MAGYAR NEMZETI BANK

MNB

Occasional Papers

39.

2005

ÁRON GEREBEN–KLÁRA PINTÉR

Implied volatility of foreign exchange options: is it worth tracking?

Áron Gereben–Klára Pintér

Implied volatility of foreign exchange options: is it worth tracking?

May 2005

Implied volatility of foreign exchange options: is it worth tracking?*
(Devizaopciókból számolt implikált volatilitás: érdemes-e vizsgálni?)

Written by: Áron Gereben–Klára Pintér
(Financial Analysis)

Budapest, May 2005

Published by the Magyar Nemzeti Bank
Publisher in charge: Missura Gábor
1850 Budapest, Szabadság tér 8–9.

www.mnb.hu

ISSN 1585-5678 (on-line)

* We wish to thank András Fülöp, György Gyomai, Zsolt Kondrát and Balázs Vonnák for their useful comments on the earlier versions of this paper. Any remaining errors or mistakes are our own.

Contents

Abstract

Market analysts and central banks often use the implied volatility of FX options as an indicator of expected exchange rate uncertainty. The aim of our study is to investigate the limits of this statistic. We present some key factors that may deviate the value of implied volatility from the exchange rate variability expected by the market. These biasing factors are linked to the simplifying assumptions of the Black-Scholes option pricing model. Our empirical results show that forint/euro implied volatilities carry useful information about future exchange rate uncertainty when the forecast horizon is shorter than one month. However, implied volatility provides a biased estimate, and does not encompass the information included in other (GARCH, ARMA) predictors of volatility calculated from historical exchange rate data. These results are in line with the findings of similar analyses of other currency pairs.

Keywords: option, volatility, exchange rate.
JEL: G13.

1. Introduction

Black-Scholes implied volatility can be calculated by computing the volatility value that equates the theoretical option price given by the Black-Scholes formula with the observed market price of the option. Analysts of financial markets often consider the implied volatility calculated with this method as a good estimate of the expected price uncertainty of the underlying product. Private market analysts, central banks and international institutions often use implied volatility calculated from currency and stock index options to characterise market uncertainty. Various publications of the Magyar Nemzeti Bank also refer to the implied volatility of forint/euro options.[1]

The aim of this study is twofold. The first aim is to examine whether it is worthwhile monitoring this indicator, i.e. whether it is possible to infer the market's expectations and the future exchange rate uncertainty from the path of the implied volatility. The second is to highlight some of the "caveats" that come with implied volatility and describe the factors analysts of implied volatility may want to pay attention to when interpreting this indicator and when drawing conclusions based on its developments.

We touch on two closely interrelated subjects. Firstly, we examine how the differences between the nature of the "real" financial markets and the simplifying assumptions of the Black-Scholes model influence the value of implied volatility. Secondly, we perform an empirical analysis to decide how forint/euro implied volatility performs in practice as the predictor of the future uncertainty of the forint/euro exchange rate.

The paper is structured as follows. In Chapter 2, we discuss the biasing factors that stem from the simplifying assumptions of the Black-Scholes model, which may result in differences between the Black-Scholes implied volatility and the market's actual expectation of future volatility. In Chapter 3 we review the empirical literature on the forecasting ability of implied volatility. In Chapter 4 we present the result of our own econometric analysis on the forecasting ability and information content of the implied volatility calculated from forint/euro foreign exchange options. Chapter 5 summarises our conclusions on how it is worthwhile using and interpreting implied volatility as an indicator of future exchange rate uncertainty.

[1] From time to time, both the *Quarterly Report on Inflation* and the *Report on Financial Stability* use the forint/euro implied volatility as an indicator of market uncertainty (see, for example, MNB [2004], MNB [2005]). The various data sources of forint/euro implied volatility, together with a broader description of the Hungarian currency option markets, are presented by Csávás and Gereben (2005).

2. Theoretical reasons behind the difference between implied volatility and expected uncertainty

As mentioned in the introduction, both financial market participants and central banks use the Black-Scholes implied volatility as a measure of future uncertainty of asset prices. In such analyses, changes in implied volatility are, in most cases, interpreted as a shift in expectations about the uncertainty of the given asset's price. Meanwhile, it is often forgotten that implied volatility, due to the simplifying assumptions of the underlying Black-Scholes pricing model, often provides a biased picture, and its value can be influenced by other factors in addition to the uncertainty related to the price of the options' underlying asset.

If the relationship between implied volatility and market expectations is carefully examined, it becomes evident that implied volatility, even theoretically, does not provide an unbiased estimate of the market's expectations regarding uncertainty, except under some very strict conditions that are not likely to occur in real-world financial markets. Therefore, prior to the empirical analysis of implied volatility's forecasting ability, it is worth elaborating on the factors which may bias the value of implied volatility.

2.1. THE RELATIONSHIP BETWEEN IMPLIED VOLATILITY AND THE SUBJECTIVE EXPECTATIONS OF THE MARKET

Why do we monitor implied volatility? Let us presume that market participants have (subjective) expectations regarding the probabilities of possible future developments in the exchange rate. It may be useful to learn about this probability distribution for several reasons. Knowing what the market as a whole expects about prices is important in itself for market participants and monetary policy-makers to evaluate the possible consequences of trading and policy decisions. Furthermore, if market participants are presumed to be rational, then the market's expectation is likely to be a good predictor of future price developments. Implied volatility is usually considered as a simple estimate of the dispersion of the market's subjective probability distribution.

The market's subjective expectations are reflected in the prices of options. If options are traded on the market and their market prices are observable, then we can use this price data to calculate the implied volatility by "reverting" the Black-Scholes formula and computing the volatility value that equates the Black-Scholes price with the price observed on the market. Chart 1 shows the relationship between the market's subjective expectation and the implied volatility calculated as described above.

The starting point for mapping the expectations is the observable data, namely the listed option price. The theory of asset pricing suggests that there exists a direct relationship between the option price and the subjective probability distribution of the market expectation. According to the theory, the market's subjective expectations about the future and the market participants' risk preferences can be combined to create a so-called "risk-neutral" probability distribution. This distribution – which is sometimes referred to as state-price density – plays a central role in option pricing. Under the theory's usual assumptions about the market structure, the knowledge of the risk-neutral probability distribution allows the precise pricing of all derivative products of a given underlying asset, including options.[2]

If the above logic worked in the other way as well, one could calculate the risk-neutral distribution from option prices, and from the risk-neutral distribution the market's subjective expectations could also be filtered out. Unfortunately, however, the reverse procedure usually cannot be done in practice. The reason is that deriving the risk-neutral distribution from option prices requires simultaneous and reliable price data for a large number of different options on the same

[2] The relationships between option prices, risk-neutral probability distribution, preferences and subjective expectations are discussed in detail by Aït-Sahalia and Lo (2000) and Jackwerth (2000).

Chart 1

The relationship between the listed option prices, the market's subjective expectations and implied volatility

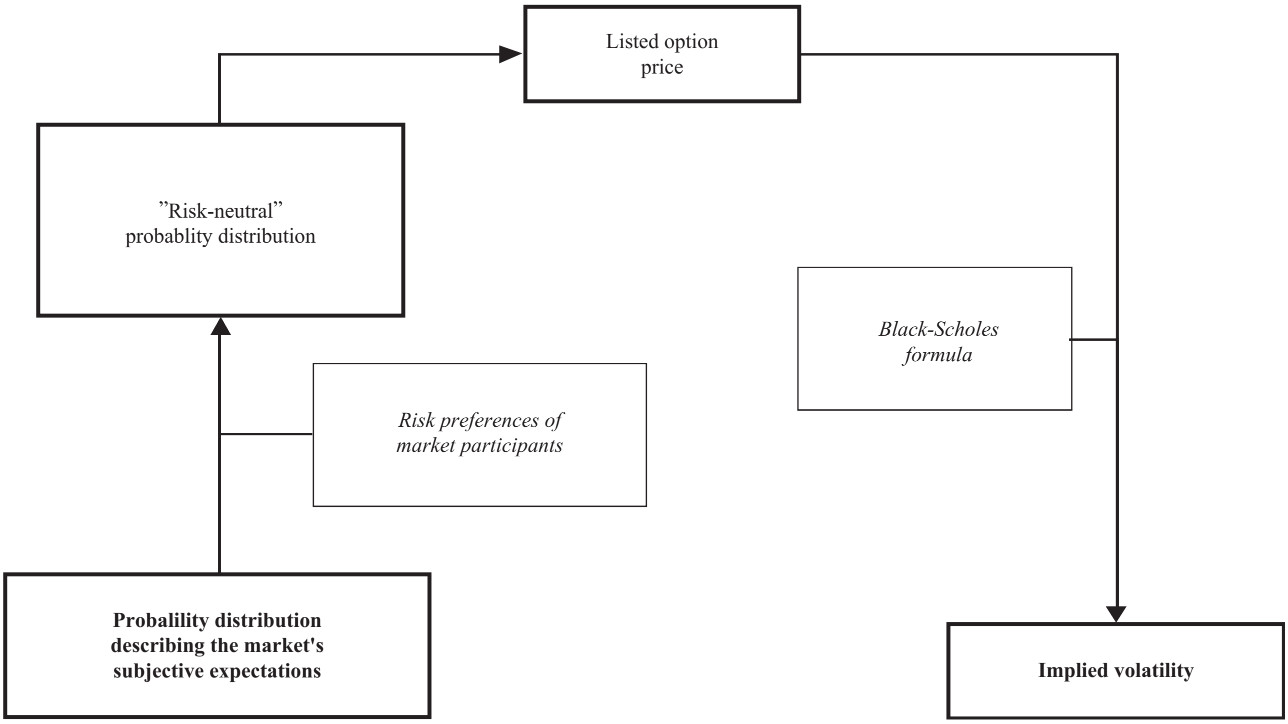

underlying asset. In practice such data are not available and the methods for approximating the missing data work only for the most liquid option markets. Even for these markets, taking the risk preferences into account creates an additional problem, as they are not observable empirically.

Analysts of financial markets can cut through the Gordian knot by using the Black-Scholes implied volatility, which allows us to obtain a quick estimate of the distribution of subjective expectations and its dispersion from relatively little data. However, in exchange for the simplicity, we must bear in mind that if we use the Black-Scholes formula, we make several strong implicit assumptions about the market environment. Although these assumptions simplify the calculations and reduce the data requirements, they often do not hold in reality, thus the estimates obtained are biased to a certain extent.

The assumptions of the Black-Scholes formula are as follows:
- the logarithmic changes in the underlying asset's price follow a normal distribution;
- the risk-free interest rate is constant;
- the volatility of the underlying asset's price is constant;
- there are no transaction costs;
- there are no arbitrage opportunities; and
- the trading of the underlying asset is continuous.

Since in reality these assumptions never hold perfectly, we cannot expect the Black-Scholes implied volatility to be identical with the actually expected uncertainty.

In addition to the deviations caused by the above assumptions, risk preferences also bring bias into the calculation of implied volatility. Under the special assumptions of the Black-Scholes model the risk-neutral and subjective probability distributions differ only in their expected values, but their shape and dispersion is identical. As a result, risk preferences do not have to be taken into account when calculating the Black-Scholes implied volatility. However, this is not true for the general case. Outside the universe of the Black-Scholes model, the shapes of the risk-neutral distribution and the distribution describing subjective expectations are not necessarily identical, and consequently their dispersion, measured by the volatility, may also be different.

In the following subsections we examine one by one the potential reasons behind the bias of implied volatility, and look into the extent and the nature of the bias they may cause in practice.[3]

2.2. DEVIATIONS FROM THE NORMAL DISTRIBUTION

The Black-Scholes model assumes that the logarithmic changes of the underlying asset's price follow a normal distribution. In reality, however, observed price changes of financial products, including exchange rates, are different from the normal distribution. The empirical distribution is usually more peaked (leptokurtic) than the normal. This also means that extreme outcomes at the distribution's tails are more frequent. If the empirical distribution is non-normal, it is plausible to presume that the distribution describing market participants' expectations is also likely to differ from the normal.

A sign of this extra leptokurtosis can also be observed on most option markets. Market data show that implied volatilities of options at different strike prices but identical maturity are usually different, while they should be identical if the normality assumption of the Black-Scholes model were true. In reality, the implied volatility of options with strike prices near the spot price (at-the-money options) is usually lower than that of in-the-money and out-of-the-money options, which have strikes farther away from the spot price. This phenomenon is called the volatility smile curve. In addition, the smile curve is usually not symmetrical, but tilted towards one direction or the other. This indicates that the underlying distribution is not symmetrical either.

It cannot clearly be decided which point of the implied volatility smile represents best the true dispersion of the distribution describing expectations. Most often the implied volatility of at-the-money forward (ATMF) options is used in analysing foreign exchange markets. However, the reason is not theoretical – since the market of at-the-money options is usually the most liquid, these data are the most reliable and the easiest to observe.[4]

The ATMF implied volatility is usually close to the lowest point of the volatility smile curve. Assuming that the standard deviation of the market expectations can be obtained as some mean value of the observed implied volatilities, the ATMF implied volatility, being at the lower section of the volatility smile curve, is biased downwards (Backus, Foresi, Li and Wu [1997]).

2.3. STOCHASTIC INTEREST RATE

The Black-Scholes model assumes that the risk-free interest rate is constant; in reality, it is not. If we calculate implied volatility from market prices, the resulting indicator will simultaneously reflect both the underlying asset's expected price volatility and the volatility of interest rates.

Merton's (1973) model releases the assumption of constant interest rate and examines how the price of the option can be calculated when interest rates are stochastic. According to his findings, the Black-Scholes formula remains valid with stochastic interest rates as well, with the difference that the volatility of the original formula should be replaced with the sum of the squares of the volatility of the underlying asset and the volatility of the risk-free asset. Formally:

$$\hat{\sigma}^2 = \left(\sigma^2 + \sigma_P^2 - 2\rho\,\sigma\,\sigma_P\right),$$

where $\hat{\sigma}$ is the volatility value brought into the BS formula, σ is the volatility of the underlying asset, σ_P^2 is the volatility of the risk-free asset, and ρ is the correlation coefficient between the prices of these two instruments.

How does this affect the value of implied volatility? The Merton model suggests that the Black-Scholes implied volatility is in fact the function of three factors: the (expected) volatility of the underlying asset, the volatility of the risk-free asset and the price correlation of the two assets. Therefore, when interest rates are stochastic, the Black-Scholes implied

[3] Bates (1996) provides a similar, but more technical analysis of the factors potentially biasing the implied volatility. Poon and Granger (2005) also provide a comprehensive summary of the possible reasons for the bias.

[4] Bates (1996) describes some alternative measures of implied volatility, which can be obtained by averaging various points of the volatility smile curve.

volatility cannot be considered an unbiased predictor of the underlying asset's volatility, as it also reflects the interest rate volatility to some extent. The size and direction of the bias depend on the magnitude of interest rate volatility and the direction of the correlation of the two asset prices.

In practice, the volatility of interest rates, especially for options with short maturity, is not significant, as the volatility of the risk-free assets is negligible compared to the volatility of the underlying assets, like equity and foreign exchange. The bias is therefore usually rather small. However, there are periods when interest rate volatility increases. During these periods it is necessary to take the effect of interest rate volatility into account when interpreting implied volatility. This is particularly true for options with longer maturities, which are more sensitive to interest rate variability. (Appendix 1 provides an analysis of the interest rate bias in the implied volatilities of long-maturity forint/euro foreign exchange options.)

2.4. STOCHASTIC VOLATILITY

The Black-Scholes model assumes that the volatility of the underlying asset is constant: if the assumptions of the original model held in the real world, the value of implied volatility should always be the same. Consequently, the fact that implied volatility is not constant in time seems to question the relevance of the indicator itself.

Nevertheless, using the Black-Scholes formula when volatility is variable is not unjustified. Merton (1973) showed that if volatility is a known, deterministic function of time, the Black-Scholes model remains valid, provided that the volatility is replaced by the average volatility of the base product for the duration of the option.

If volatility follows a stochastic process, however, the Black-Scholes formula does not hold anymore. Hull and White (1987) provided a pricing formula for a particular case of stochastic volatility, when the volatility and the price of the underlying asset do not correlate. Under the assumptions of the Hull-White model, the at-the-money implied volatility calculated using the Black-Scholes model is slightly biased downwards, while in the case of in-the-money and out-of-the money options the bias goes upwards (Bates [1996], Poon and Granger [2005]). Fleming (1998) shows that for short-maturity at-the-money options the Black-Scholes model is nearly linear with regard to volatility; therefore implied volatility calculated from such options is an almost unbiased predictor of the volatility expected until the option's maturity, if the conditions of The Hull-White models hold.

Since the appearance of the Hull-White model, several other option pricing techniques have been developed under the assumption of stochastic volatility (see, for example, Heston [1992]). From our perspective, is worth highlighting the GARCH option pricing models (Duan 1995), which assume that the underlying asset's volatility follows a generalised autoregressive process of conditional heteroskedasticity (GARCH). The GARCH option models are particularly interesting due to their good empirical performance. As will be discussed in later chapters of this study, most empirical analyses find that implied volatility calculated from the at-the-money options usually overestimates the realised volatility. This contradicts the conclusions of the traditional stochastic volatility models, but at the same time it is in line with the predictions of the GARCH option pricing models.

2.5. THE DIFFERENCE BETWEEN THE RISK-NEUTRAL DISTRIBUTION AND THE DISTRIBUTION OF THE MARKET'S EXPECTATIONS

General equilibrium models of option pricing (e.g. Cox, Ingersoll and Ross [1985]) suggest that the probability distribution obtainable directly from the option prices is a so-called risk-neutral distribution. The risk-neutral distribution can be considered as a mixture of the distribution of the market's expectations and the market participants' risk preferences. The Black-Scholes model is a special case in a sense that under its assumptions the dispersions of the risk-neutral distribution and the distribution of subjective expectations are identical.

In reality, where the assumptions of the Black-Scholes model do not hold, the dispersions of the risk-neutral distribution and of the distribution of subjective expectations are not necessarily identical. Implied volatility may thus be biased, as it may reflect risk preferences as well as the market's expectations.

As to the magnitude of this bias, earlier empirical studies attempting to estimate the risk-neutral distribution with the help of option prices for various instruments (Melick and Thomas [1997], Bahra [1997], Malz [1997], Gereben [2002], etc.) assumed that it was negligible. Most of these works refer to Rubinstein's study (1994), which asserts that the risk-neutral and the subjective distributions differ mainly in their expected value, and the difference between their dispersion and higher-order moments is insignificant.

However, according to more recent studies, the difference may be important for practical applications as well. Using a more general model than the Black-Scholes case, Bliss and Panigirtzoglou (2002) find that for the S&P 100 equity index options the dispersions of risk-neutral and subjective distributions differ only slightly. However, this difference is sufficient to affect the forecasting ability of options-based forward-looking indicators. Breuer (2003) draws a similar conclusion for Hong Kong dollar/US dollar foreign exchange options.

2.6. MARKET IMPERFECTIONS AND INVESTORS' PREFERENCES

The Black-Scholes model involves strict assumptions about the market's structure. If these do not hold – e.g. transaction costs exist and arbitrage opportunities cannot completely be exploited – option prices observed in the market do not reflect the expectations perfectly. If this is the case, demand and supply conditions of the option market can also influence the prices, beyond the market expectations, at least to a certain extent. This causes an additional bias factor. This bias is less important for highly liquid markets, where bid-ask spreads are narrow. For the forint/euro foreign exchange options, however, which cannot be considered too liquid,[5] the bias caused by market imperfections is likely to be reflected, even in the implied volatility of the most liquid at-the-money options.

Prices may also become biased if the risk-taking ability of individual market participants is limited and different market participants have different risk preferences. For example, on markets for equity index options it is often observed that the price – and the implied volatility – of put options with low strike prices is much higher than the theoretically fair price suggested by option pricing models. Numerous studies, e.g. Franke, Stapleton and Subrahmanyam (1998) or Bates (2001), explain this phenomenon by market segmentation. Large investment funds like buying such options to insure themselves against large falls in equity prices. At the same time, there are very few market participants that are able to sell large quantities of out-of-money put options, since for most financial institutions the internal or external risk regulations do not allow the excessive accumulation of such exposures. Due to the different preferences and the risk-taking limits, the price of these options is higher than what would be justified by the expected probability distribution of equity prices, while regulations and differing risk preferences prevent market participants exploiting the arbitrage opportunity.

Similar phenomena arise in foreign exchange option markets, too. According to the triennial global foreign exchange market survey of the BIS, option market-makers sell more options than they buy (BIS [2005]). This suggests that the market of foreign exchange options is also segmented: the market makers' clients tend to buy options, while the supply is provided by market-makers themselves. As they are likely to have limited risk-taking abilities, and hedging open option positions with other instruments is costly, the extra risk taking is likely to be reflected in the option prices.

Other forms of supply constraints, arising from the segmented markets, may further bias the value of implied volatility in foreign exchange option markets. For example, hedging strategies connected to some exotic options often require very large transactions on the market for regular, plain vanilla currency options. In these cases it is often empirically observable that sudden, temporary shifts in demand or supply, caused by these hedges, affect the value of implied volatilities calculated from regular at-the-money options (Malz [1995], Csávás and Gereben [2005]).

2.7. THE EFFECT OF AN EXCHANGE RATE BAND

In the case of the forint-euro implied volatilities, the role of the exchange rate band should be also mentioned.[6] The fact that the exchange rate of the forint does not float freely, but can only fluctuate between the limits of the declared exchange rate band, affects the values of implied volatility, and in certain cases may cause significant bias.

[5] While the bid-ask spread of euro-dollar options is around 10-40 basis points, this value for the forint-euro options fluctuates between 100 and 300 basis points (Csávás and Gereben [2005]).

[6] Naszódi (2004) examines the link between the Hungarian forint's exchange rate band and option prices from a different aspect.

Due to the existence of the exchange rate band, the probability distribution of exchange rate changes differs from the normal. If the band is perfectly credible, the exchange rate cannot have values out of the band. Even in the case of a not fully credible exchange rate regime, the probability of out-of-the-band exchange rate outcomes is usually smaller than that of within-the-band values. As a result, if the exchange rate approaches the edge of the band, the future probability distribution will be tilted compared to the normal distribution, which is symmetrical.

In an exchange rate band, the expected volatility of the exchange rate depends on the position of the exchange rate within the band. In Krugman's (1991) exchange rate target zone model, which assumes a completely credible exchange rate band, it can be shown that the closer the exchange rate to the edge of the band, the smaller the dispersion of the distribution of the expected future exchange rate (Svensson [1991]).

This result can also be extended to some models in which the band is not completely credible, i.e. there is a positive probability of giving up the band. In those target zone models where the probability of giving up the band is determined independently of the exchange rate developments (models with exogenous realignment risk, e.g. Bertola and Svensson [1993]), the expected volatility of the exchange rate declines when the exchange rate approaches the edges of the band, similarly to the case of the credible target zone.

However, in models where the target zone's maintenance depends on the same fundamentals as the value of the exchange rate (so-called endogenous realignment models, e.g. Bertola and Caballero [1992]), the situation is the opposite: when approaching the edge of the band, the expected volatility of the exchange rate increases.

All in all, the existence of the band affects the value of the expected exchange rate volatility. Even if all other factors remain unchanged, the distance from the edge of the band has an impact on the expected exchange rate uncertainty; and the magnitude and direction of this effect depends on the credibility of the band.

This attribute can also be used for verifying the credibility of the exchange rate band. Campa and Chang (1998), for example, test the credibility of the Exchange Rate Mechanism (ERM) of the European Monetary System based on the implied volatilities calculated from option prices on the member countries' currencies. They found that during times when exchange rates were close to the band limits, implied volatilities usually rose. This is consistent with the endogenous realignment models and indicates that the ERM had lost its credibility in the market's eye well before the emergence of exchange rate crises.

A relevant question is how well does Black-Scholes implied volatility measure the expected exchange rate volatility in target zone systems, where the exchange rate's distribution is clearly not normal. In the following we illustrate with a numerical example the magnitude of the potential bias resulting from the existence of an exchange rate band.

De Jong, Drost and Werker (2001) propose an exchange rate model with a credible target zone, which model is similar to the Krugman model, but easier to handle in terms of calculation. They also derive an option pricing formula consistent with the model. Using this framework, we examined the magnitude and direction of the bias which may be caused by using the Black-Scholes implied volatility in a credible band.

The results of this experiment are shown in Table 1. In our example, the parameters of the model were set in a way such that the edges of the band correspond to the Hungarian framework. Implied volatility around the middle of the band is set at an annual level of 5 per cent.[7]

Given these assumptions, the table shows the developments in the one-year ahead volatility of the expected exchange rate approaching the edge of the band, and its estimate indicated by the Black-Scholes implied volatility, calculated backwards from the price of one-year ahead options, priced with the framework's own option pricing model. The last row of the table shows the percentage difference between the two volatilities.

[7] We would like to call our readers' attention to the fact that our calculations are illustrative only. We have not tested whether the assumptions of the model hold in practice, nor are the parameters based on empirical estimates. Our aim here is to show, within the framework of a given model, how large a bias the existence of the band can bring into the value of implied volatility. We do not intend here to provide even a rough representation of the Hungarian exchange rate regime.

Table 1

The performance of implied volatility in a target zone

Spot exchange rate	282.36 (mid-point of the band)	250	245	242	241	240.5
Model-based expected volatility (%)	5.00	3.45	2.54	1.66	1.23	0.96
Black-Scholes implied volatility (%)	5.29	3.59	2.58	1.56	1.09	0.84
Bias (%)	5.68	3.95	1.53	-5.99	-11.92	-16.19

It can be seen from the table that in this model, which assumes a perfectly credible exchange rate band, both realised and implied volatilities decline fast as the spot exchange rate approaches the edge of the band (which is at HUF 240.01). It can also be observed that the Black-Scholes volatility is biased slightly upwards in the middle of the band. Once approaching the edge of the band, however, the direction of the bias changes and the Black-Scholes implied volatility is somewhat lower than the model-based volatility. The magnitude of this bias is relatively low; even very near the edge of the band it is comparable to the magnitude of the bid-ask spread normally observed on the forint/euro options market.

2.8. CONCLUSIONS

The above examples illustrate that several factors may cause bias in the implied volatility, relative to future expected volatility. The simplifying assumptions of the Black-Scholes formula provide only an approximate description of real-world financial markets, therefore the implied volatility cannot be expected to be a perfect predictor.

However, it would be early to conclude that the Black-Scholes implied volatility does not carry useful information. Firstly, biases may very often be non-significant for practical applications. Secondly, if bias factors are relatively stable over time, changes in implied volatility will indicate changes in expected exchange rate uncertainty, even if the level is biased. Consequently, the dynamics of implied volatility may carry important information, even in the presence of forecasting bias.

In the next part of the study we examine the forecasting performance of implied volatility from an empirical aspect. First, previous results with regard to various equity and foreign exchange markets will be reviewed. Then, in the light of these results, the forecasting ability of the forint/euro implied volatility will be examined.

3. The information content of implied volatility – an overview of the literature

Numerous studies discuss the relationship between the implied volatility derived from option prices and the realised volatility of the underlying assets' prices. Research in this area usually covers two topics. Some studies simply examine whether implied volatility can forecast future realised volatility. Other studies also raise the question whether implied volatility contains all the relevant information, in addition to providing a forecast of future volatility. In other words, these studies also test whether there are other variables beyond the implied volatility, e.g. past price data, that could improve the forecast.

The above-described biasing factors may affect the different markets to different extents, therefore it would not be surprising if the information included in implied volatility differed markedly in different asset markets. However, as we will see, empirical results prove to be roughly similar, irrespective of the financial instruments examined.

As the first analyses of implied volatility's information content has been carried out on stock index options market data, and a considerable part of the literature continues to focus on the equity markets, we will begin our (non-comprehensive) overview with the literature on equity options. Then the results for foreign-exchange option markets, more important for our purposes, will be presented.

3.1. EQUITY OPTIONS

As mentioned above, initial research on the subject focussed on the equity markets and analysed mainly the volatility derived from the most liquid stock index options. Most of these studies concluded that implied volatility contains limited information on future volatility.

Analysing individual stocks, Lamoureux and Lastrapes (1993) concluded that time series models based on historical data provide a better forecast than implied volatility. Using data on the S&P 100 equity index options, Canina and Figlewski (1993) found that in the period between 1983 and 1989 the implied volatility gave practically no information on future realised volatility. Their results held irrespective of the options' strike price and maturity.

Later studies, however, show that option prices – at least on liquid markets – have significant, non-negligible information on future volatility. For the S&P 100 index, Christensen and Prabhala (1998) noted that the conclusions drawn by Canina and Figlewski partly result from the inappropriateness of the methodology they used. In their study, they point out that in time series of implied volatility, data on consecutive days reflect forecasts for overlapping periods. In the case of exchange traded options these periods also differ in the forecast horizon. Usual estimation methods and test statistics are not applicable in such cases, and disregarding these features of the data leads to false conclusions for an empirical analysis.

The authors also carried out the analysis with a more suitable methodology. The conclusion was that in the period following the 1987 crisis, implied volatility was an unbiased forecasting indicator of future realised volatility. They also found that it was also providing an efficient forecast in the sense that past volatilities – calculated in this case from past realised yields with a 60-day rolling window – do not contain additional information above implied volatility.

Subsequently, a number of other studies confirmed the forecasting power of implied volatility; for examples see Fleming (1998), Blair, Poon and Taylor (2001), Poteshman (2000), Koopman, Jungbacker, Hol (2004) or Corrado and Miller (2004). These works generally found that although implied volatility forecasted future volatility with an upward bias, it contained more information than historic data. The results of Fleming (1999) suggest that even the bias is insignificant, especially if one cleans the data from the effects of transaction costs. These results are further confirmed by Christensen and Strunk Hansen (2002), who extend the analysis of Christensen and Prabhala (1998) to a longer time period, and conclude that implied volatility is an unbiased indicator of the future realized volatility.

Conclusions differ whether implied volatility includes all the information that can be obtained from past data. Ederington and Guan (2002), and Martens and Zein (2002) came to the conclusion that ARCH models based on past volatilities contained more information on future volatility than implied volatility. In contrast, Blair, Poon and Taylor (2001), Christensen and Prabhala (1998), Christensen and Strunk Hansen (2002) and Fleming (1998) concluded that time series based models of volatility had no additional information content relative to implied volatility.

In the case of small and less liquid stock markets, it is often concluded that the forecasting power of implied volatility is negligible, or at least it does not provide additional information relative to the simple time series models (see Gonzalez Perez [2004], on Spanish equity index options). For Swedish equity index options, Frennberg and Hansson (1996) found that implied volatility was not able to outperform the random walk in forecasting future volatility.

3.2. CURRENCY OPTIONS

As to the currency markets, most studies find that implied volatility is a good forecaster of future volatility.

Jorion (1995) analyses the forecasting power of implied volatility derived from the exchange-traded options on the deutschmark/dollar, the yen/dollar and the Swiss franc/dollar currency pairs in the period between 1985 and 1992. He found that implied volatility could explain a significant part of variations in future realised volatility, both at one day and one month forecast horizons. He pointed out that, although backward-looking indicators, such as GARCH(1,1) and MA(20) forecasts, also have explanatory power, they do not provide more information than implied volatility. Taylor and Xu (1995) draw similar conclusions from an analysis of deutschmark/dollar options between 1985 and 1991. The ARCH model they used as a benchmark did not provide additional information relative to the implied volatility.

Martens and Zein (2003) analysed the implied volatility of yen/dollar options between 1996 and 2000. They, too, found that in comparison to the GARCH(1,1) model, implied volatility was an efficient predictor over a horizon of 1 to 20 days. The projections could not be improved by time-series based models of daily data. By using intraday data and taking volatility indicators' long-memory properties into account, the forecast could, however, be improved on a few horizons. Nevertheless, in most cases implied volatility also had the information provided by the intraday indicators.

Neely (2002) analyses the dollar/deutschmark, the dollar/yen, the Swiss franc/dollar and the dollar/pound between 1987 and 1998. He uses ARIMA, long-memory ARIMA and GARCH(1,1) models to serve as bases of comparison for the information content of implied volatility. In the case of the yen/dollar exchange rate, his results support the view that implied volatility is an efficient forecast of future realised volatility. None of the alternative models provided extra information relative to the implied volatility, irrespective of whether daily or intraday data were used for measuring realised volatility. For intraday data, implied volatility has proven to be an efficient predictor for the Swiss franc/dollar exchange rate volatility, too.

Christoffersen and Mazzotta (2004) analysed OTC option data on the dollar/euro, yen/euro, pound/euro and yen/dollar currency pairs between 1998 and 2002. (Prior to 1999, the euro was replaced by the German mark). One of their most significant conclusions is that when using OTC data instead of exchange-traded options, implied volatility provides better forecasts of future realised volatility. This is due to the fact that OTC option data poses fewer methodological problems than exchange-traded ones.

Their results confirm that implied volatility usually provides a better forecast for future volatility than the simultaneously analysed time series models, such as historic volatility, exponentially weighted historic volatility and the GARCH(1,1) model. This has proven to be true over a one-month as well as a three-month horizon. However, implied volatility seemed to be an inefficient predictor: time-series models contained some additional predictive information. The method of measuring realised volatility – i.e. whether it was calculated from daily or intraday data – did not affect the conclusions.

Taylor and Xu (1997) also compare implied volatility with the forecasts that rely on intraday exchange rates, besides the forecasts based on historical exchange rate data at daily frequency. Their results suggest that if the information set is extended to historical intraday exchange rate fluctuations, implied volatility is no longer efficient as a predictor. Intraday exchange rate data provides additional information on future exchange rate volatility, therefore their addition to the models can improve the forecasts. Neely (2002) came to the same conclusion when he analysed the dollar/deutschmark and pound/dollar exchange rates.

Pong, Shackleton, Taylor and Xu (2004) analysed implied volatilities derived from OTC options on the dollar/deutschmark, pound/dollar and dollar/yen rates. They compared the implied volatilities with the forecasts based on the daily frequency GARCH(1,1) model, and the ARMA and ARFIMA models based on intraday data. They found that in the case of 1-day-ahead and 1-week-ahead forecasts, the intraday indicators performed better. For the 1-month and 3-month forecast horizons, however, the implied volatilities were the best predictors. Furthermore, they found that with every forecast horizon, implied volatilities contained some unique information that could not be extracted merely from past data.

As to whether implied volatility is an unbiased predictor of the future realised volatility, results are less favourable. The majority of analyses find evidence for a bias: in estimated equations explaining future volatility, the regression coefficients of implied volatility are usually lower than 1, which suggests that implied volatility contains an upward bias. This result was found, for example, by Jorion (1996) for all three currency pairs he examined, by Neely (2002) for the dollar/deutschmark, dollar/Swiss franc and pound/dollar rates; and by Christoffersen and Mazzotta (2004) for some currencies they analysed.

Neely (2002), however, found the implied volatility of the dollar/yen to be an unbiased predictor. Furthermore, Christoffersen and Mazzotta (2004) find that implied volatilities of post-1999 euro cross rates are also unbiased predictors of realised volatility calculated from daily data.

In contrast to stock markets, implied volatilities calculated from foreign exchange options provide useful information on future volatility on the smaller markets, too. Aguilar (1999) compared the implied volatilities calculated from Swedish krone/deutschmark, Swedish krone/US dollar and Australian dollar/US dollar OTC options with forecasts based on the GARCH(1,1) and EGARCH(1,1) models. His results revealed that between 1995 and 1998, the implied volatility of Swedish krone/deutschmark and Australian dollar/US dollar options forecasted future volatility without a bias on most horizons. In the case of the Australian dollar/US dollar rate, implied volatility proved to be the best predictor, while the EGARCH(1,1) model provided the most accurate forecast for the volatility of the Swedish krone/deutschmark exchange rate, although the implied volatility also provided additional information.

Results are favourable for the emerging markets, too. Using OTC options data, Cincibuch and Bouc (2001) reveal that in the case of the Czech koruna, implied volatility predicts exchange rate volatility without a bias on a 5-day horizon. In comparison to a simple historic average it is also an efficient predictor. Canesso de Andrade and Tabak (2001), as well as Tabak, Chang and Canesso de Andrade (2002) analysed the implied volatility calculated from exchange-traded options of the Brazilian real/US dollar exchange rate. Their results suggest that irrespective of the forecast horizon, implied volatility provides good, although biased, forecasts on future volatility. A comparison of implied volatility and forecasts based on the GARCH(1,1) and MA(20) models revealed that in an information set that included daily data, option-based implied volatilities contained all the relevant information. Once the information set was complemented with the daily closing and opening prices, implied volatility became an inefficient predictor; nevertheless the additional information provided by the time-series models has proven to be small compared to the information built in option prices.

3.3. IMPLIED VOLATILITY AS AN INDICATOR OF MAJOR MARKET TURBULENCES

The key question of the literature summed up so far is whether implied volatility predicts future volatility. Future volatility is, however, frequently not the most pivotal issue for practical applications. For the risk management of financial institutions, or for central bank policy-making, forecasting normal-sized fluctuations in future volatility, which are part of the usual market conditions, is often a secondary issue. The main question is whether we can predict drastic and major shocks to market prices.

Surprisingly, very little is written in the literature about the forecasting power of implied volatility in this regard. As far as we know, a study by Malz (2000) is the single example of an explicit attempt at answering the question whether the Black-Scholes implied volatility is suitable for predicting major market turbulences.

In his study, Malz analysed options-implied volatilities of 11 assets, including stocks, bonds, crude oil, gold and the exchange rates of diverse currencies. Using Granger causality tests and contingency tables, he examined whether changes in implied volatilities and historic volatility indicators can predict major shocks in the market of the specific asset.

His results suggest that implied volatility performs well in this respect: in numerous cases it helped predicting major shocks, even when historic volatility indicators have not changed. As a result, the author concluded that the implied volatility was a useful forward-looking indicator for risk management.

3.4. CONCLUSIONS

On the basis of the extensive international literature on the forecasting performance of implied volatility calculated from equity and foreign exchange options, the following conclusions can be drawn:

- In the majority of cases, implied volatility provides useful information on the expected volatility; its changes usually suggest changes in future volatility.
- Nevertheless, implied volatility usually predicts future volatility with a bias: its level tends to be higher than the actual future volatility.
- Implied volatility has not always proven to be an efficient predictor. Other indicators of past changes in the exchange rate – especially intraday data for exchange rate volatility – often provide additional information, which helps to improve the forecast.
- In the case of foreign exchange markets these promising results could be extended to smaller, less liquid currency markets as well.
- For the purposes of central bank policy-making it needs to be highlighted that implied volatility has proven to be a useful predictor of major shocks to asset prices.

4. The information content of the forint/euro implied volatility – an econometric analysis

The empirical section of our study discusses whether the conclusions drawn in the international literature are also applicable to the forint/euro options. In other words, we will examine whether the forint/euro implied volatility predicts future exchange rate volatility, and whether this prediction is unbiased and/or efficient. First we provide an introduction to our data set, then we discuss the methodology applied, and finally we discuss the results of the econometric analysis.

4.1. DATA

a. Measuring realised volatility. The very first question one faces in an empirical analysis of the predictive power of implied volatility is how to measure the subject matter of our forecasts. Exchange rate volatility is a feature of the actual data generation process, which cannot be directly observed.[8] Its adequate approximation is a key issue in any empirical analysis.

Several approaches have been developed to model and measure volatility. They are discussed in detail in the study of Poon and Granger (2003). Earlier studies focussed on data at daily frequency, and approximated the volatility with absolute or quadratic changes in asset price changes. Later, the availability of intraday data added new tools to the measurement of volatility. Andersen and Bollerslev (1998) point out that daily absolute or quadratic price changes are an unbiased, but noisy estimator of the latent volatility. Simultaneously, they demonstrate that volatility can be more accurately measured by the sum of quadratic intraday yields. If yields are calculated at sufficiently frequent intervals, the measurement error of volatility – i.e. the noise – can be reduced considerably. Volatility can practically be treated as an observable variable, if yields can be measured at a high enough frequency. Andersen, Bollerslev, Diebold and Labys (2001, 2003) provided theoretical verification for this conclusion in a no-arbitrage framework.

In our analysis, realised volatility is also approximated with intraday yields. Our time series of realised volatility are generated from Reuters D2000 quotes, which are available at a 2-minute frequency. We cover the period between 2 January 2002 and 26 May 2005.[9]

Theoretically, the more frequently one can measure the price, the better volatility can be approximated. However, as is highlighted by the studies of Alizadeh, Brandt and Diebold (2002) and Brandt and Diebold (2003), with increasing sampling frequency, market microstructure noise may increasingly influence price dynamics, and this effect may outweigh the benefit of high-frequency sampling. As a result, considerable attention has been given in the literature to find the optimal sampling frequency for intraday yields. The problem was analysed, for instance, in a study by Aït-Sahalia and Mykland (2003), as well as by Bandi and Russel (2003). In empirical analyses, 5 to 30-minute yields are the most widely used.

We performed the analysis using 30-minute and 60-minute intraday yields. The frequency of sampling did not affect our conclusions, therefore we decided to report only the results based on the 30-minute yields. Initially, we approximated daily realised volatilities by the daily quadratic sum of the 30-minute yields. For this purpose, the following formula was used:

$$RV_t = \sqrt{\sum_{h=1}^{H} r_{t,h}^2} \quad ,$$

[8] In this study we will not discuss the possibility of data generation processes where the concept of volatility cannot be interpreted.

[9] Transaction price data are also available through Reuters. However, this time series is only available for a shorter period, therefore we have decided to use quotes instead.

where RV_t is the estimated (realised) volatility on day *t*; and $r_{t,h}$ is the (logarithmic) yield realised in the *h*-th 30-minute period on day *t*. To ensure synchroneity with implied volatility data, volatility for any given day is calculated from the yields measured between 10.30 am on that day and 10.30 am on the following day.

By using this technique, however, the time series of realised volatility proved to be too volatile: outliers were frequent and their presence was independent of the frequency of yield measures within a day. The time series of intraday logarithmic exchange rates reveals that a possible underlying reason for these outliers may be the presence of jumps in the intraday prices (See the graphs in Appendix 2).

A study by Barndorff-Nielsen and Shephard (2004) offers an alternative solution for measuring daily variance. This is also consistent with stochastic volatility models and is relatively robust in cases when major jumps are seen in the observed exchange rate.[10] The recommended alternative indicator is the bipower variation. It is similar to realised volatility, but instead of using quadratic intraday yields, the absolute values of yields in consecutive intervals are multiplied together, before being added up. By using this technique, the volatility impact of a major price change in a specific interval decreases if it is followed by a smaller change in the next period. The realised "bipower variation" is calculated by the following formula:

$$\sigma_t = \sqrt{\sum_{h=1}^{H-1} |r_{t,h}| \cdot |r_{t,h+1}|},$$

where σ_t stands for volatility on day *t*; and $r_{t,h}$ is the (logarithmic) yield realised in the *h*–th 30-minute period on the day *t*; while $r_{t,h+1}$ is the (logarithmically measured) yield realised in the *h+1*-th 30-minute period on the day t. Due to its beneficial properties, in this study we approximated daily realised volatility with this measure.

Daily volatilities were then used to calculate daily, weekly (5-day), monthly (21-day) and 3-month (63-day) volatilities, and then annualised to match the quoting conventions of the implied volatility used. Thus the realised volatility time-series we used in the regressions were calculated using the following formula:

$$V_{t,H} = \sqrt{\sum_{i=1}^{H} \sigma^2_{t+i-1}} \cdot \sqrt{\frac{252}{H}},$$

where *H* is the number of days in the given time horizon.

Some outliers due to shocks related to the exchange rate regime have been removed. These shocks include the excessive exchange rate fluctuations seen during the speculative attack against the strong edge of the forint's exchange rate band in January 2003 (16 and 17 January 2003), the excessive volatility subsequent to the target zone realignment in June 2003 (4-10 June 2003), and large fluctuations related to the exchange rate shock in early December 2003 (1 and 2 December, 2003).

b. Implied volatility. The source of our implied volatility time series *(IV)* is the Royal Bank of Scotland's Reuters page (RBVN), where the bank lists prices for 1-week, 1, 2, 3, 6, 9 and 12-month forint/euro options. In the analysis, we used the average of the bid and ask volatilities of one-week *(IW^{1W})*, one- and 3-month *(IW^{1m}, IW^{3m})*, European-style at-the-money (ATM) options. The time series spans the period from 31 December 2002 to 25 May 2005.

We excluded options with maturity longer than 3 months, as the market of longer-term forint/euro options is rather illiquid. On certain maturities, data are available from May 2001. However, again, insufficient liquidity made us question the usefulness of the pre-2003 data, thus this was excluded from the final analysis.

[10] If jumps are not present, both the realised volatility and the alternative indicator discussed in the study converge to the integrated variance; thus both are consistent estimates of the latter.

Although option data is quoted directly in the form of implied volatility, this does not mean that it is independent of the option pricing model or measurement errors. According to the relevant market conventions, quotes are given as implied volatilities calculated using the Black-Scholes model.[11] Despite numerous potential biasing factors linked to theoretical reasons and market imperfections, we used these implied volatilities directly in their raw form.

As both market and central bank analyses usually refer to the Black-Scholes implied volatility directly, it is the relevant indicator to study, rather than a more complicated, more realistic but less widely used one. Furthermore, the use of more complex models assuming stochastic volatility requires the estimation of additional parameters, which may introduce further uncertainties and measurement errors. Moreover, the literature suggests that the model choice does not affect the results regarding the predictive power of implied volatility. In the case of foreign exchange options, Neely (2002) compares implied volatilities calculated by three option pricing models – Heston (1993), Barone-Adesi and Whaley (1987) and Black (1976) – and reveals that they produce highly similar descriptive statistics.

c. Alternative projections. The predictive power of implied volatility is compared to two alternative, time-series based forecasts.

Our first alternative forecast is the GARCH(1,1) volatility forecast, which is widely popular, both in the academic literature and in practical applications. We created forecasts using this model on daily, weekly (5 days), one-month (21 days) and three-month (63 days) horizons for the period between 31 December 2002 and 25 May 2005. The forecasts are ex ante out-of-sample forecasts, made at any given day with a model that was estimated using a sample covering the period from 2 January 2002 to the given day.

The second time-series-based alternative forecast was created by modelling the daily time series of volatility using intraday data. Andersen, Bollerslev, Diebold and Labys (2001) analysed the distribution of realised volatility for exchange rates, and found that realised volatility followed a process close to a unit root, with slowly declining autocorrelation. They found that this structure could be best modelled by fractionally integrated ARFIMA models. Moreover, they also pointed out that the distribution of logarithmic realised volatility was closer to normal, and thus modelling in log form results in more efficient estimates and better forecasts. Pong, Shackleton, Taylor and Xu (2004) compare the forecasts by the ARFIMA model and forecasts resulting from an ARMA(2,1) model,[12] and conclude that this latter predicts exchange rate volatility similarly well.

The realised volatility of the forint exchange rate behaves similarly. The distribution of logarithmic volatility is closer to the normal in terms of skewness and kurtosis than that of the non-logarithmic data. Slowly decaying autocorrelations indicate the strong persistence of the process. At the same time, the process is stationary – both the ADF and the Phillips-Perron tests reject the null hypothesis of the root of unit, and the KPSS test also confirms this result.

Thus, in addition to the forecasts based on the GARCH model, we also compare the performance of implied volatility to the forecasts obtained by fitting the ARMA(2,1) model to the realised volatility time series – or more precisely, to its logarithm. Similarly to the GARCH model, forecasts were made for the period between 31 December 2002 and 25 May 2005, on daily, weekly (5 days), one-month (21 days) and three-month (63 days) horizons. ARMA forecasts are also ex ante out of sample forecasts. Forecasts on a specific day were made on the basis of a model estimated on a sample from 2 January 2002 to the day of the forecast.

In order to maintain unbiasedness, forecasts of realised logarithmic volatility were transformed into forecasts on (daily) volatility by the use of the following formula:

$$\widehat{V}_{t+j} = \exp\left[\ln\left(\widehat{\sigma}_{t+j}\right) + \frac{1}{2}\operatorname{var}\left(\ln\left(\widehat{\sigma}_{t+j}\right)\right)\right]$$

[11] The quoting conventions of currency option markets are discussed in Csávás-Gereben (2005).

[12] The model choice was triggered by the results of Gallant, Hsu and Tauchen (1999), who concluded that the sum of two AR(1) processes was suited to describe the powerful persistence that characterised exchange rate volatility. The sum of two AR(1) processes can, in turn, be described in an ARMA(2,1) model.

where $\widehat{V}_{t+j}$ is the forecast of volatility on the $t+j$ day; $\ln(\widehat{\sigma}_{t+j})$ is the forecast of the logarithm of the volatility, and $\mathrm{var}(\ln(\widehat{\sigma}_{t+j}))$ is the variance of this latter.

4.2. METHODOLOGY

In line with the majority of the empirical literature on the information content of implied volatility, we use two types of regressions to find out whether options-based volatility is an unbiased and/or efficient predictor of realised volatility.

a. Testing forecasting ability and unbiasedness. Following the literature, we will test implied volatility's forecasting ability and unbiasedness using regression analysis (Day and Lewis (1992).

If implied volatility reflects market participants' expectations of future volatility, and the participants' expectations are rational, then the expected value of realised volatility corresponds to implied volatility. Formally,

$$E\left[\sigma_{t,T} \middle| \Phi_t\right] = IV_{t,T},$$

where Φ_t is the information set available for the market participants at time t. By definition, the expected deviation from the expected value equals zero under this information set, therefore:

$$\sigma_{t,T} = E\left[\sigma_{t,T} \middle| \Phi_t\right] + \varepsilon_t\,, \quad E\left[\varepsilon_t \middle| \Phi_t\right] = 0\,.$$

It follows from the above two equations that if the initial assumptions hold, then in the following linear regression between realised and implied volatilities

$$\sigma_{t,T} = \alpha + \beta_1 \cdot IV_{t,T} + \varepsilon_t \tag{1}$$

the parameters will have to meet the conditions $\alpha=0;\ \beta_1=1$.

The question whether implied volatility can forecast future realised exchange rate volatility can be decided by testing the significance of the β_1 parameter in a regression based on equation (1), and by examining the overall explanatory power of the regression (R^2). The question of unbiasedness can be analysed by testing the conditions $\alpha=0;\ \beta_1=1$.

b. Tests of forecast efficiency. Besides analysing whether implied volatility is an unbiased (rational) predictor of future realised volatility, we will also look at whether it contains all the relevant information available at a given moment of time. The efficiency test is carried out using the encompassing regression technique of Fair and Shiller (1990).

In the case of rational expectations, the forecast error (ε_t) is independent of the information set. In other words, the variables included in the information set cannot help explaining the error term, and cannot improve the forecast. As a consequence, the above-mentioned parameter restrictions are valid for all rational forecasts. The "quality" of the information set, i.e. the amount and significance of the included information, is reflected in the explanatory power of the regression.

Let us consider the case of analysing two forecasts, where the information set underlying the second forecast is a subset of the information set of the first forecast. In other words, the forecast $F^1_{t,T}$ includes all the information included in the second forecast $F^2_{t,T}$. We call $F^1_{t,T}$ an efficient predictor on the given information set, and $F^2_{t,T}$ has no role in explaining future volatility.

Formally, if $F^1_{t,T}$ is an efficient forecast, then in the regression

$$\sigma_{t,T} = \alpha + \beta_1 F^1_{t,T} + \beta_2 F^2_{t,T} + \varepsilon_t \tag{2}$$

the condition $\alpha=0;\ \beta_1=1;\ \beta_2=0$ must be met.

To reveal whether it is the implied volatility or the alternative variables that provide better forecasts we will compare the explanatory power *(R^2)* of the regressions based on equation 1. Then we will use the coefficients in regressions based on equation 2 to study whether the alternative forecasts have additional information *relative to one another*.

c. Methodological problems. Two conditions must be met for implied volatility to be a good and unbiased predictor of future volatility. Firstly, it needs to reflect market expectations. As we have seen in Chapter 2, a number of factors may prevent this. Secondly, expectations must also be rational. Our analysis is a joint test of these two hypotheses. If we reject this joint hypothesis, the econometric results alone will not be able to tell which of the assumptions was violated leading us to the rejection.

When assessing the results, it is important to note that our data sometimes have overlapping forecast horizons. The analysis was carried out over 4 different forecast horizons (1 day, 1 week and 1 and 3 months), using data of daily and weekly frequency. As a consequence, regression residuals will be autocorrelated in those estimations using daily data where the forecast horizon is one week or longer, and those using weekly data where the forecast horizon is one month or longer. In the case of daily data, for example, realised future volatility with 1 week horizon equals the sum of the volatilities of the next 5 days. As a result, consecutive observations are not independent of each other, and forecast errors will be serially correlated.

When the problem of overlapping forecast periods is present, ordinary least squares (OLS) parameter estimates continue to be unbiased; however, the estimate of the covariance matrix will be biased. In small samples, the problem may result in non-efficient estimation and – especially in the case of the time series based forecasts – an upwardly biased explanatory power.[13]

The literature offers two solutions for handling this problem. We can either take autocorrelation into account through imposing an appropriate structure for the covariance matrix, or reduce the sampling frequency to match the forecast horizon. As our sample is relatively short, we cannot choose the latter solution for every horizon, so we performed the analysis on data of daily and weekly frequencies[14] and adjusted the covariance matrix. In those cases where the forecast horizon exceeded the frequency of sampling, we used the autocorrelation-consistent covariance matrix proposed by Newey and West (1987).

As to the 1-day forecast horizon, 1-day-ahead implied volatility data was not available. For this horizon we used implied volatilities of the two shortest maturities: 1 week and 1 month. Results were similar in both cases, thus we present here only the results obtained by using the 1 month implied volatilities.[15]

An additional methodological issue is due to the fact that – as we pointed out earlier – both the dependent variable (realised volatility) and the explanatory variable (implied volatility) can only be observed with measurement errors. If the measurement error is not independent of the information set, the OLS estimation of parameters can be biased. To handle this problem, we also performed the estimations with the generalised method of moments (GMM), using the lagged values of explanatory variables as instruments. However, this alternative estimation method has not influenced the results.

4.3. FINDINGS

a. Forecasting ability and unbiasedness. In order to establish the forecasting ability and unbiasedness of implied volatility, we estimated equation (1) on daily and weekly data, using implied volatility as the explanatory variable and actual realised volatility as the dependent variable. The results are shown in Table 2.

[13] In the case of exchange-traded options, further difficulties may arise due the fact that the maturity of the options are not permanent, and thus implied volatility corresponds to fixed-event forecasts. Ignoring these characteristics of the time series could also lead to significant biases and could also change qualitative conclusions. Christensen and Prabhala (1998) point out that this was partly the reason why early literature found that implied volatility had no information content relative to time series based forecasts.

[14] Weekly time series were based on Tuesday data with the smallest number of missing observations.

[15] The explanatory power of regressions was nearly identical in both cases. However, the estimated parameter of 1 month implied volatility was somewhat closer to 1.

Table 2

The forecasting ability of forint/euro implied volatility

Forecast horizon	(daily data)			(weekly data)		
	R^2	α	β_1	R^2	α	β_1
1 day	0.37	-1.24* (-2.24)	0.75** (10.11)	-	-	-
1 week	0.39	1.07* (2.22)	0.5** (7.93)	0.37	1.14* (2.29)	0.49** (8.41)
1 month	0.18	2.53** (3.88)	0.34** (4.82)	0.17	2.54** (2.84)	0.33** (3.70)
3 months	0.00	5.08** (5.53)	0.05 (0.49)	0.00	5.19** (3.06)	0.03 (0.20)

*Adjusted t-values of the estimated parameters are shown in parentheses. Parameters significant at 5% are marked with *, while parameters significant at 1% are marked with **.*

Looking at the data, we can conclude that in the majority of forecast horizons, implied volatility seems to contain useful information related to future realised volatility, but it is not an unbiased predictor.

The coefficients of implied volatility differ from zero up to the 1-month horizon. The explanatory power of the models is between 40 and 17 per cent. Therefore, looking at both daily and weekly data, we can establish that implied volatility contains forecasting power up to the 1-month horizon. At the 3-month forecasting horizon, however, the coefficient of implied volatility is not statistically different from zero.

With the lengthening of the forecast horizon, the coefficients decrease. They are also significantly smaller than unity in every case. Furthermore, the constant generally differs from zero at a 5 per cent significance level, with the exception of the 1-day horizon. The joined test of the two parameters results in the rejection of the hypothesis of unbiasedness in every case. In line with the majority of the literature, we found that implied volatility is an upwardly biased predictor of future realised volatility over all horizons. Our findings are independent of whether we used daily or weekly data.

It is worth comparing our findings with the results of other studies in the literature. Table 3 shows a few results for various currency pairs.

Table 3

The explanatory power of the estimations (R^2) – international comparison

Source	Foreign currency pairs	R^2 1 day	1 week	1 month	3 months
Gereben and Pintér (2005)	forint/euro (daily data)	0.37	0.39	0.18	0.00
Pong at al. (2003)	pound/dollar deutschmark/dollar yen/dollar	0.39 0.45 0.43	0.38 0.44 0.42	0.4 0.34 0.46	0.19 0.17 0.4
Aguilar (1999)	Swedish krone/deutschmark Swedish krone/dollar deutschmark/dollar Australian dollar/US dollar	n. a. n. a. n. a. n. a.	0.3 0.16 0.29 0.37	0.35 0.33 0.34 0.52	0.19 0.23 0.18 0.54
Cincibuch and Bouc (2001)	Czech koruna/euro	n. a.	n. a.	0.26-0.48	n. a.

The table shows that at shorter maturities (1 day and 1 week) the explanatory power of the regressions estimated on the Hungarian data is similar to the results of other currency pairs. Over the 1 month forecast horizon, however, we can already see that the explanatory power is lower for the forint/euro than for other currency pairs. For the 3 months forecasting horizon, the forint/euro implied volatility does not provide information on future volatility, which is in contrast with the results of the other studies.

In summary, our findings indicate that the informational content of implied volatility on the Hungarian foreign exchange market decreases faster with the lengthening of the forecast horizon than on other markets. This is most likely due to the fact that the forint/euro options market is relatively young and undeveloped compared to the other foreign currencies examined, and the liquidity and turnover of longer-maturity options has been, so far, too low for quoted prices to accurately reflect market expectations.

b. Efficiency and alternative forecasts. In the second part of our empirical analysis we examine whether the alternative forecasting models contain extra information, in addition to the implied volatility, in those forecasting horizons where implied volatility is capable of forecasting future realised volatility to some extent. To answer this question we examine the coefficients of the encompassing regressions based on equation (2). The results are shown in Table 4.

Over the 1-day horizon, the ARMA model's forecasts have proven to be the best. In those cases where the regression contains additionally the implied volatility or the GARCH model forecasts, their coefficients are not significantly different from zero. Neither the implied volatility, nor the GARCH model forecasts contain information in addition to the forecasts of the ARMA model. Nevertheless, implied volatility contains more information than the GARCH forecast, and the coefficients of implied volatility in the regressions are higher. At the same time, however, when compared to the GARCH model, it is still not true to say that implied volatility contained all the information in itself – coefficients for both variables are significant at a 5% significance level.

The superior performance of the ARMA model over the 1-day horizon is partly due to the fact that implied volatility time series reflect expectations for the average volatility of the next 1 month horizon, as opposed to the volatility of the following day.

Table 4

Forecast efficiency of implied volatility – "encompassing" regressions

Forecast horizon	R^2	α	α_1 (implied)	β_2 (ARMA)	β_3 (GARCH)
1 day	0.37	-1.24* (-2.24)	0.75** (10.11)		
	0.59	-1.13** (-3.24)		1.22** (14.77)	
	0.31	1.35** (4.39)			0.52** (10.09)
	0.60	-0.94* (-2.34)	-0.15 (-1.3)	1.42** (8.53)	
	0.39	0.03 (0.07)	0.47** (4.6)		0.15* (1.98)
	0.59	-1.12** (-3.18)		1.24** (10.77)	-0.02 (-0.23)
	0.60	-0.91* (-2.31)	-0.2 (-1.67)	1.41** (8.3)	0.06 (0.81)
1 week	0.39	1.07* (2.22)	0.5** (7.93)		
	0.36	1.2** (2.93)		0.79** (9.2)	
	0.28	2.10** (4.62)			0.98** (6.29)
	0.41	0.78 (1.74)	0.3* (2.31)	0.39* (2.17)	
	0.39	1.03* (2.08)	0.44** (5.4)		0.18 (1.09)
	0.38	1.00* (2.16)		0.61** (5.2)	0.37 (1.82)
	0.41	0.77 (1.63)	0.26* (2.05)	0.38* (2.13)	0.13 (0.77)
1 month	0.18	2.53** (3.88)	0.34** (4.82)		
	0.19	3.04**(7.09)		0.46** (6.73)	
	0.09	3.57** (6.16)			1.10** (6.3.31)
	0.20	2.75** (4.06)	0.1 (0.61)	0.35 (1.89)	
	0.18	2.56** (3.83)	0.37** (4.55)		-0.15 (-0.43)
	0.19	3.09** (5.74)		0.48** (4.47)	-0.08 (-0.18)
	0.2	2.80** (3.91)	0.13 (0.9)	0.35 (1.83)	-0.20 (-0.52)
3 months	0.00	5.08** (5.53)	0.05 (0.49)		
	0.05	4.53** (11.12)		0.19** (3.23)	
	0.01	6.22** (8.20)			-0.21 (-1.02)

*Adjusted t values attached to estimated parameters are shown in parentheses. Parameters significant at 5% are marked with *, while parameters significant at 1% are marked with **.*

Over the 1-week horizon, however, implied volatility has proven to be an efficient predictor. Neither the ARMA nor the GARCH models improved the forecasts, and their coefficients did not differ from zero at any usual significance levels.

Over the 1-month horizon the GARCH forecasts do not provide any additional information relative to the ARMA forecast and the implied volatility. It is interesting to note that the latter two forecasts have explanatory power in themselves; however, when they appear together in a regression, none of the coefficients are significant. This indicates that these two forecasts are very strongly correlated, and have very similar information content.

Over the 3 month time horizon implied volatility has no significant explanatory power. The ARMA forecast outperforms here the other two candidates.

In summary, implied volatility cannot be considered to be an efficient predictor of future volatility over most time horizons. Thus, in addition to implied volatility, time series forecasts based on intraday data – and the indicator based on the ARMA model in particular – should also be taken into account, when one attempts to forecast future volatility. These results correspond to the earlier findings by other authors on a wide range of currency pairs.

5. Final conclusions

Both the findings of international literature and our analysis based on Hungarian data suggest that the Black-Scholes implied volatility is a useful indicator of future volatility, especially on the foreign exchange market. Although it is usually biased, and does not necessarily contain all past information, it can be stated that changes in the implied volatility usually lead to changes in the future realised volatility as well.

Naturally, the Black-Scholes implied volatility is not a "perfect" indicator, and its constraints must be taken into account for practical applications. First, the information content of longer-maturity options on less developed option markets, such as the Hungarian one, has been relatively low so far. Secondly, as implied volatility is upwardly biased, this bias must be adjusted when forecasting the level of future volatility. And finally, to create efficient forecasts of future volatility, the information from intraday exchange rate data is also worth considering, in addition to the implied volatility.

For the accurate interpretation of changes in implied volatility, the possible biasing factors described in section 2 of our study should be taken into account. Occasionally these biasing factors, which result from the simplifying assumptions of the Black-Scholes model, may cause changes in the value of implied volatility, without changes in the market expectations of future exchange rate uncertainty. On the forint/euro option market, such bias can arise from the existence of the exchange rate band, from occasionally large interest rate volatility, and low liquidity. These bias factors should always be taken into account when drawing conclusions from changes in implied volatility.

Literature

AGUILAR, J. (1999), GARCH, implied volatilities and implied distributions: an evaluation for forecasting purposes, *Sveriges Riksbank Working Paper* No. 88.

AÏT-SAHALIA, Y. & A. LO (2000), Nonparametric risk management and implied risk aversion, *Journal of Econometrics*, Vol. 94, 9-51.

AÏT-SAHALIA, Y. & P. A. MYKLAND (2003), How often to sample a continuous-time process in the presence of market microstructure noise, *NBER Working Paper* 9611.

ALIZADEH, S., M. BRANDT & F. X. DIEBOLD (2002), Range-based estimation of stochastic volatility models, *Journal of Finance* Vol. 57, 1047-1091.

ANDERSEN, T. G. & T. BOLLERSLEV (1998), Answering the sceptics: yes standard volatility models do provide accurate forecasts, *International Economic Review*, Vol. 39, 885-905.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD & P. LABYS (2001), The distribution of exchange rate volatility, *Journal of the American Statistical Association*, Vol. 96, 42-55.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD & P. LABYS (2003), Modeling and forecasting realised volatility, *Econometrica*, Vol. 71, 579-625.

BACKUS, D., S. FORESI, K. LI & L. WU (1997), Accounting for biases in Black-Scholes, manuscript, *Stern School of Business*, New York University.

BAHRA, B. (1997), Implied Risk-neutral Probability Density Functions From Option Prices: Theory and Application, *Bank of England Working Paper* No 66

BANDI, F. M. & J. R. RUSSEL (2004), Microstructure noise, realised volatility, and optimal sampling, *Econometric Society 2004 Latin American Meetings* 220, Econometric Society.

BARNDORFF-NIELSEN, O. E. & N. SHEPHARD (2004), Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics*, Vol. 2, 1-36.

BATES, D. S. (1996) Testing Option Pricing Models, in G.S. Maddala and C. R. Rao, (ed.), Statistical Methods in Finance (Handbook of Statistics, v. 14), Amsterdam, Elsevier, 567-611.

BATES, D. S. (2001), The Market for Crash Risk, *NBER Working Papers* No. 8554.

BATES, D. S. (2003), Empirical Option Pricing: A Retrospection, *Journal of Econometrics* Vol 116, 387-404.

BERTOLA, G. & L. E. O. SVENSSON (1993), Stochastic Devaluation Risk and the Empirical Fit of Target-Zone Models, *Review of Economic Studies*, Vol. 60, No. 3, 689-712.

BERTOLA, G. & R. J. CABALLERO (1992), Target Zones and Realignments, *American Economic Review*, Vol. 82, No. 3, 520-36.

BIS (2005), Central Bank Survey of Foreign Exchange and Derivatives Market Activity 2004 – Final Results, Bank for International Settlements, Monetary and Economic Department.

BLACK, FISCHER & MYRON S. SCHOLES (1973), The pricing of options and corporate liabilities, *Journal of Political Economy*, Vol. 81, 637-654.

BLAIR, B. J., S. POON & S. J. TAYLOR (2001) Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high frequency index returns, *Journal of Econometrics* 105, 5-26.

BLISS, R. & N. PANIGIRTZOGLOU (2001), Recovering risk aversion from options, *Federal Reserve Bank of Chicago Working Paper* No. 2001-15

BRANDT, M. & F. X. DIEBOLD (2003), A no-arbitrage approach to range-based estimation of return covariances and correlation (second version), *PIER Working Paper* No. 03-013.

BREUER, P. (2003), How Does the Volatility Risk Premium Affect the Informational Content of Currency Options?, manuscript.

CAMPA, J. M. & K. P. H. CHANG (1998), ERM Realignment Risks and Its Economic Determinants as Reflected in Cross-Rate Options, *The Economic Journal*, Vol. 108, No. 449, 1046-1066.

CANESSO DE ANDRADE, S. & B. M. TABAK (2001), Is it worth tracking dollar-real implied volatility?, *Banco Central do Brasil Working Paper*, March 2001.

CANINA, L. & S. FIGLEWSKI (1993), The informational content of implied volatility, *Review of Financial Studies* 6, 659-681.

CHRISTENSEN, B. J. & C. STRUNK HANSEN (2002), New Evidence on the Implied-Realised Volatility Relation, *European Journal of Finance* 8, 187-205.

CHRISTENSEN, B. J. & N. R. PRABHALA (1998), The relation between implied and realised volatility, *Journal of Financial Economics* 50, 125-150.

CHRISTOFFERSEN, P. & S. MAZZOTTA (2004), The informational content of over-the-counter currency options, *ECB Working Paper* No.366.

CINCIBUCH, M. & P. BOUC (2001), Interpretation of Czech FX options, *CNB Working Paper* No. 36.

CORRADO, C. J. & T. MILLER (2004), The forecast quality of CBOE volatility indexes, J*ournal of Futures Markets* (forthcoming).

COX, J., J. INGERSOLL, & S. ROSS (1985), An Intertemporal General Equilibrium Model of Asset Prices, *Econometrica*, Vol. 53, 363-384.

CSÁVÁS, CS. & GEREBEN, Á. (2005) Traditional and Exotic Options on the Hungarian Foreign Exchange Market, *MNB Occasional Papers* 35.

DAY, T. & C. LEWIS (1992), Stock market volatility and the information content of stock index options, *Journal of Econometrics*, Vol 52, 267-287.

DE JONG, F., F. C. DROST & B. J. M. WERKER (2001) A Jump-Diffusion Model for Exchange Rates in a Target Zone, *Statistica Neerlandica*, Vol. 55, 269-299.

DUAN, J. C. (1995), The GARCH Option Pricing Model, *Mathematical Finance*, Vol. 5,13-32.

EDERINGTON, L. H. & W. GUAN (2002), Is Implied Volatility an Informationally Efficient and Effective Predictor of Future Volatility? *The Journal of Risk*, Vol. 4, No. 3, 29-46.

FAIR, R. C. & R. J. SHILLER (1990), Comparing information in forecasts from econometric models, *American Economic Review* Vol. 80, 375-389.

FLEMING, J. (1998), The quality of market volatility forecasts implied by S&P 100 index option prices, *Journal of Empirical Finance* 5, 317-345.

FLEMING, J. (1999), The economic significance of the forecast bias of S&P 100 index option implied volatility, *Advances in Futures and Options Research* 10, 219-251.

FRANKE, G., R. C. STAPLETON & M. G. SUBRAHMANYAM (1998), Who Buys and Who Sells Options: The Role of Options in an Economy with Background Risk, *Journal of Economic Theory*, 82, 89-109.

FRENNBERG, P. & B. HANSSON (1996), An Evaluation of Alternative Models for Predicting Stock Volatility: Evidence From a Small Stock Market, *Journal of International Financial Markets*, Institutions and Money, 5, 117-134.

GALLANT, R., C-T. HSU & G. TAUCHEN (1999), Using daily range data to calibrate volatility diffusions and extract the forward integrated variance, *The Review of Economics and Statistics* 81, 617-631.

GEREBEN, Á. (2002), Extracting market expectations from option prices: an application to over-the-counter New Zealand dollar options, Reserve Bank of New Zealand Discussion Paper Series DP2002/04.

GONZALEZ PEREZ, M. T. (2004), Información contenida en la Volatilidad Implícita del IBEX-35, preliminary version submitted for the conference of Universitat Pompeu Fabra XII. Foro de Finanzas.

HULL J. & A. WHITE (1987), The pricing of options on assets with stochastic volatilities, Journal of Finance, Vol. 42, 281-300.

JACKWERTH, J. C. (2000), Recovering Risk Aversion from Option Prices and Realised Returns, *Review of Financial Studies*, Vol. 13, No. 2, 433-51.

JORION, P. (1995), Predicting volatility in the foreign exchange market, J*ournal of Finance* 50, 507-528.

KOOPMAN, S. J., B. JUNGBACKER, E. HOL (2004), Forecasting Daily Variability of the S&P 100 Stock Index using Historical, Realised and Implied Volatility Measurements, T*inbergen Institute Discussion Papers* 04-016/4.

KRUGMAN, P. (1991), Target zones and exchange rate dynamics, *Quarterly Journal of Economics*, Vol. 56, 669-682.

LAMOUREUX, C. G. & W. D. LASTRAPES (1993), Forecasting stock-return variance: toward an understanding of stochastic implied volatilities, *Review of Financial Studies* 6, 293-326.

MALZ, A. M. (1995), Currency Option Markets and Exchange Rates: A Case Study of the U.S. Dollar in March 1995, Current Issues in Economics and Finance, *Federal Reserve Bank of New York*, Vol. 1, No. 4.

MALZ, A. (1997), Option-implied Probability Distributions and Currency Excess Returns, Federal Reserve Bank of New York, *Staff Report* No. 32.

MARTENS, M. & J. ZEIN (2004), Predicting financial volatility: high-frequency time series forecasts vis-à-vis implied volatility, *Journal of Futures Markets* Vol. 24, No. 11, 1005-1028.

MELICK, W. & THOMAS, C. (1997), Recovering an Asset's Implied PDF from Option Prices: An application to Crude Oil During the Gulf Crisis, *Journal of Financial and Quantitative Analysis* Vol. 32, 91-115.

MERTON, R. C. (1973). Theory of rational option pricing, *Bell Journal of Economics and Management Science*, Vol. 4, No. 1, 141-183.

MNB (2004), Report on Financial Stability, December.

MNB (2005), Report on Inflation, February.

NEELY, C. J. (2002), Forecasting foreign exchange volatility: why is implied volatility biased and inefficient? And does it matter?, *Federal Reserve Bank of St. Louis Working Paper*, 2002-107D.

NEWEY, W. K. & K. D. WEST (1987), A simple positive-definite heteroscedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703-708.

PONG, S., M. B. SHACKLETON, S. J. TAYLOR & X. XU (2004), Forecasting currency volatility: a comparison of implied volatilities and AR(FI)MA models, *Journal of Banking and Finance* 28, 2541-2563.

POON, S. H. & C. W. J. GRANGER (2003), Forecasting volatility in financial markets: a review, *Journal of Economic Literature* Vol. 41, 478-539.

POTESHMAN, A. M. (2000), Forecasting future volatility from option prices, AFA 2001 New Orleans.

RUBINSTEIN, M. (1994), Implied Binomial Trees, *Journal of Finance*, Vol. 49, 771-818.

SVENSSON, L. E. O. (1991), Target zones and interest rate variability, *Journal of International Economics*, Vol. 31. No. 1-2, 27-54.

TABAK, B. M., E. J. CHANG & S. CANESSO DE ANDRADE (2002), Tracking Brazilian Exchange Rate Volatility, Econometric Society 2004 Far Eastern Meetings 487, Econometric Society.

TAYLOR, S. J. & X. XU (1995), Conditional volatility and the informational efficiency of the PHLX currency options markets, *Journal of Banking and Finance* 19, 803-821.

TAYLOR, S. J. & X. XU (1997), The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance* 4, 317-340.

Appendix 1

During 2003, the Magyar Nemzeti Bank (MNB, the central bank of Hungary) carried out two large increases in its key policy rate. The two rate hikes occurred in June and December. The level of short term interest rates, together with the euro-forint interest rate differential, suddenly increased in both cases. During this period, the implied volatility of the forint-euro exchange rate also moved together with the interest rate differential (see Chart 1). The data suggest that the value of the forint-euro implied volatility had been influenced and biased by the significant changes in the interest rates – or the interest rate differential – in this period. In other words, it is possible that the increase in implied volatility in the period in question was only partly due to greater uncertainty surrounding the exchange rate, and the other part of the increase may be attributed somehow to the change in the interest rate (differential).

Chart 2

Interest rate differential and implied volatility

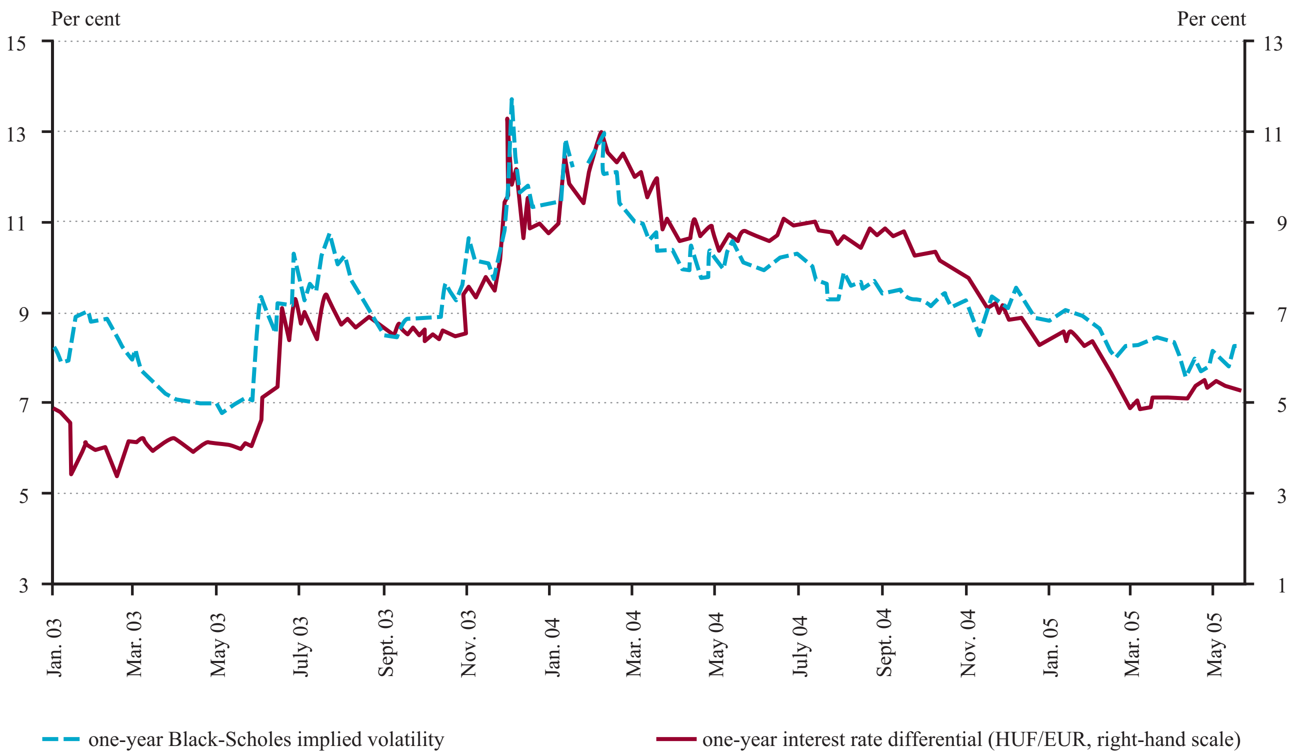

In this appendix we present a modified version of the Black-Scholes implied volatility based on Merton's option pricing model for stochastic interest rates (Merton, 1973) . We show how the use of this alternative measure of implied volatility alters the outlook regarding the future uncertainties surrounding the exchange rate.

Merton's model extends the original Black-Scholes model, by incorporating stochastic interest rates (or stochastic price for the risk-free bond) into the option pricing formula. Let us denote the time-t value of a risk-free bond paying one unit at time T as $P(t, T)$. If the expected yield until time T – $R(t, T)$ – is known, then $P(t, T)=e^{-R(t, T)}$. Let us assume that $P(t, T)$ follows the stochastic process

$$\frac{dP(t,T)}{P(t,T)} = \mu_P dt + \sigma_P dz_P \quad ,$$

where dz_P is a Wiener process. Keeping all the other assumptions of the Black-Scholes model unchanged, Merton showed that in this case the price of call options on stocks can be calculated according to the following formula:

$$c = SN(d_1) - P(t,T)XN(d_2)$$, where

$$d_1 = \frac{\ln(S/X) - \ln P(t,T) + \hat{\sigma}^2 (T-t)/2}{\hat{\sigma}\sqrt{T-t}},$$

$$d_2 = d_1 - \hat{\sigma}\sqrt{T-t}$$, and

$$\hat{\sigma}^2 = \left(\sigma_S^2 + \sigma_P^2 - 2\rho\sigma_S\sigma_P\right) \qquad \text{(A1.1)}$$

In the above formula, ρ denotes the contemporaneous correlation of the bond price and the share price, σ_S stands for the volatility of the stock price, and σ_P is the bond price volatility. As we can see, Merton's formula is very similar to the original Black-Scholes formula. The only difference involves the method of calculating the volatility parameter entering to the formula. In Merton's model the value of this volatility depends on the volatility of both the stock price and the bond price, as well as the correlation of the two prices.

Merton's formula can easily be extended to foreign exchange options as well. For foreign exchange options, the discount factor is determined by the interest rate differential; the bond should be replaced with a virtual bond with a yield equal to the interest rate differential.

It follows from the Merton model that if we calculate implied volatility ($\hat{\sigma}$) based on the Black-Scholes formula, the indicator we obtain is indeed dependent on three factors:

1. Expected exchange rate volatility (the "true" implied exchange rate volatility, σ_S);
2. the volatility of the bond yielding the interest rate differential *(σ_P)*; and
3. the correlation between the exchange rate and the bond price *(ρ)*.

If the second and third factors are not negligible, implied volatility calculated from the Black-Scholes formula is a biased predictor of the realised exchange rate volatility.

As a first step, let us quantify the price of a 1-year bond paying the interest rate differential. The time series of the 1 year interest rate differential can easily be calculated using zero-coupon yield curves, and the price of the virtual bond can be obtained through the $P=e^{-R}$ formula.

Let us now attempt to filter the effects of interest rate volatility from the Black-Scholes implied volatility indicator using formula (A1.1). The Black-Scholes implied volatility *($\hat{\sigma}$)* is known. Unfortunately, as there are no traded options on the interest rate differential, we have quantified the volatility of the interest rate differential-paying bond with the historical volatility of the virtual bond price. We have also calculated the correlation of the virtual bond and the exchange rate *(ρ)* on the basis of historical correlation

Using the formula, and by substituting historical bond volatility we obtain

$$\hat{\sigma}^2 = \left(\sigma_S^2 + \sigma_{P,hist}^2 - 2\rho\sigma_S\sigma_{P,hist}\right),$$

where $\hat{\sigma}$ is the observed Black-Scholes implied volatility. Once the substitutions are made, the value of implied volatility filtered from interest rate volatility effect (σ_S can be determined (see Chart 3).

Chart 3

Original and filtered implied volatility

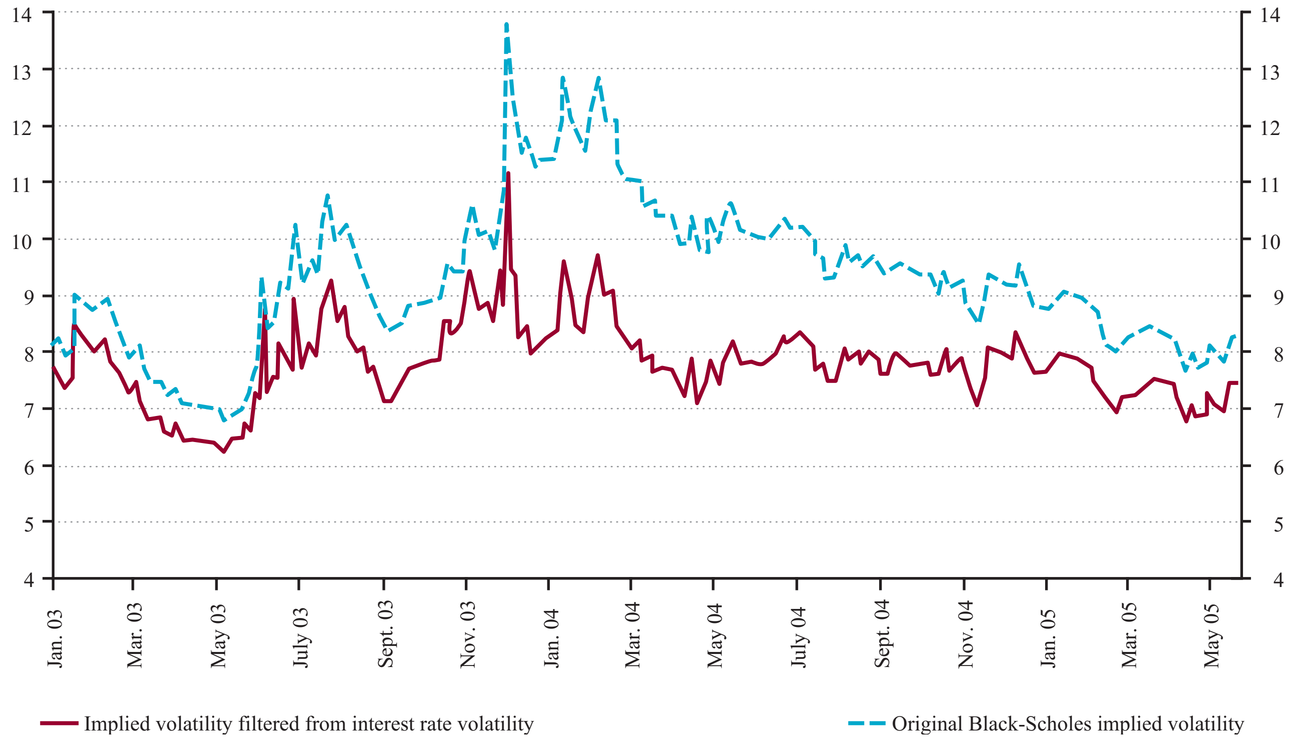

It is clear from the chart that the filtered implied volatility indicator shows significantly lower expected volatility than the Black-Scholes value. The difference is significant only when large shifts in the interest rate are present, or in other words when the interest rate volatility increases are significantly greater.

It seems that at times of large changes in the interest rates, the Black-Scholes implied volatility may be biased due to increased interest rate volatility. This should be taken into account when interpreting implied volatility as an indicator of future exchange rate uncertainty. In such cases, it would be worthwhile to use the transformed indicator, filtered from interest rate volatility.

Appendix 2

Chart 4

30 minute log yields

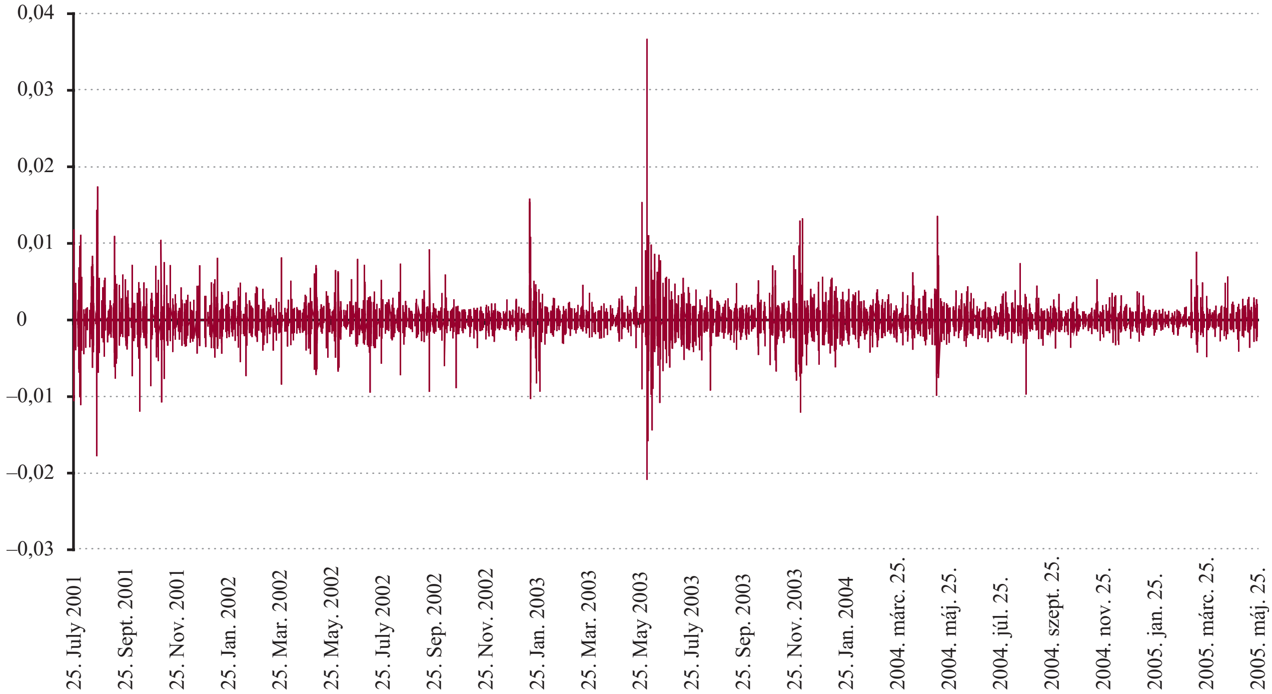

Chart 5

Logarithm of 30 minute exchange rate

Chart 6

Realised volatility and "bipower variation"

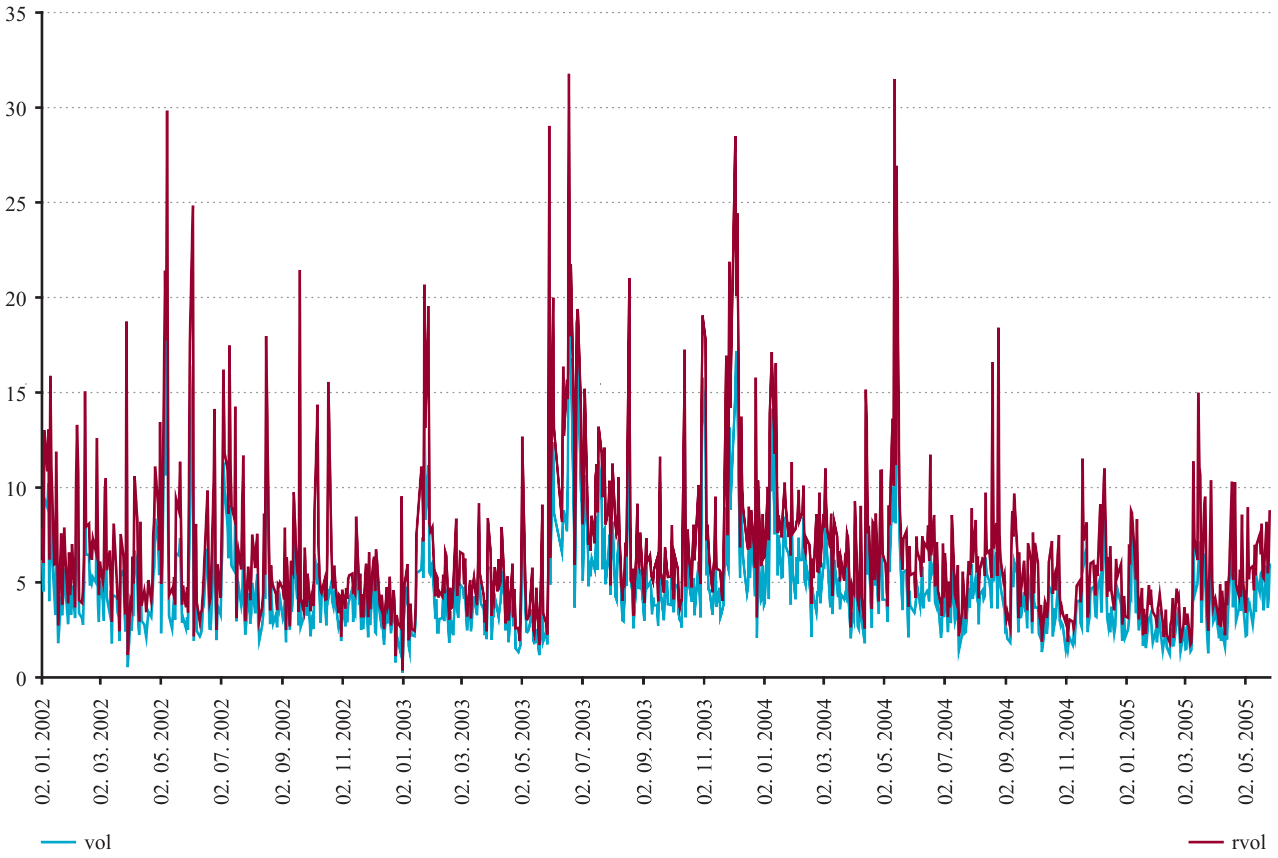

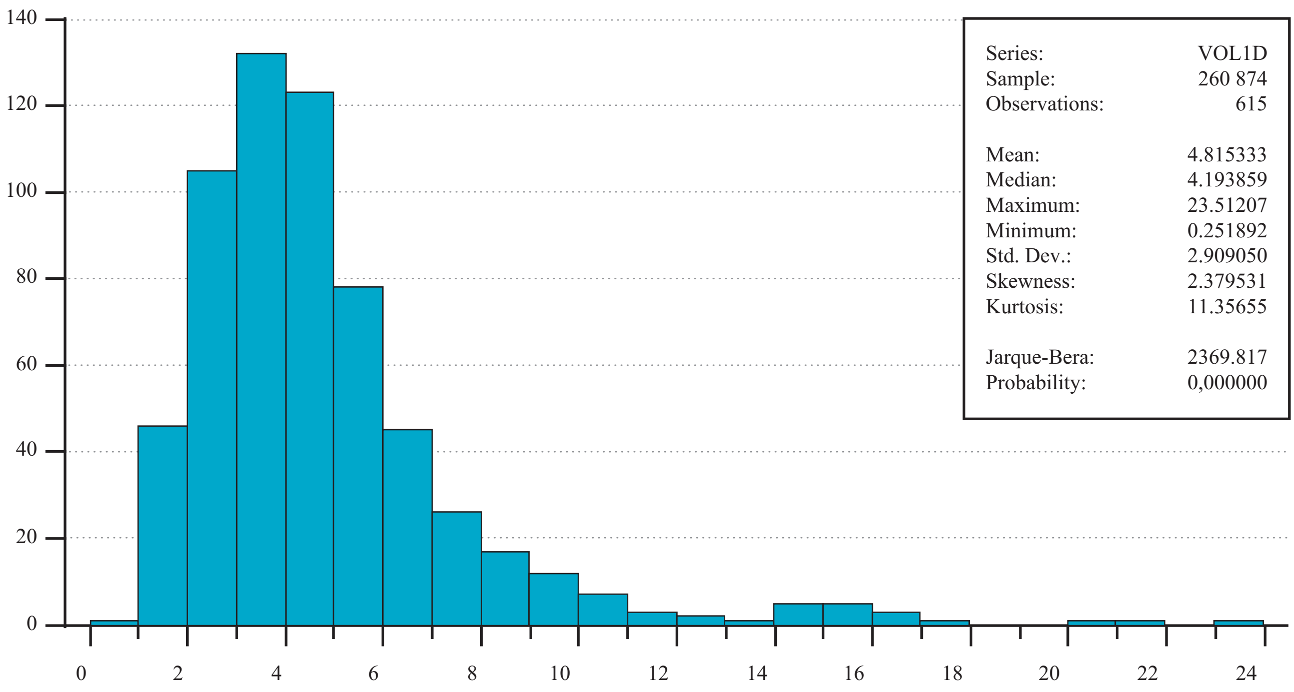

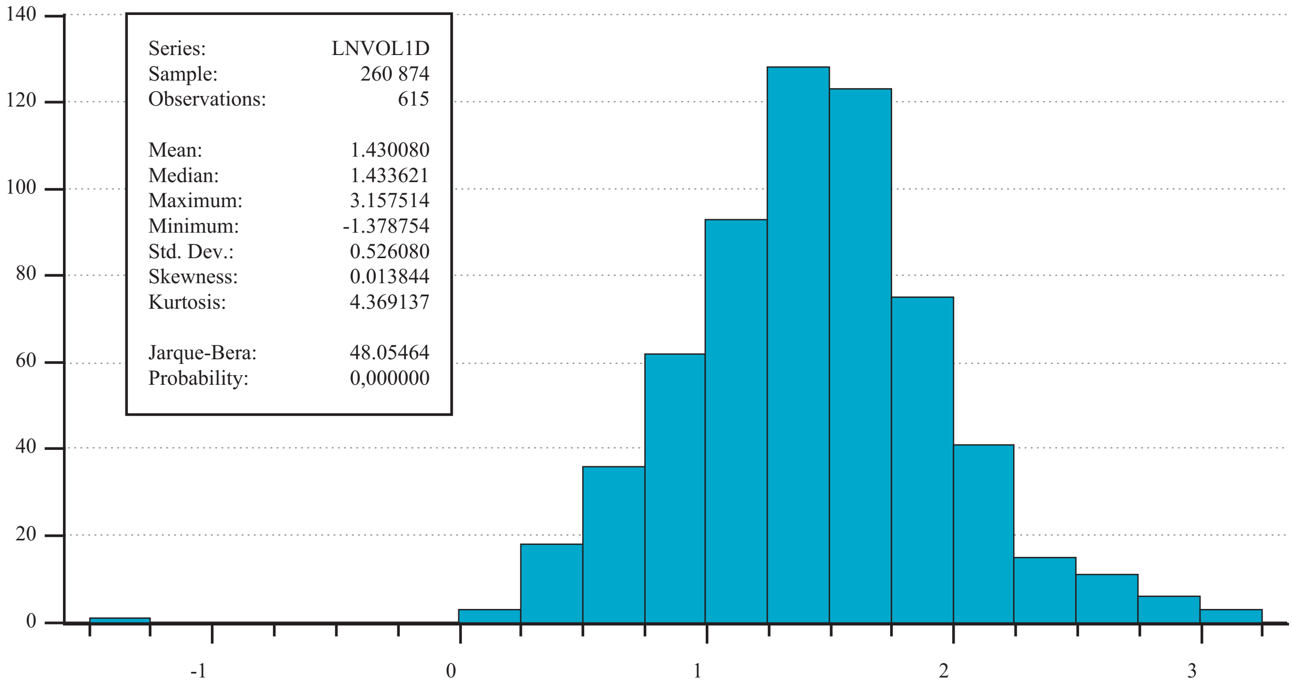

Chart 7

Realised future volatility and alternative forecasts over 1 day, 1 week, 1 and 3 months horizon

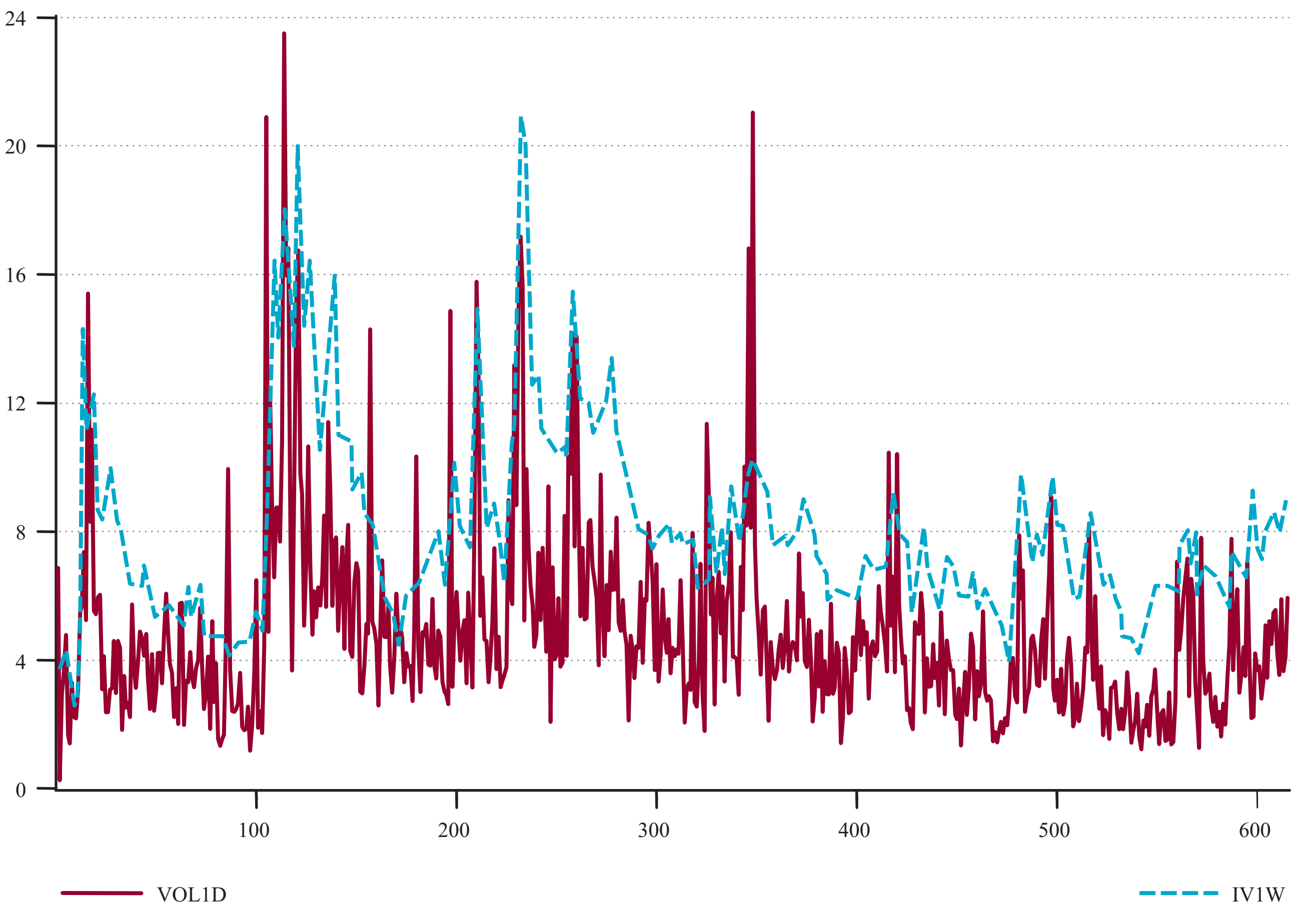

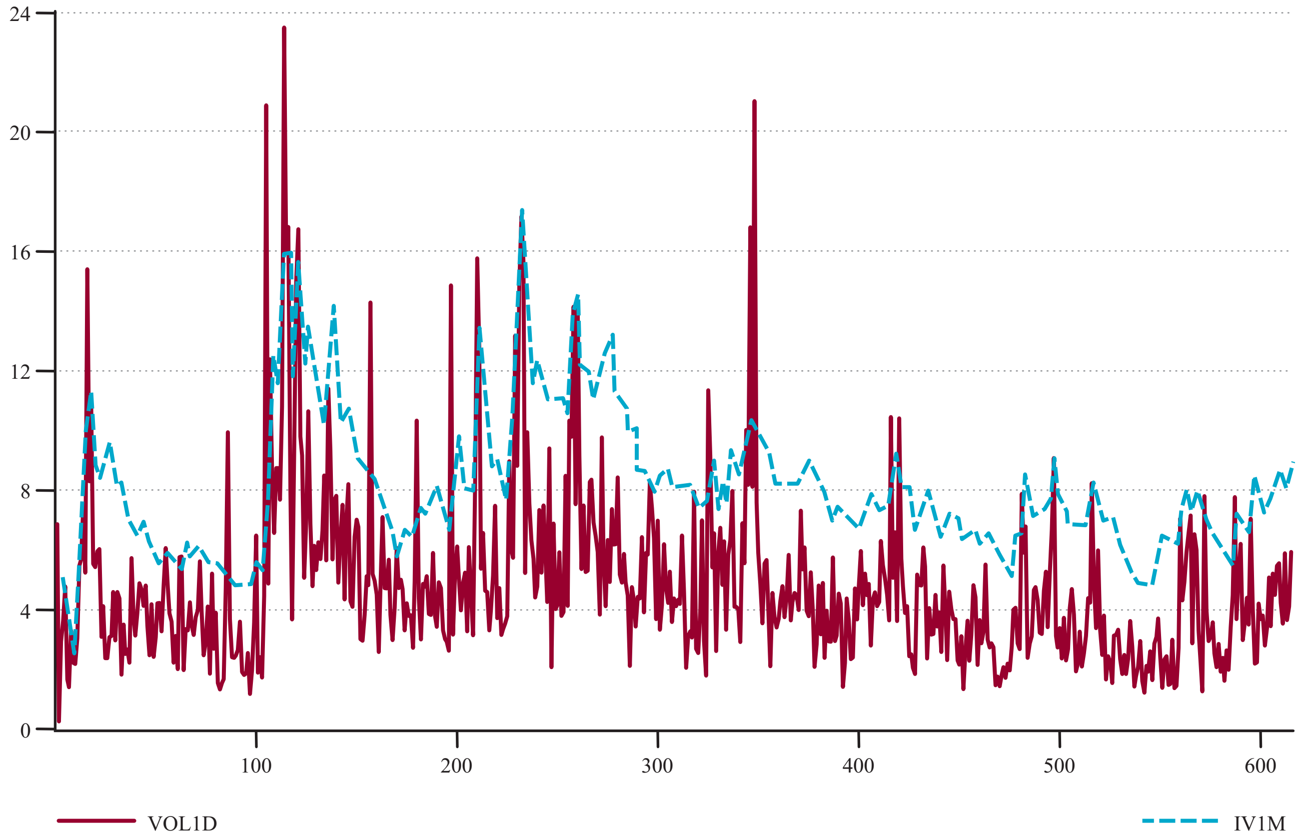
24
20
16
12
8
4
0
100
200
300
400
500
600
VOL1D
IV1M

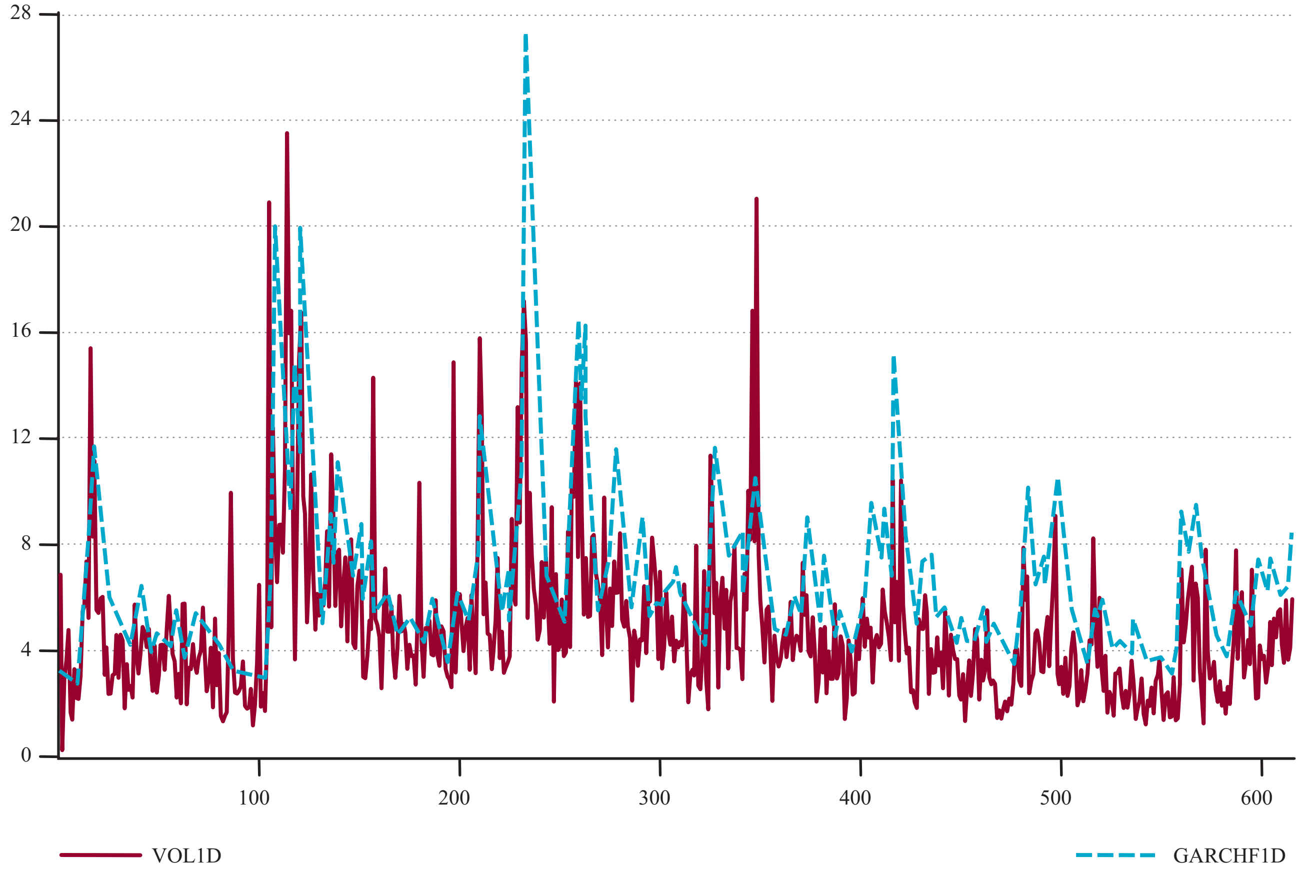
28
24
20
16
12
8
4
0
100
200
300
400
500
600
VOL1D
GARCHF1D

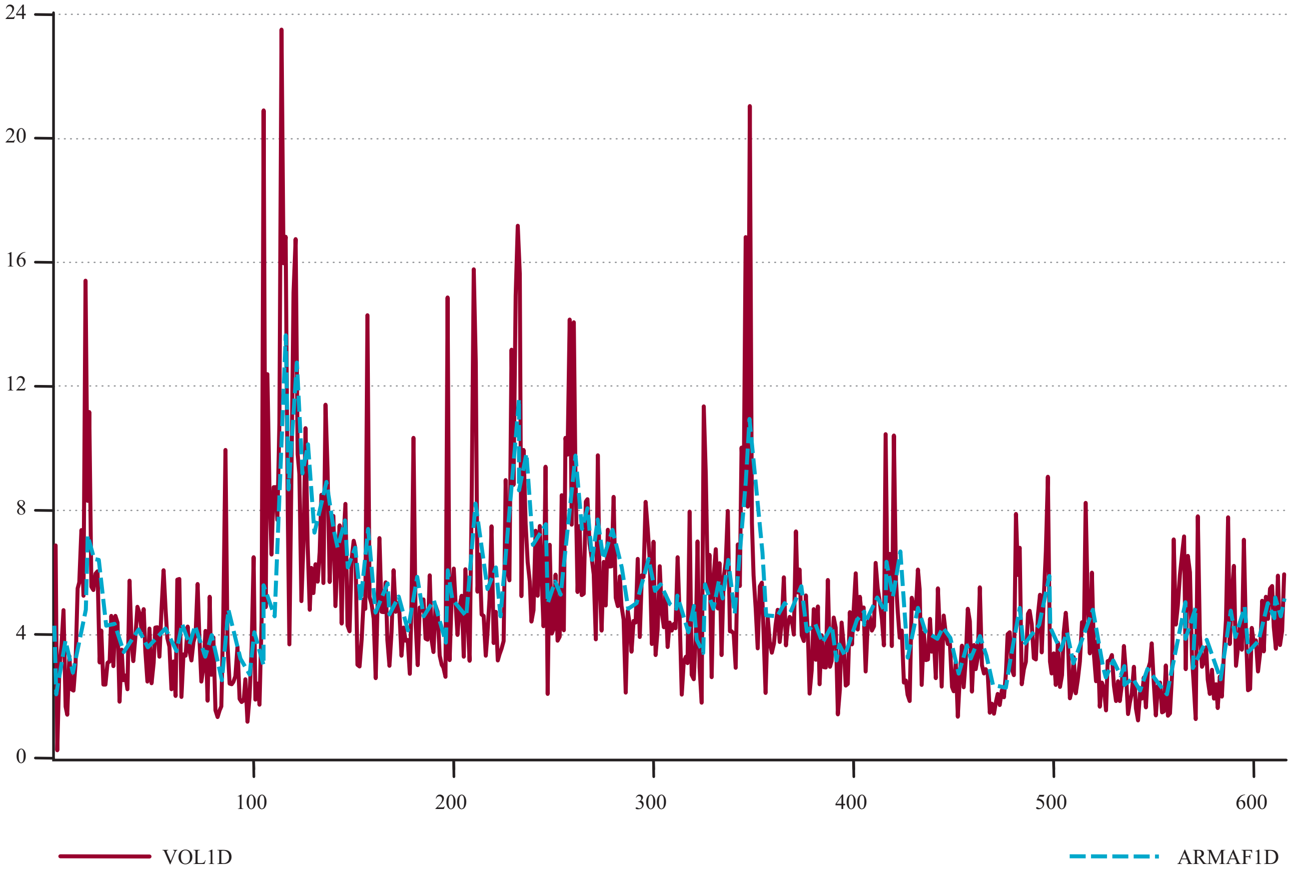
24
20
16
12
8
4
0
100
200
300
400
500
600
VOL1D
ARMAF1D

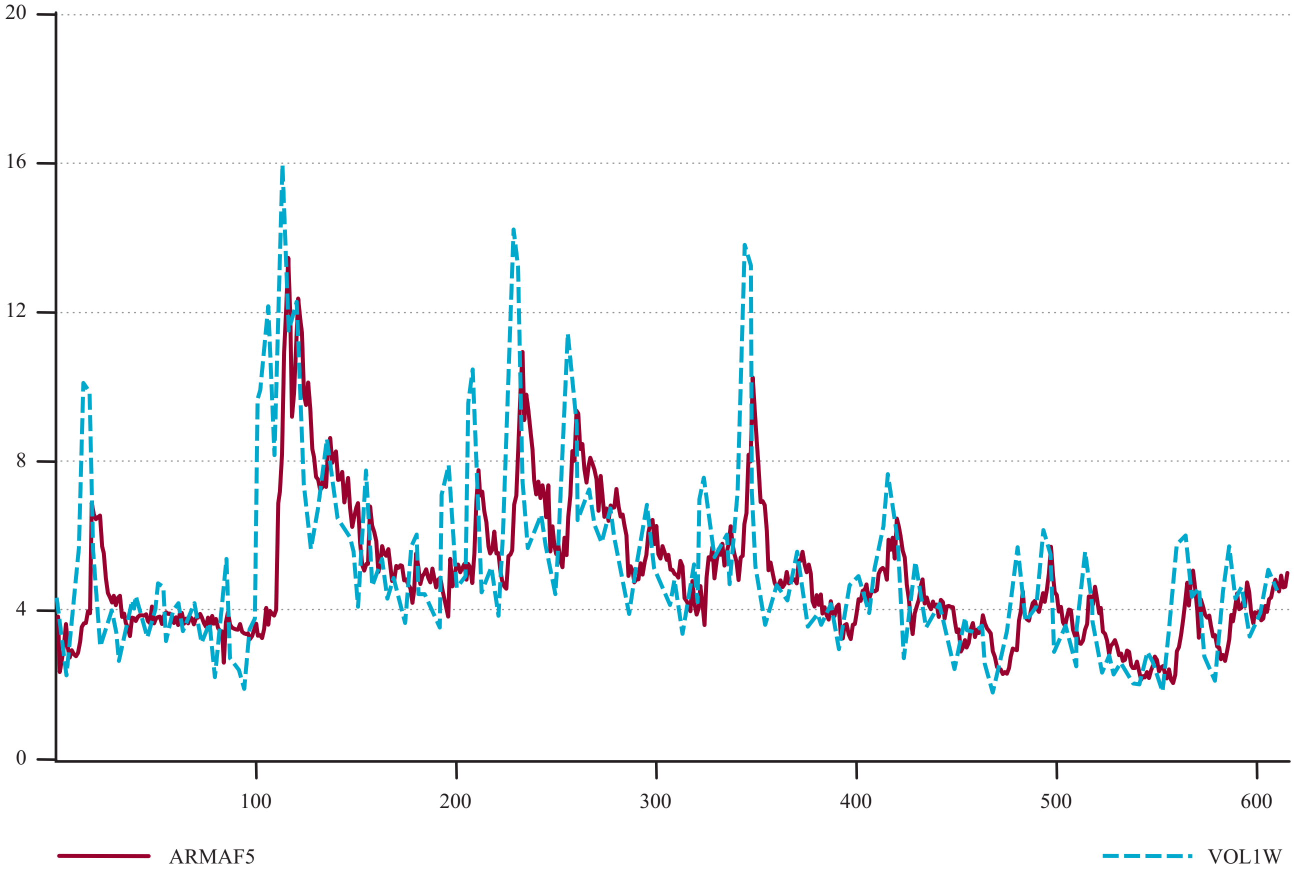
20
16
12
8
4
0
100
200
300
400
500
600
ARMAF5
VOL1W

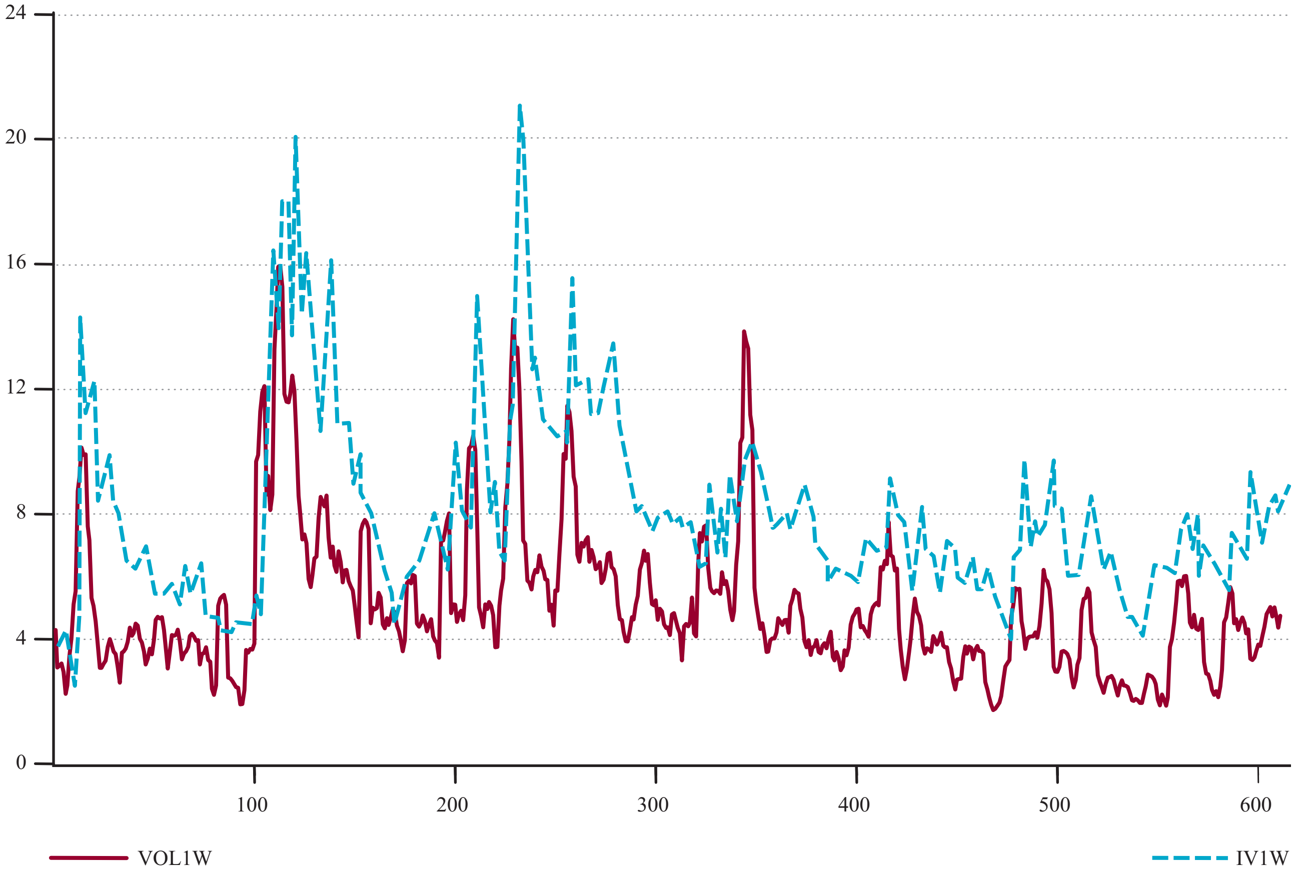
24
20
16
12
8
4
0
100
200
300
400
500
600
VOL1W
IV1W

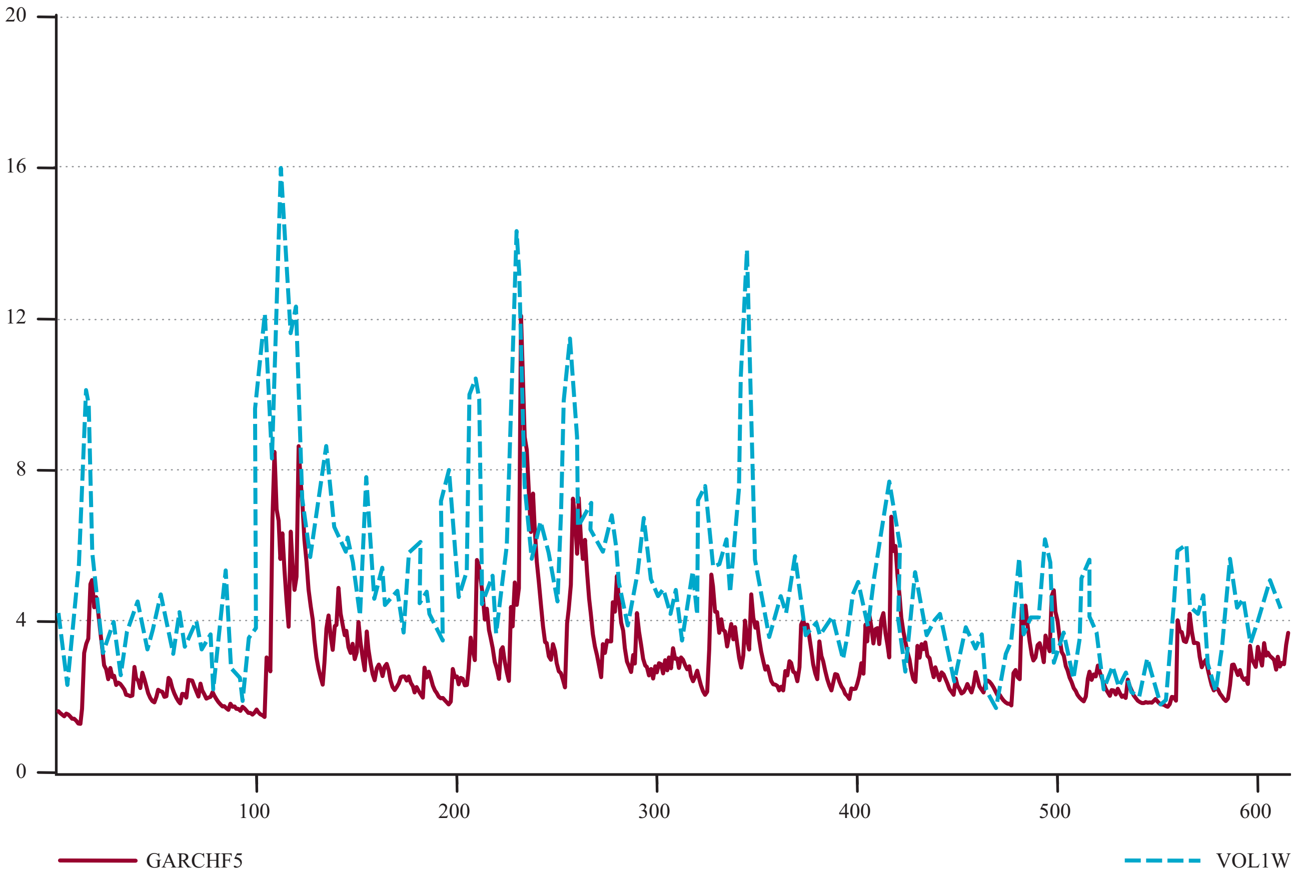
20
16
12
8
4
0
100
200
300
400
500
600
GARCHF5
VOL1W

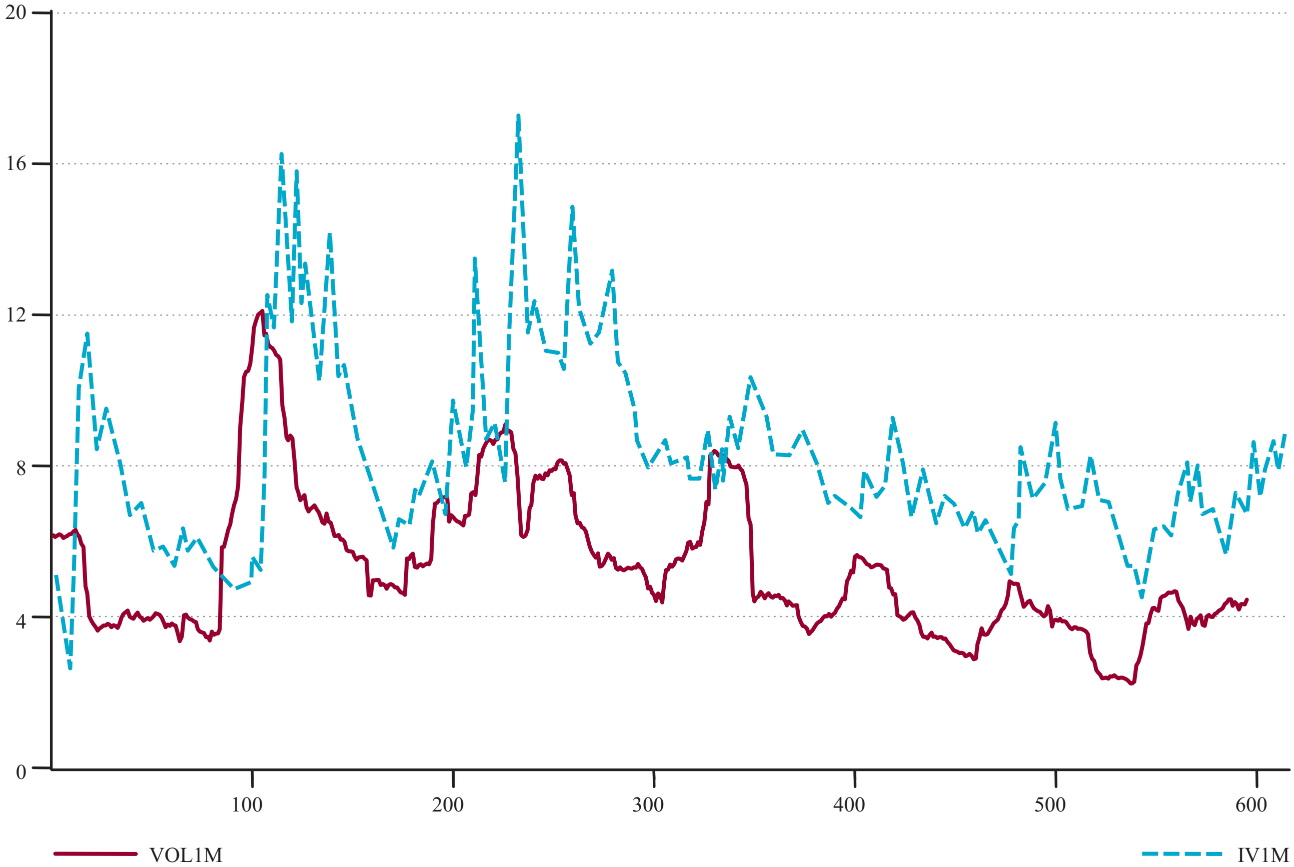
20
16
12
8
4
0
100
200
300
400
500
600
VOL1M
IV1M

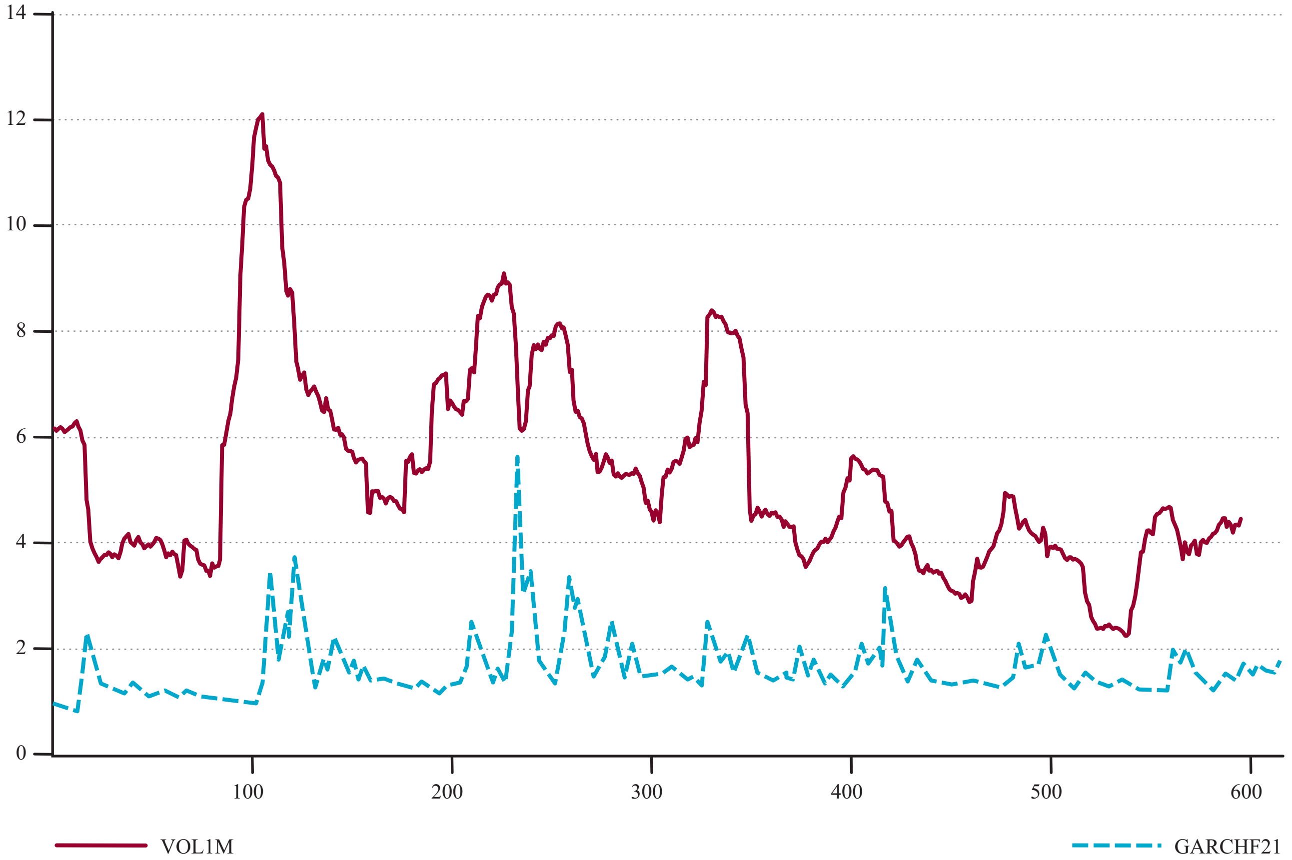
14
12
10
8
6
4
2
0
100
200
300
400
500
600
VOL1M
GARCHF21

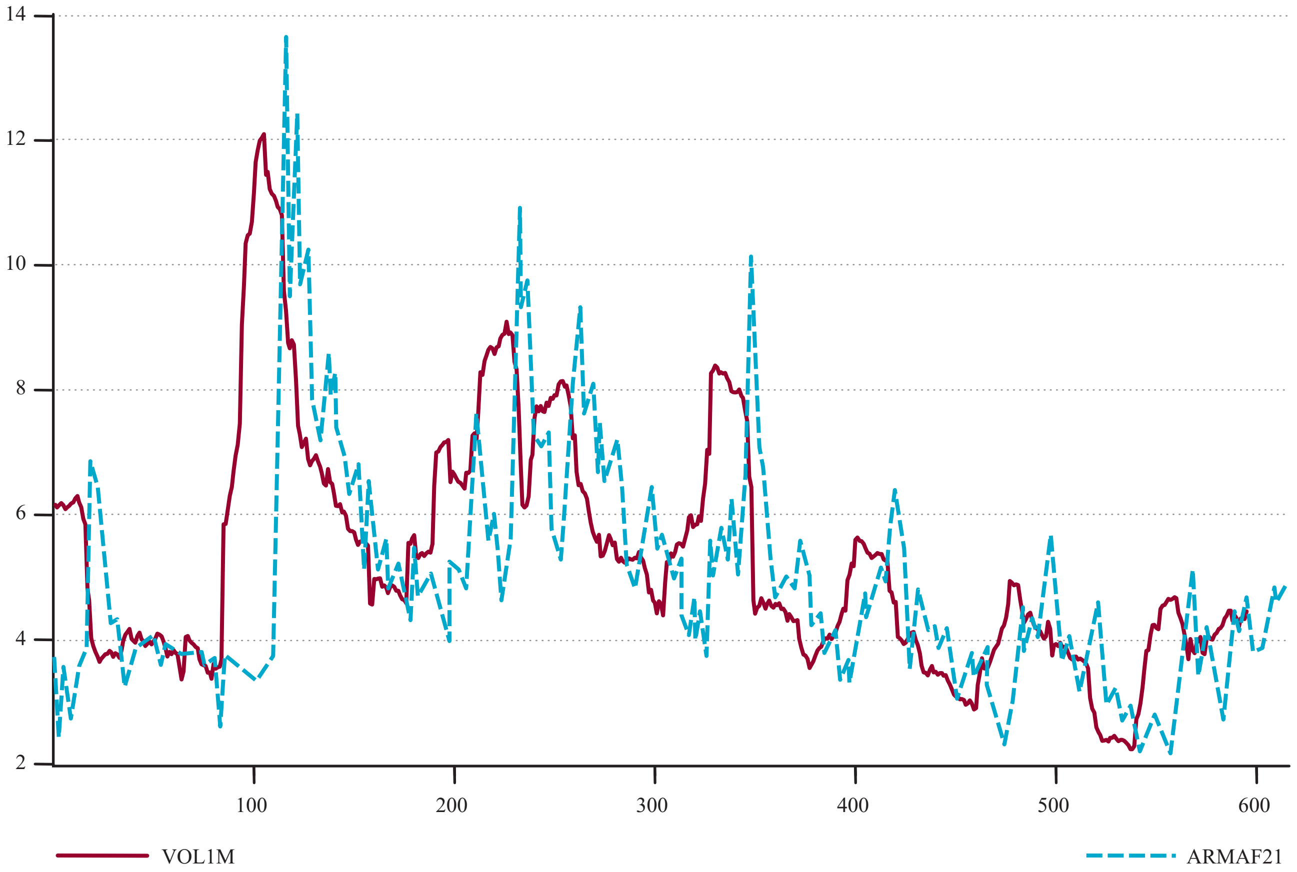
14
12
10
8
6
4
2
100
200
300
400
500
600
VOL1M
ARMAF21

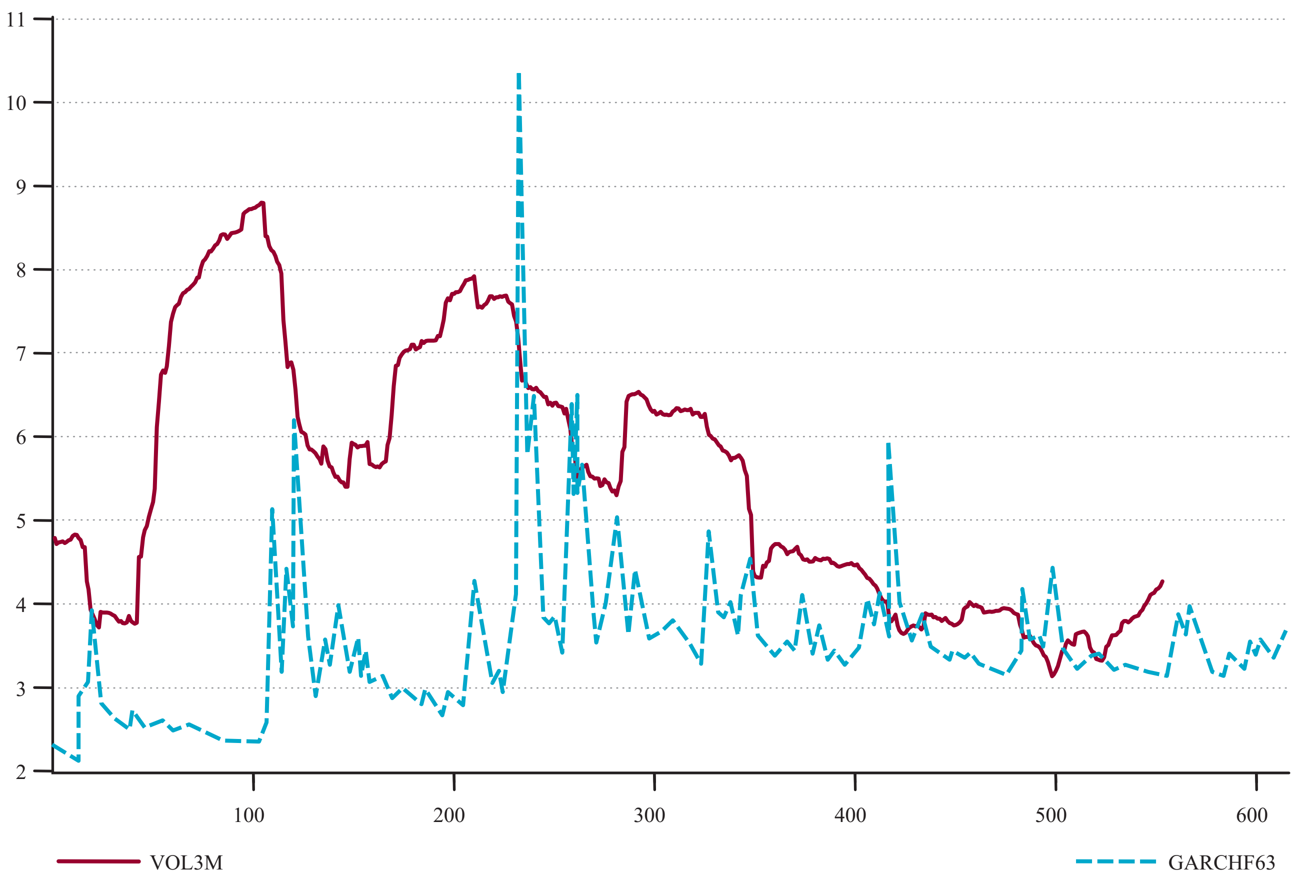
11
10
9
8
7
6
5
4
3
2
100
200
300
400
500
600
VOL3M
GARCHF63

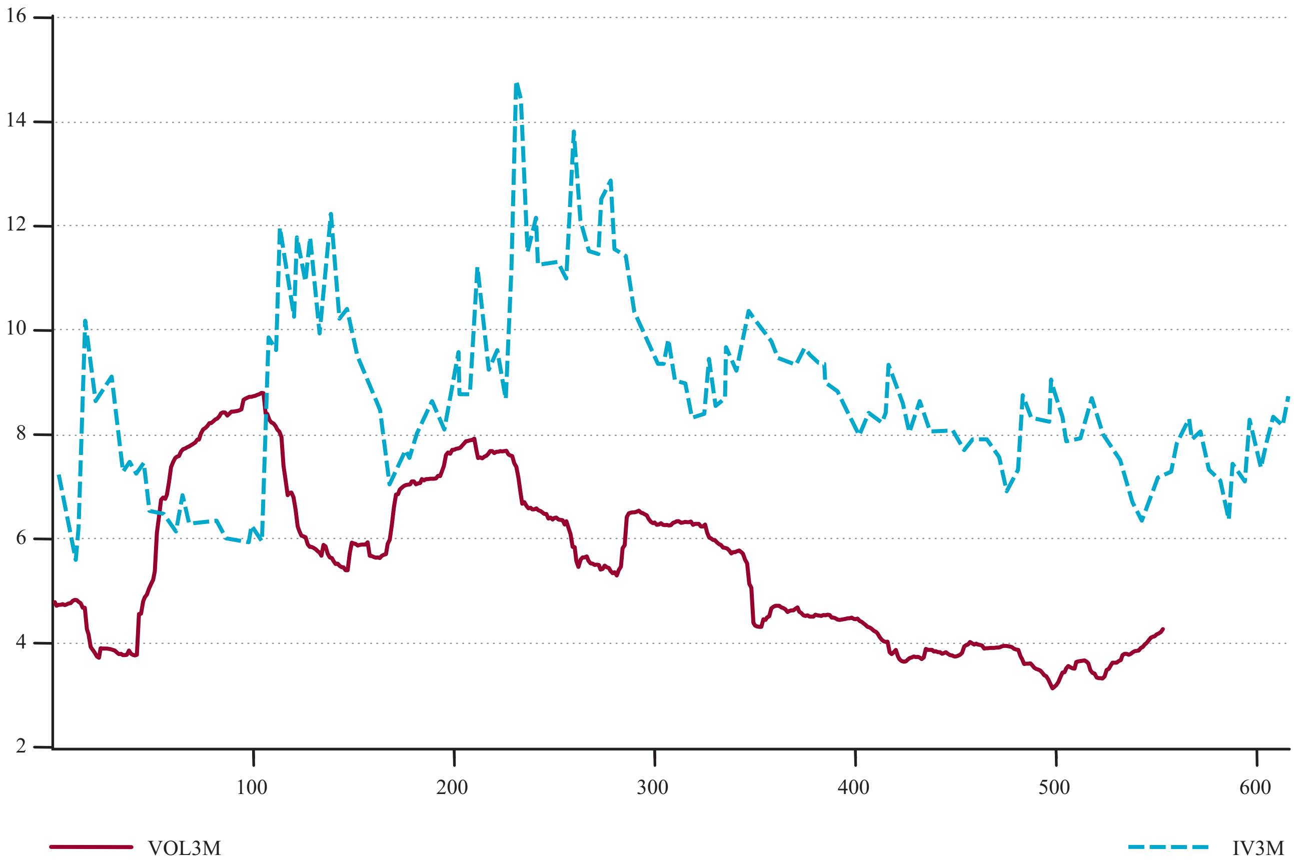
16
14
12
10
8
6
4
2
100
200
300
400
500
600
VOL3M
IV3M

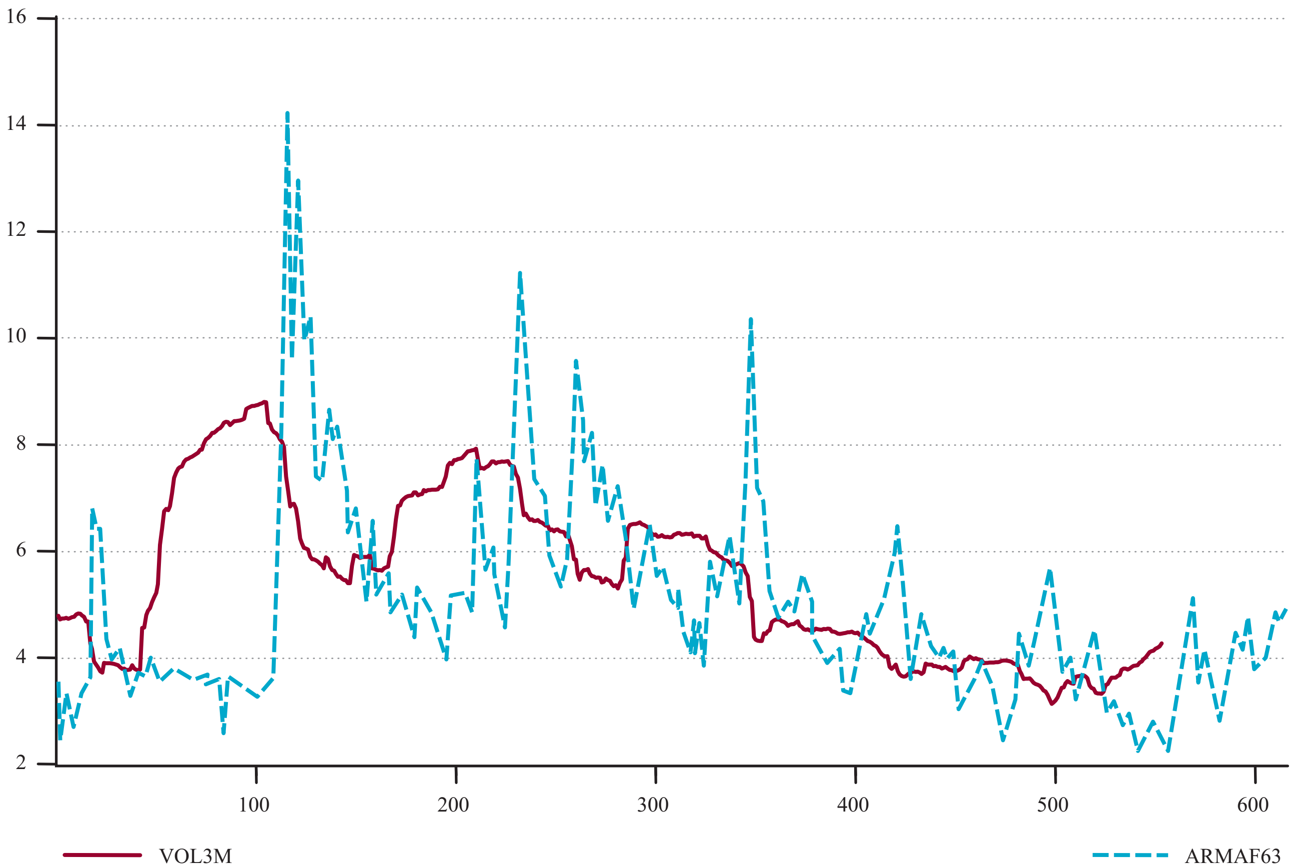
16
14
12
10
8
6
4
2
100
200
300
400
500
600
VOL3M
ARMAF63

MNB Occasional Papers 39.
Implied volatility of foreign exchange options: is it worth tracking?
May, 2005

Print: D-Plus
H–1033 Budapest, Szentendrei út 89–93.