number of traded contracts on options exceeded 3.5 billions, the number of traded futures contracts is below 100 millions. See <http://eng.krx.co.kr> for details.

Although the three markets are very different, our investigation gives the same results for all of them, allowing us to interpret these findings as stylized facts of the joint dynamics of option and underlying markets.

2 Theoretical framework

Under the real world probability measure, we assume that stock price dynamics are given by a stochastic volatility diffusion as in the equations below :

$$dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_t^1 \quad (1)$$

$$dV_t = \nu dt + \zeta dW_t^2 \quad (2)$$

$$d \langle W^1, W^2 \rangle_t = \rho dt, \quad (3)$$

where we make the usual (but not at all trivial) assumption that the price C_t of an option on the underlying S is a smooth function of time, spot and underlying variance, i.e. $C_t = C(t, S_t, V_t)$. Therefore, the Itô formula gives us the dynamics of this option price :

$$\begin{aligned} dC_t &= \left[\frac{\partial C}{\partial t} + \mu S_t \frac{\partial C}{\partial S} + \frac{1}{2} V_t S_t^2 \frac{\partial^2 C}{\partial S^2} + \nu \frac{\partial C}{\partial V} + \frac{1}{2} \zeta^2 \frac{\partial^2 C}{\partial V^2} + \sqrt{V_t} \zeta \rho S_t \frac{\partial^2 C}{\partial S \partial V} \right] dt \\ &+ \sqrt{V_t} S_t \frac{\partial C}{\partial S} dW_t^1 \\ &+ \zeta \frac{\partial C}{\partial V} dW_t^2. \end{aligned} \quad (4)$$

Stochastic terms dominate drift terms as their order of magnitude is \sqrt{dt} whereas it is dt for the latter. Additionally, we will work with detrended data and assume that this enable us to ignore the drift term in the above equation, leaving us with the stochastic part only. So we end with the detrended option price process :

$$dC_t = \sqrt{V_t} S_t \frac{\partial C}{\partial S} dW_t^1 + \zeta \frac{\partial C}{\partial V} dW_t^2. \quad (5)$$

Option price is then, as expected, governed by two sources of randomness : the spot, and its stochastic volatility. Notice however that, as this stochastic volatility is itself correlated with the spot, part of its own randomness can be attributed to the noise governing the spot price itself. To isolate the pure spot noise one can write :

$$dC_t = \underbrace{\left(\sqrt{V_t} S_t \frac{\partial C}{\partial S} + \zeta \rho \frac{\partial C}{\partial V} \right)}_{\text{pure spot risk}} dW_t^1 + \underbrace{\zeta \sqrt{1 - \rho^2} \frac{\partial C}{\partial V}}_{\text{pure volatility risk}} dZ_t, \quad (6)$$

where now W_t^1 and Z_t are independent brownian motions.

Recalling that the detrended spot price is $dS_t = \sqrt{V_t} S_t dW_t^1$, we rewrite the above equation in terms of dS_t so as to emphasize the option and spot price dynamics relationship :

$$dC_t = \underbrace{\left(\frac{\partial C}{\partial S} + \frac{\zeta \rho}{\sqrt{V_t} S_t} \frac{\partial C}{\partial V} \right)}_{\text{Effective delta}} dS_t + \underbrace{\zeta \sqrt{1 - \rho^2} \frac{\partial C}{\partial V}}_{\text{pure volatility delta}} dZ_t. \quad (7)$$

The so-called *effective delta* in the above equation carries the total first order sensitivity of the option price to spot movements in this framework. This is the classical delta $\frac{\partial C}{\partial S}$ corrected by an adjustment factor accounting for the correlation between the spot and its volatility. The orthogonal term carries the volatility specific noise, hence the name *pure volatility delta*.

To summarize, we end with a regression relationship between option and spot price increments :

$$dC_t = \beta dS_t + \alpha dZ_t. \quad (8)$$

Obviously, the coefficients of the above regression relationship are time dependent and may be stochastic. In (Mykland and Zhang, 2008) the authors study the problem of inference of such an instantaneous β in a generalized one factor framework of implied volatility dynamics and its implication on option hedging. For our purpose, we assume that the regression coefficients do not vary a lot within one day, which is the time window that we consider to begin our study. The regression of option and underlying price increments can then be conducted, quantifying the underlying spot influence on the option price dynamics. The residual of this regression carries the *pure volatility risk*, for which we will explore, if any, the relationship with the realized variance of the underlying. This is done in the next section.

3 Intraday joint dynamics of option and underlying prices

When dealing with statistical time series analysis, a trade off has to be made between the stationarity assumption on the one hand, which is necessary for the estimation of different quantities of interest such as correlations, volatilities, regressions and so on.. and the length of the data set on the other hand, as this has to be sufficiently long in order to assume that the estimators have converged.

This problem is harder when we deal with option price time series. Indeed, an option fundamentally changes from day to day. For instance, today's ATM 30-day maturity option becomes a 15-day maturity option in two weeks, with a non negligible theta effect that may be hard to quantify without many modeling assumptions. Moreover, the option may no longer be considered ATM due to the spot movements, which completely change its characteristics. So, comparing option prices spanning several days may be misleading as we are not talking about the same financial product (unlike a stock price which is generally the same from day to day). Consequently, some care have to be taken in the study of the joint dynamics of option and underlying prices.

In order to study raw option prices without relying on any modeling hypothesis, we look at intraday dynamics. In the case of intraday price changes, we can consider that the maturity of the option does not change, and generally, spot movements leave the option in the same moneyness zone. With high frequency data, we still have long enough price series to get reliable estimations.

However, there is another issue with high frequency data, not specific to option prices. Indeed, if we want to make a regression study of two or more price time series, we begin by sampling them at a given frequency to synchronize them, and then conduct our regression. In an intraday study with high frequency data, we are then tempted to sample the data at the highest possible frequency in order to obtain the longest possible time series and ensure a better convergence rate of our statistical estimators. Unfortunately, the ideal sampling frequency is not necessarily the greatest possible, because of, among other phenomena, the Epps effect.

3.1 Epps effect

Epps effect (Epps, 1979) is the observed decrease of estimated correlation of price time series as the sampling frequency increases. It is mainly due to the asynchronicity of trading, as well as discretization effects (Münnix et al., 2010). It is well documented for stock/stock price time series. We observe the same effect for the asset and its derivative product.

Table 1 presents an illustration of this effect. We take trades and quotes data for a 22-day maturity call option on Kospi 200 and the future² on the same index, both maturing the same day on March 2008. We consider mid prices for both securities. Estimation of the daily correlation of price returns, sampled at different time interval lengths, is shown in Table 1.

Sampling period in seconds	Estimated correlation	95% Conf interval	99% Conf interval
300	97.10%	[95.5% , 98.1%]	[94.8% , 98.4%]
180	96.48%	[95% , 97.5%]	[94.5% , 97.7%]
150	95.83%	[94.3% , 96.9%]	[93.7% , 97.2%]
120	95.19%	[93.7% , 96.3%]	[93.1% , 96.6%]
90	93.41%	[91.6% , 94.8%]	[91% , 95.2%]
60	93.50%	[92.1% , 94.6%]	[91.6% , 95%]
45	91.56%	[90% , 92.8%]	[89.5% , 93.2%]
30	90.84%	[89.5% , 92%]	[89.1% , 92.3%]
20	87.63%	[86.2% , 88.9%]	[85.7% , 89.3%]
10	80.57%	[79.1% , 81.9%]	[78.6% , 82.4%]
5	69.94%	[68.4% , 71.4%]	[67.9% , 71.8%]
2	52.63%	[51.3% , 53.9%]	[50.9% , 54.3%]

Table 1: Correlation vs Sampling period for a 22-day Call option on Kospi 200 vs the index future. Confidence intervals construction relies on the Fisher Transformation of the sample correlation coefficient, as well as a bivariate normality assumption.

As shown by the confidence intervals of Table 1, asymptotic correlation is reached in few minutes for pairs of index future vs option considered. The order of magnitude of the phenomenon in stock/stock case is of 20 to 30 minutes. This shortening of the time scale can be explained by delta effect of future's price on option's price: this mechanical effect leads to a fairly fast adjustment of option quotes when future prices move.

There exist some techniques to deal with asynchronicity of data in covariance estimation. We can cite (Hayashi and Yoshida, 2005) who propose an unbiased and consistent estimator of realized covariance of two diffusions observed asynchronously and in discrete time, that does not require any data pre-processing. Another variance/covariance estimator using all available data without extra manipulation is proposed in (Mancino and Malliavin, 2002). Tests of the first estimator with our data give satisfactory results, in

²For all our studies in the sequel of the joint dynamics of option and underlying prices, we take the future price as a proxy for the underlying index. This is a good proxy for many reasons : the future is tradable and is the main hedging instrument for options on the index. As we are not interested in trends but in the noise governing the dynamics, we can assume it is the same for the future and for the spot itself. Even if the future has in general a stochastic interest rate component, this can be ignored as we consider short maturities, enabling to assume that interest rates are deterministic. In fact, future/index correlations for the maturities we consider are always greater than 99%.

accordance with a few minutes data sampling period. As we will need to conduct regression studies (and even multidimensional regressions), we rely on time series sampling. Doing that with a period of 5 minutes seems to be a good trade off between the time needed for the Epps effect to disappear, and the need to retain enough data for our estimators to give reliable results, as we work in an intraday setting. This 5-minute sampling allow us also to assume that microstructure noise has disappeared.

3.2 Regression analysis

As usual in high frequency data, the regression analysis is conducted on price increments of the option and the future. The regression relationship established in section 2 is well observed in the data. Typical option price versus future price increment plots look like the Figure 1.




Figure 1: Evidence of the linear regression relationship of option price increments and future price increments as observed in the data (here Kospi 200). As expected, the slope is positive for the Call and negative for the Put.

We perform this regression on our universe of options and futures. Every day, the actively traded (for kospi 200) or quoted (for Dax and Eurostoxx) options are considered, as well as the front maturing future. Even if the three markets considered are very different by multiple aspects -as noticed in the data description- the results are roughly the same: a very good linear regression relationship (nearly 0 p-values) with a very high coefficient of determination R^2 .

3.2.1 R^2 of option / future increments regression

The observed coefficient of determination R^2 is very high, and is most of the time very close to 1. In Figure 2, we plot as an example the calculated R^2 for all the options on Eurostoxx and Dax that we consider. Every point in the graph is the R^2 obtained for a 1-day regression of option price increments on future price increments.

Even if some of the regressions in the Figure 2 reveal a bad quality of fit as shows the small R^2 value, further investigation allows to discard them. Indeed, we inspect individually the days and options giving an R^2 which is less than 0.5. In all these cases, the individual investigation reveals the presence of an outlier that highly impacts the regression. In Figure 3, such a situation is presented for a Dax option: in this case, at two different times, errors in the fed best ask quote were responsible for a large move in the option mid price which feeds our regression analysis. When these outliers are taken away (and they are indeed corrected in the market data), the R^2 takes a very high value, which is clear on the graph.

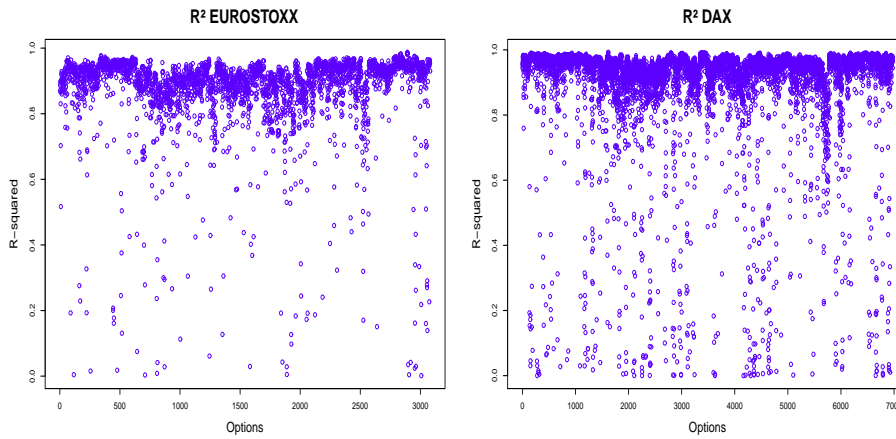


Figure 2: Observed R^2 of option/future increments regression for Eurostoxx and Dax.

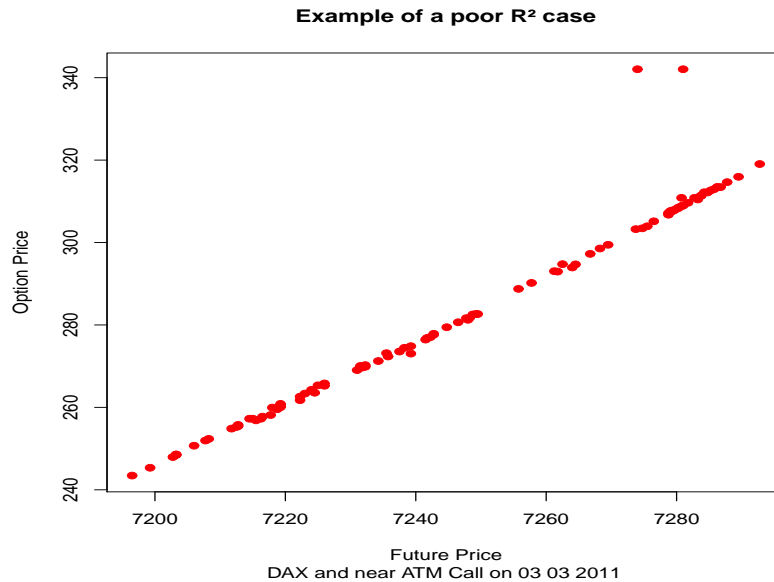


Figure 3: Example of a poor R^2 case : presence of the two outliers largely accounts for the bad quality of fit denoted by a small R^2 .

As a conclusion, most of option price variance is explained by the underlying variance, but an orthogonal factor does exist, hence the strictly less than one R^2 value. Motivated by the stochastic volatility model intuition, we seek this orthogonal factor in the realized variance time series.

3.2.2 Multiple linear regression on the future and its realized variance

Recall that the purpose of the regression we are conducting is to relate the option price dynamics to the underlying dynamics, as well as the variance of underlying dynamics, as suggested by stochastic volatility models. The second step of our intraday study is then to calculate the intraday realized variance time series. To estimate these variance time series, we sample the futures price data for the three indices every second, and

then use the classical variance estimator on different time windows, namely 15, 20 and 25 minutes:

$$V_t = A \times \sum_{i=1}^N \ln\left(\frac{S_{i+1}}{S_i}\right)^2, \quad (9)$$

where A is the adequate annualization factor for each time window.

This gives us three time series of realized variance, one for each time window considered. Motivated by the above noted considerations on the Epps effect, we then construct 5-minute spaced time series of realized variance. For example, for variance time series constructed with a 25 minute window, begin at 9:30, take the future prices between 9:05 and 9:30 sampled every second and calculate the above estimator, obtaining V_1 , then move to 9:35, take the future prices between 9:10 and 9:35 sampled every second and compute the same estimation obtaining V_2 and so on and so forth... We pick also the future prices and option prices at the times of variance calculation to do our regressions.

As an example, in Figure 4, we plot the spot of Eurostoxx against its realized volatility (defined as the square root of variance, and calculated in a 20 minute window here). Notice the decreasing volatility during the middle of the day due to the decreasing market activity during lunch time. Notice also the overall negative correlation of volatility movements and price movements, but near the end of the day, where intense activity before the closing raises the volatility even if the price goes up.




Figure 4: This figure ensures the plausibility of the realized volatility estimated : observe the lunch time effect in the middle of the day.

We then make a multiple linear regression of option price increments on spot and variance increments. Over-all R^2 of the total regression remains unchanged compared to the one factor regression of the preceding subsection, as well as the slope of the spot factor. However, the regression hypothesis on the realized variance is clearly rejected with p-values exceeding 0.5.

Even if Figure 4 advocates the reliability of our realized variance estimator, it can be argued that our high sampling frequency of future prices used to estimate the realized variance time series induces spurious autocorrelations of returns that can be due for example to order splitting, and that inflates the realized variance estimator.

To avoid such pitfalls, we perform the same analysis with another realized variance estimator that is less sensitive to the market microstructure noise, namely the Garman-Klass variance estimator (Garman and Klass, 1980):

$$V_t = \frac{1}{2} (h_t - l_t)^2 - (2 \ln(2) - 1) (c_t - o_t)^2, \quad (10)$$

where h_t , l_t are respectively the highest and lowest log-price for the considered time window. And o_t (respectively c_t) is the opening (respectively the closing) log-price for this time window.

Using this estimator leads to the same results. The intraday realized variance does not have any effect on the option price dynamics. Option price variance is almost entirely explained by the underlying price variance. On some days, their correlation seem to be perfect. Observe for instance Figure 5 graph of a Eurostoxx Put option vs the future:




Figure 5: The anticorrelation of intraday put price and future price seems perfect.

These empirical observations are in total contradiction with the intuition behind stochastic volatility models. In fact, when one sees Figure 5, one wonder if the stochastic volatility is worth modeling at all! Does an R^2 of 0.9 let any place for another factor for modeling the option price dynamics? How does this quantitatively compare to stochastic volatility model predictions?

4 Stochastic volatility model predictions

4.1 implied R^2

We recalled in section 2 that the option price dynamics in the stochastic volatility framework are driven by two factors: one coming from the stock and the second being what we identified as a pure volatility risk, namely:

$$dC_t = \underbrace{\left(\sqrt{V_t} S_t \frac{\partial C}{\partial S} + \zeta \rho \frac{\partial C}{\partial \sigma} \right)}_A dW_t^1 + \underbrace{\zeta \sqrt{1 - \rho^2} \frac{\partial C}{\partial V}}_B dZ_t. \quad (11)$$

Empirical investigation reveals a very high weight of the spot price in the option price dynamics, which explains the very high observed R^2 . But what do stochastic volatility models quantitatively predict? Do they equally spread option price variance over the underlying and its variance, or do they project the risk on one factor preferably to the other?

To answer this question, we have to quantify the weights (A and B) in Equation (11) relatively to each other. These weights depend on unobservable quantities, such as the volatility of variance ζ or the underlying/variance correlation ρ . Fortunately, these quantities are not altered by an absolutely continuous change of probability. So, the calibration of a pricing model specified under the risk neutral measure should make them observable, giving us orders of magnitude for the weights governing the spot process and the pure variance process in the option price evolution, and allowing us to compare model predictions to the data.

In order to design a benchmark, we choose to calibrate Heston model to the observed option prices. Under the risk neutral measure, this model writes

$$dS_t = rS_t dt + \sqrt{V_t} S_t dW_t^1 \quad (12)$$

$$dV_t = \kappa(\theta - V_t) dt + \zeta \sqrt{V_t} dW_t^2 \quad (13)$$

$$d \langle W^1, W^2 \rangle_t = \rho dt. \quad (14)$$

This choice is naturally explained by the analytical tractability of Heston model. Indeed, the model being affine, its characteristic function is known in closed form and the pricing by FFT as described in (Carr and Madan, 1999) is straightforward, allowing a fast calibration, pricing and greeks calculation.

To design the calibration set, we rely on historical data for rate curves (EUR and KRW), and collect liquid option mid prices for our three indices. This leaves us with the next month maturity options for the Kospi 200³, and with maturities less than 18 months for the Dax and Eurostoxx. At the end of every trading day of our three indices (precisely half an hour before the close), mid-prices of actively traded options are observed for every traded maturity. Put-Call parity relation enables us to determine the forward price of each index. These markets being very liquid, this is a reliable method to calculate the forward. Indeed, the values of the forward we find by applying the relation to different option strikes generally differ by less than 0.05%, this difference being explained by asynchronicity of price contributions and market microstructure noise. We then take the median forward given by the Put-Call parity relation applied to near the money strikes.

We define our objective function as the sum of squared differences between model and market prices as in (Cao et al., 1997), and conduct the minimisation procedure thanks to (Lourakis, Jul. 2004).

In Table 2, we give mean calibration results and their standard deviations for the three indices for the period of March to April 2011. Globally, orders of magnitude of the parameters correspond to what is generally reported in the literature, see for example values in (Gatheral, 2006). The only strange value concerns the κ of Kospi 200: as we have only one smile, the speed of mean reversion is really huge.

³liquidity is concentrated in contracts with the nearest expiration in the Kospi 200 options market. This can be explained by the huge proportion of individual investors, which is due to many reasons like the small size of the contracts and individuals' appetite and enthusiasm for trading. An analysis done in 2001 but still relevant by many aspects, can be found here : <http://www.futuresindustry.org/fi-magazine-home.asp?a=994>

Index	σ_0	κ	θ	ζ	ρ
Eurostoxx	0.2255 (0.0289)	1.8893 (0.6072)	0.0778 (0.0066)	0.5739 (0.1653)	-0.9072 (0.0847)
Dax	0.2027 (0.0274)	1.9887 (0.6128)	0.0677 (0.0402)	0.68311 (0.2164)	-0.7698 (0.0455)
Kospi 200	0.1737 (0.0303)	9.5877 (6.7695)	0.1990 (1.0397)	0.9892 (0.3067)	-0.5813 (0.0854)

Table 2: Mean and standard deviation of the calibrated Heston parameters on the set of actively contributed options for the months of March and April 2011.

Recall that our aim is to have orders of magnitude of the weights in Equation (11) relative to each other, so after the calibration of the models, we have to calculate the derivatives of the price with respect to the spot and to the volatility, to obtain the quantities A and B of Equation (11). These quantities are easily calculated within the Heston model by the same Fourier transform method leading to the option price, where some algebra is applied to calculate derivatives of the characteristic function with respect to the spot and to the volatility. Notice that even a calculation by finite differences is fast enough. The values obtained for A and B enable us to calculate the ratio $\frac{A^2}{A^2+B^2}$ which is the ratio of the variance carried by the underlying dynamics to the total variance of option price increments. It is then the R^2 of a regression of option prices on spot prices implied by our stochastic volatility model.

Figure 6 shows a plot of this implied R^2 . To our surprise, the results are not far from the really observed ones: the implied R^2 in a stochastic volatility framework is still very high (around 0.9). Individual investigation of poor values reveal poorly calibrated options by the model and can be safely ignored. The result is therefore in accordance with the empirical finding. Stochastic volatility model attributes the major part of option price variance to the underlying price movements, in accordance with observation in the data. Option price variance due to the underlying volatility dynamics is far smaller in these models. To push the investigation further, we see in the next section if the model predicts the same slope as seen in the data in the option/underlying regression relationship.

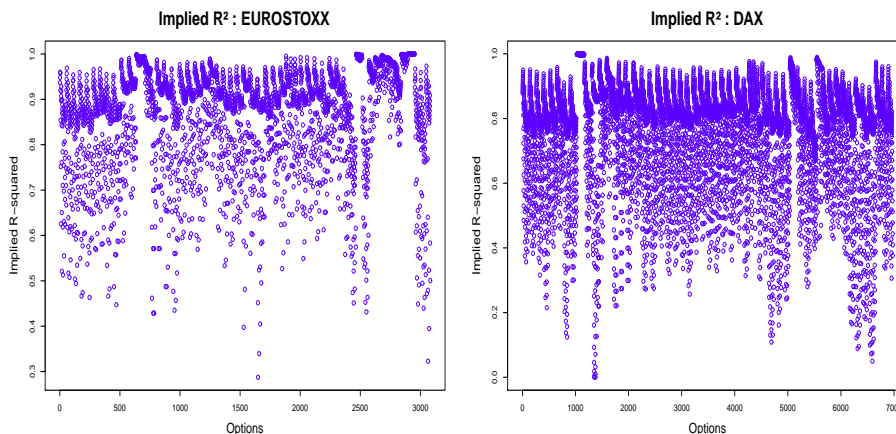


Figure 6: implied R^2 corresponding to the ratio of the variance explained by the spot process to the total variance of the option price process in a stochastic volatility framework, where we calibrate a Heston model to the data to calculate values of the different parameters in equation (11).

4.2 Regression slope versus model delta

Slopes of the regressions in Section 3 are always less than 1 in absolute value and, as expected, positive for Calls and negative for Puts. Recall that these slopes represent the total spot sensitivity of the option in our stochastic volatility framework, and were interpreted as being the *effective delta* of the option in the regression analysis. If our intuition is correct, these slopes give a model free hedge ratio of the option allowing to hedge all the spot dependence. A legitimate question is then if this *model-free* delta is a robust hedge ratio, and how does it compare to the Heston model total delta, namely :

$$\Delta_{Heston} = \left(\frac{\partial C}{\partial S} + \frac{\zeta \rho}{\sqrt{V_t} S_t} \frac{\partial C}{\partial V} \right) \quad (15)$$

4.2.1 Hedging performance of the regression slope

In order to quantify the robustness of this model-free hedge ratio, we perform the following experiment :

- For every index, we collect the options that have been actively traded for at least 20 days since their inception. Even if we do not require these days to be consecutive, it is actually the case in practice, as the option is actively traded when it is relatively near the money, and virtually disappears when it is deeply IN our OUT of the money, or actually disappears when it matures, which are the two retained cases for closing the position as explained later. That leaves us with 975 options on Kospi 200, 527 options on Eurostoxx, 748 options on Dax, for the period between January 2007 and April 2011.
- For every option, we consider that the trader takes a short position on the first day of data availability. The retained price C_0 is the median of best BID quotes of the last 15 minutes of the trading day sampled every second. At the same time, the trader performs the regression $dC_t = \beta dS_t + dZ_t$ of option price increments on future price increments during the trading day and uses the calculated *effective delta* to enter a future position (on the front maturing future). The retained price of the futures is the median mid price of the last 15 minutes of the trading day sampled every second. As the spread on future prices for such liquid indices is very tight, working with mid price instead of bid or ask price depending on the sign of the position simplifies the study without impacting the results. As no cash is needed to enter the futures position, the cash position at inception of the portfolio consists in C_0 . Notice also that we take median prices instead of mean prices as taking the median automatically gets rid of outliers and errors in the data.
- At the end of every trading day, the cash position earns the risk free rate, that we calculate based on historical Euribor rate curves for Eurostoxx and Dax, and Koribor rate curves for the Kospi 200. The trader re-performs the regression and updates the delta position, earning $\Delta_{t-1} \times (S_t - S_{t-1})$ that he adds to his cash position.
- Close the option position if the option goes out the forward moneyness region $[0.85, 1.15]$ as it will no longer be actively traded, or on the eve of the maturity day, by buying back the option. The retained price C_T for the option is the median of best ASK quotes of the last 15 minutes of the trading day sampled every second. Closing the option at the eve of its maturity day avoids us to enter in the subtleties of the regulations of the last trading day of options and futures. That simplifies the analysis without impacting the conclusions.
- For each option, we obtain then a final PnL, that we transform into a dimensionless number, by scaling it by the initial investment, i.e. relative pNL = $\frac{\text{PnL}}{C_0}$. This will make comparison between different options and even between indices easier and allow aggregation of data.

In Table 3, we report means and standard deviations of calculated relative PnL for the three indices under consideration. The three of them are positive, gaining around 20% of the initial investment. This confirms that the *effective delta* is indeed a good hedging ratio. The strictly positive sign of the final PnL in

the three cases is not at all surprising : banks take margins! Notice that these margins seem to be greater in Kospi 200 option market. This can be explained by the great number of non professional agents in this market.

Index	mean relative PnL	std dev
KOSPI 200	0.2489727	0.4992516
EUROSTOXX 50	0.2057907	0.3349121
DAX 30	0.2141580	0.2456583

Table 3: Means and standard deviations of the relative PnL of the model-free hedging strategy.

For completeness, we also plot in Figure 7 the density of relative realized PnL of this model-free hedging strategy.

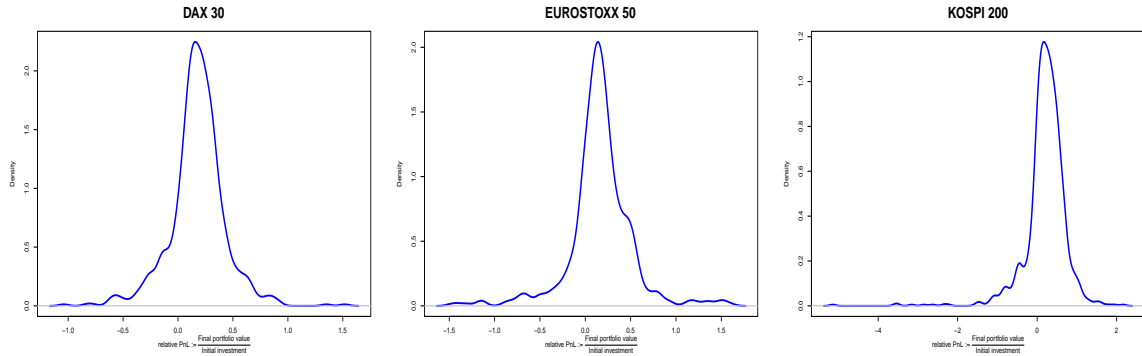


Figure 7: Density of the relative PnL of a model-free-delta-hedged option position.

We conclude that this model-free hedge ratio is robust enough. We compare it in the sequel to other hedge ratios obtained by pricing models.

4.2.2 model free vs model deltas

The regression analysis of the option price dynamics revealed the option “realized delta”. We stated in section 4.2.1 that this is indeed a robust model-free hedge ratio. In this section, we compare this hedge ratio to other model-based deltas. We obtain typically figures looking like Figure 8 :

Candidate hedge ratios are listed below:

- **Nd1** is the delta obtained by simply inverting the Black Scholes formula to obtain the implied volatility Σ of the option, and then calculating the Black and Scholes delta: $\Delta = N(d_1)$ where N is the standard normal cumulative distribution function, and $d_1 = \frac{\ln(\frac{S}{K}) + (r + \Sigma^2)T}{\Sigma\sqrt{T}}$.
- **Local vol** is the delta obtained for a local volatility model *a la Dupire* that we calibrate to the observed Dax implied volatility surface.
- **BS Delta** is the total Black Scholes delta, which takes into account the implied volatility smile: $\Delta = \frac{\partial C}{\partial S} + Vega \frac{\partial \Sigma}{\partial S}$.
- **Heston** is the local minimum variance hedge ratio in the Heston model : $\Delta = \frac{\partial C}{\partial S} + \frac{\zeta \rho}{\sigma S} \frac{\partial C}{\partial \sigma}$.
- **Realized** is the regression coefficient of option price increments on future price increments as in section 4.2.1.

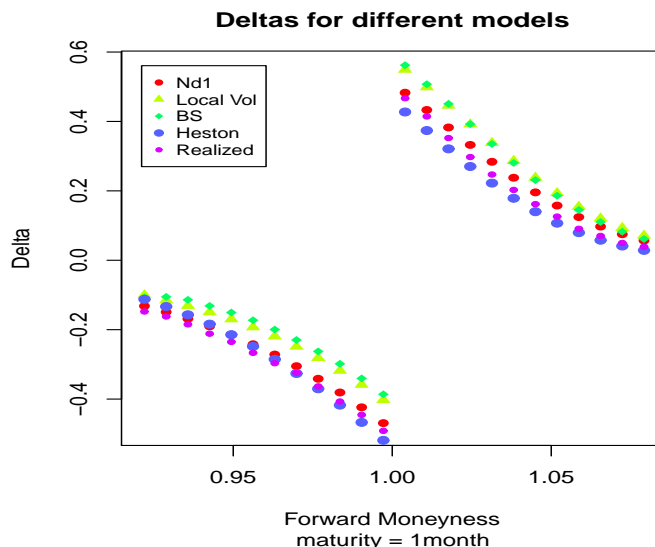


Figure 8: Comparison of different delta measures on one month Dax options.

Observe that **Local vol** delta and **BS Delta** are close to each other, as both of them take into account the smile without adding a dynamical part (dependence of the smile itself on the spot). Heston model yields the closest results to the model-independent effective delta, which proved to be a robust hedge ratio in the study of section 4.2.1. This is conform to the findings of (Cao et al., 1997) that modeling stochastic volatility yields the best dynamical model behavior, i.e. the best hedge performance.

Notice however that practitioners, when using local volatility or Black Scholes volatility, generally add a dynamical term in the smile, letting the volatility surface move as the spot moves. In this case, implied volatility surface is parametrized as $\Sigma = \Sigma(S, m)$ where m is the moneyness, say $m = \frac{K}{S}$.

The first dependence in S isolates the dynamics of the implied volatility, whereas the second reflects the mechanical volatility change when the spot moves, changing the moneyness of the option. We then have :

$$d\Sigma = \frac{\partial \Sigma}{\partial S} dS + \frac{\partial \Sigma}{\partial m} dm \quad (16)$$

where $\frac{\partial \Sigma}{\partial m}$ is the skew of the smile, and $dm = \frac{\partial m}{\partial S} dS = -\frac{K}{S^2} dS$.

So finally :

$$C = C\left(S, \Sigma(S, m(S))\right) \quad (17)$$

And the total delta writes :

$$\frac{dC}{dS} = \Delta_{BS} - \underbrace{\text{Vega}_{BS} \times \text{Skew} \frac{K}{S^2}}_{\text{static correction}} + \underbrace{\text{Vega}_{BS} \times \frac{\partial \Sigma}{\partial S}}_{\text{dynamic correction}} \quad (18)$$

Comparison, at the end of the trading day, of the total delta with the effective delta calculated by regression can give an idea of the quality of the dynamical correction. We will return to these considerations in the sequel.

We conclude from this section that stochastic volatility models do well in capturing dynamical option price properties. This answers positively the question raised in the previous section of whether this feature is worth modeling. The answer is fortunately yes, and the implied R^2 study as well as the delta study of this section

reinforce the well established results of empirical research dealing with robust dynamical model performances of stochastic volatility models.

However, stochastic volatility models tell a different story of option price formation and evolution than that found in the data in section . In the model, the underlying stochastic variance is a factor driving option prices, which could not be observed in the data.

A criticism that can be made to the study in section is that of time scales. In fact, small time scales considered can result in a noisy, or non convergent realized volatility estimator. Larger time scales, however, will result in a fewer data points within a trading day, making any intraday statistical study worthless.

Recall that one problem when dealing with options is that the nature of a given option changes from day to day. Its maturity as well as its moneyness are not the same, and hence, so are its dynamical properties. In order to be able to enlarge the time scale of the study, we must rely on a maturity- and moneyness-independent quantities, enabling us to treat it as being the same financial quantity from day to day. This will be done in the next section.

5 Several day analysis : implied volatility perspective

The previous section made us more confident in the stochastic volatility paradigm as a reduced form description of the option price dynamics. With this regained confidence in this framework, we push its analysis further in order to come with a quantity that characterizes option price dynamics, without depending on a maturity nor a moneyness measure.

Recall that in Heston model the instantaneous variance writes

$$dV_t = \kappa (V_\infty - V_t) dt + \zeta \sqrt{V_t} dW_t, \quad (19)$$

and the expected instantaneous variance in this model obeys a simple first order ODE, whose solution is

$$\mathbb{E}[V_T] = V_\infty + (V_0 - V_\infty) e^{-\kappa T}. \quad (20)$$

The expected total variance from time 0 to time T, which in a deterministic time dependent volatility model would be the Black Scholes implied variance, as established for example in (Lee, 2005) is :

$$\mathbb{E} \left[\frac{1}{T} \int_0^T V_t dt \right] = V_\infty + (V_0 - V_\infty) \frac{1 - e^{-\kappa T}}{\kappa T}. \quad (21)$$

This functional form can then be applied to fit the term structure of the implied volatility. Implied volatility expansions as in (Lewis, 2000), give us a second order polynomial form of the smile as a function of the moneyness. We then apply these functional form as well as the term structure parametrization to fit option smiles.

We end with

$$\Sigma_{K,T} = \Sigma_{ATM} + Smile \times \ln \left(\frac{K}{S} \right) + Curve \times \ln \left(\frac{K}{S} \right)^2 \quad (22)$$

$$\Sigma_{ATM} = \sqrt{V_\infty + (V_0 - V_\infty) \frac{1 - e^{-\kappa T}}{\kappa T}} \quad (23)$$

Notice that in practice, some care is taken to avoid volatility explosions on the wings and the polynomial parametrization is tweaked for very high and low strikes in order to obey non arbitrage conditions established in (Lee, 2004).

This is a very pleasant parametrization of the smile that results from stochastic volatility dynamics. Not only does it describe the smile in a simple polynomial form, but also the presence of three parameters is coherent

with empirical studies of dynamical behavior of volatility surfaces. Indeed, (Cont and da Fonseca, 2002) find that daily variations of implied volatility surfaces can be satisfactorily explained by three factors. The first of them presents the overall shift in the volatility surface, and accounts for more than 90% of the variability. This can be read in the quantity V_0 of the suggested parametrization. Indeed, V_0 characterizes the overall level of the smile surface. Moreover, this is a maturity and moneyness independent quantity. It is therefore a good candidate for our study.

In (Durrleman, 2008), the author establishes the convergence of at the money implied volatility to the spot (continuous part) volatility in a general model with a Brownian component and a jump component of finite variation. In our continuous diffusion setting, this gives a nice financial interpretation of V_0 which is the zero-maturity limit of at the money implied variance, and then, according to (Durrleman, 2008), the spot variance.

We then translate the question on the dynamics of option prices in terms of this quantity : what drives V_0 dynamics? More precisely, we expect to see a spot component in the dynamics of V_0 , and we will seek for a stochastic volatility component as done previously with raw prices. Independence of V_0 with respect to maturity and moneyness will enable us to aggregate several day data in our observation. We will not be constrained to estimate realized volatility in small time windows, circumventing the problem raised about the time scale of the study of section 4.2.2.

The first task is to isolate the spot component in the V_0 dynamics.

5.1 Spot dependence

V_0 is defined in a Heston stochastic volatility setting as the limit of at the money implied variance for short maturities. As such, it is equivalent to the spot instantaneous variance according to (Durrleman, 2008). It is naturally a stochastic quantity.

In the Heston model framework, (Gatheral, 2006) gives a formula for the conditional expectation of the instantaneous variance :

$$\mathbb{E}[V_t|x_T] \approx \hat{V}_T + \rho\zeta \frac{x_T}{W_T} \int_0^T \hat{V}_s e^{-\lambda(T-s)} ds, \quad (24)$$

where :

$$x_t = \ln \frac{S_t}{K} \quad (25)$$

$$W_T = V_\infty T + (V_0 - V_\infty) \frac{1 - e^{-\kappa T}}{\kappa} \quad (26)$$

$$\lambda = \kappa - \frac{1}{2}\rho\zeta \quad (27)$$

$$\hat{V}_s = \left(V_0 - V_\infty \frac{\kappa}{\lambda} \right) e^{-\lambda s} + V_\infty \frac{\kappa}{\lambda}. \quad (28)$$

The interesting conclusion of these equations is that when V_0 is calibrated for $K = S_0$, it evolves linearly with respect to $x = \ln \frac{S}{S_0} \approx \frac{S}{S_0} - 1$, when S is not far from S_0 .

In order to check in the data if this linear approximation is plausible, we fit every five minutes our parametrized implied volatility model to option mid quotes, then we plot the so obtained pairs (S, V_0) . The Figure 9 shows a typical kind of graph, here obtained for EUROSTOXX. Modeling linearly the response of V_0 to the spot movements is clearly a good approximation. Unsurprisingly, the slope of the regression is negative, reflecting the leverage effect of spot movements in implied volatility dynamics. That was clear in the equations above, as the slope is proportional to the spot/volatility correlation ρ of Heston model, which is usually a negative



Figure 9: Graph of intradaily calibrated V_0 versus spot (for Eurostoxx, on April 12th 2011) : Evidence of the linear relationship as a good approximation.

parameter (as found in our estimations).

Therefore :

$$V_0 = V_0^0 + \beta \times \ln \frac{S}{S_0} \text{ where } V_0^0 = V_0(S_0)$$

Thus, the fixed strike volatility parametrization (near the money, to avoid arbitrage problems on the wings) becomes :

$$\Sigma_{K,T} = \Sigma_{ATM}(S) + Smile \times \ln \left(\frac{K}{S} \right) + Curve \times \ln \left(\frac{K}{S} \right)^2 \quad (29)$$

$$\Sigma_{ATM}(S) = \sqrt{V_\infty + \left(V_0^0 + \beta \times \ln \frac{S}{S_0} - V_\infty \right) \frac{1 - e^{-\kappa T}}{\kappa T}}. \quad (30)$$

Notice also that now, the option delta writes

$$\Delta = \frac{dC}{dS} \quad (31)$$

$$= \frac{\partial C}{\partial S} + \frac{\partial C}{\partial \Sigma} \frac{\partial \Sigma}{\partial S} \quad (32)$$

$$= \Delta_{BS} + Vega_{BS} \times \left(\frac{\partial \Sigma_{ATM}(S)}{\partial S} - \frac{Smile}{S} - Curve \times \frac{2}{S} \times \ln \left(\frac{K}{S} \right) \right), \quad (33)$$

where $Vega_{BS} \times \frac{\partial \Sigma_{ATM}(S)}{\partial S}$ accounts for the before mentioned dynamical correction of the delta.

5.2 Pure volatility effect on V_0

In order to isolate the spot effect in V_0 , we have to come with an estimate of β .

β is itself stochastic and depends on market conditions. A careful look at the implied volatility parametrization reached so far tells us that a zero β indicates a pure sticky delta regime. We can also calculate a value $\beta(S)$

for which the smile would be sticky strike⁴. For our parametrization this can not be fully achieved and these are rather good news, as a fully sticky strike smile is arbitrable as established in (Daglish et al., 2007). So finally β characterizes the regime of volatility which depend on market conditions, as described in (Derman, 1999).

An exhaustive analysis is then conducted :

- The smile is calibrated to a market snapshot every 5 minutes, giving as series of V_0
- The β process is given by a daily β estimation.

All this preparation enable us to study the spot orthogonal component of V_0 . Recall that our purpose is to seek if realized volatility has an effect on the dynamics of options prices that we capture here thanks to V_0 .

In order to seek this effect, we calculate the daily realized volatility. We estimate it thanks to the quadratic variation estimator as in section 4.2.2. We then estimate the following regression :

$$V_0^{close} - V_0^{open} - \underbrace{Slope \times \frac{\Delta S}{S}}_{\text{spot effect}} = \alpha \times \underbrace{(realizedVol - V_0^{open})}_{\text{orthogonal movement}}$$

On the right left hand side of the above equation is the daily evolution of V_0 from which we substracted the spot effect as explained by the β . In the right hand side, the algebraic quantity $(realizedVol - V_0^{open})$ captures the realized volatility effect. It is written as a difference between actually realized volatility, and V_0 at the open of the market. The intuition behind that is that V_0 being an implied volatility, it is the guess of market agents of future realized volatility. We are therefore asking the question if market agents, at the end of the day, update their new volatility guess (V_0^{close}) based on the volatility that was actually realized, raising it when realized volatility raises, and decreasing it in the opposed case.

Figure 10 and in Figure 11 are scatter plots illustrating this regression relationship. The graphics do not exhibit a clear linear dependence. Even if we plotted a regression line, its slope is very close to 0.

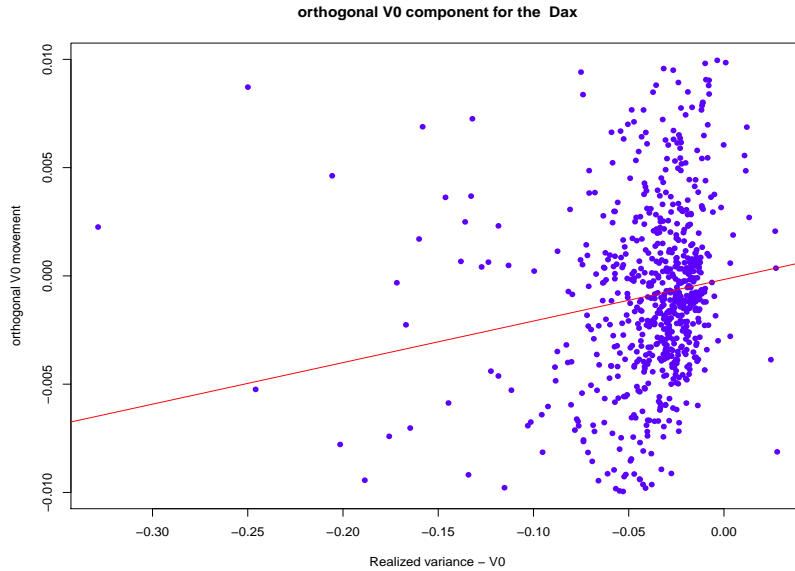


Figure 10: Effect of Realized Volatility on V_0 dynamics for the DAX.

⁴it is the value of β for which $\frac{\partial \Sigma(K,T)}{\partial S} = 0$




Figure 11: Effect of Realized Volatility on V_0 dynamics for the EUROSTOXX.

The criticism made to the intraday study that realized variance estimator may not have converged cannot be made here, as we estimate realized variance on the hole day. There is no dependence of option prices on realized quadratic variation in the data. All is in the spot, and even if a second factor is needed to capture dynamical properties of option prices, the realized quadratic variation does not help to predict it.

6 Conclusion

As a conclusion, we see that while there is a great temptation to see the stochastic volatility models as structural ones, deducing option prices and their dynamics from fundamental market quantities which are the spot and its instantaneous variance, with the latter made observable thanks to the availability of high frequency data and estimation models, the reality of the options markets completely denies that.

Stochastic volatility models are just reduced form models. They meet a lot of success in the description of the overall market behavior, but do not give the structural mechanism governing option price formation and its evolution.

References

- Torben G. Andersen and Luca Benzoni. *Handbook of Financial Time Series*, chapter Realized Volatility. New York, NY: Springer Verlag, 2009.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3):637–654, May-June 1973.
- Charles Quanwei Cao, Gurdip S. Bakshi, and Zhiwu Chen. Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5), December 1997.
- Peter Carr and Dilip B. Madan. Option valuation using fast fourier transform. *The Journal of Computational Finance*, 2(4):61–73, 1999.
- Rama Cont and José da Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1):45–60, 2002.
- Toby Daglish, John Hull, and Wulin Suo. Volatility surfaces: theory, rules of thumb, and empirical evidence. *Quantitative Finance*, 7(5):507–524, 2007.
- Emmanuel Derman. Regimes of volatility. 1999. Quantitative Strategies Research Notes, Goldman Sachs, New York, NY. Also in Risk Magazine, April, 1999.
- Bruno Dupire. Pricing with a smile. *Risk*, (7):18–20, 1994.
- Valdo Durrleman. Convergence of at-the-money implied volatilities to the spot volatility. *Journal Of Applied Probability*, 45(2):542–550, 2008.
- Thomas W. Epps. Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, 74(366):291–298, Jun 1979.
- Mark B. Garman and Michael J. Klass. On the estimation of security price volatilities from historical data. *The Journal of Business*, 53(1):67–78, January 1980.
- Jim Gatheral. *The Volatility Surface : A Practitioner’s Guide*. Wiley Finance, 2006.
- Takaki Hayashi and Nakahiro Yoshida. On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli*, 11(2):359–379, 2005.
- Steven L Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–43, 1993.
- Roger W. Lee. The moment formula for implied volatility at extreme strikes. *Mathematical Finance*, 14(3):469–480, 2004.
- Roger W. Lee. Implied volatility: Statics, dynamics, and probabilistic interpretation. *Recent Advances in Applied Probability*, 2005.
- Alan L. Lewis. *Option Valuation under Stochastic Volatility*. Finance Press, 2000.
- M.I.A. Lourakis. levmar: Levenberg-marquardt nonlinear least squares algorithms in C/C++. [web page] <http://www.ics.forth.gr/~lourakis/levmar/>, Jul. 2004. [Accessed on 31 Jan. 2005.].
- Maria Elvira Mancino and Paul Malliavin. Fourier series method for measurement of multivariate volatilities. *Finance and Stochastics*, 6(1):49–61, 2002.
- Per Aslak Mykland and Lan Zhang. Inference for volatility-type objects and implications for hedging. *Statistics and Its Interface*, 1:255–278, 2008.
- Michael C. Münnix, Rudi Schäfer, and Thomas Guhr. Impact of the tick-size on financial returns and correlations. *Physica A*, 389(21):4828–4843, 2010.

Riccardo Rebonato. *Volatility and Correlation, The Perfect Hedger and The Fox*, chapter Stochastic-Volatility Processes. John Wiley & Sons Ltd, 2. ed. edition, 2004.

Paul Wilmott. *Paul Wilmott On Quantitative Finance*, chapter Stochastic Volatility. John Wiley & Sons Ltd, 2. ed. edition, 2006.