

Aguilar, Javiera

Working Paper

GARCH, Implied Volatilities and Implied Distributions: An Evaluation for Forecasting Purposes

Sveriges Riksbank Working Paper Series, No. 88

Provided in Cooperation with:

Central Bank of Sweden, Stockholm

Suggested Citation: Aguilar, Javiera (1999) : GARCH, Implied Volatilities and Implied Distributions: An Evaluation for Forecasting Purposes, Sveriges Riksbank Working Paper Series, No. 88, Sveriges Riksbank, Stockholm

This Version is available at:

<https://hdl.handle.net/10419/82454>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

GARCH, Implied Volatilities and Implied Distributions: An Evaluation for Forecasting Purposes¹

Javiera Aguilar

Abstract

Volatility implied in option prices reflects the market participant's beliefs about future volatility and incorporates information that is not historical. Implied volatility is therefore widely believed to perform better as an indicator of future volatility than other forecasts based on historical time-series. In this study, I investigate the information content and predictive power of implied volatility from currency options traded on the OTC-market. Furthermore, I evaluate implied volatility both against other forecasts based on option prices and against volatility forecasts from models that are strictly historical by nature such as different GARCH models. I find that implied volatility has predictive power in forecasting future volatility, at least for shorter forecast horizons, although in most cases the forecasts are not unbiased. Furthermore, for some currencies GARCH volatility forecasts outperform implied volatility.

1. Introduction

There are several reasons why one should be interested in trying to predict future volatility. Expectations about future volatility plays a crucial role in finance theory. For instance, it is of major importance for market participants to make accurate predictions of future volatility since volatility is an important input in asset pricing, in

¹ I am most grateful to Paul Söderlind, Anders Vredin and Clas Wihlborg for comments. I would specially like to thank Peter Hördahl for very helpful comments and support. Comments from seminar participants at Sveriges Riksbank were also helpful. I would also like to thank participants at 1998 Bank of England workshop "Extracting Information from Financial Markets", and in particular Bhupinder Bahra, Kevin Chang and Gabriele Galati.

hedging and in portfolio management. The ability to predict future volatility is also important to central banks as an input in the monetary policy making process.

The majority of models that have been developed for forecasting volatility generally rely on the past behaviour of asset prices and are therefore backward looking by nature. Option valuation, on the other hand, is forward looking since option prices depend on expected future volatility. Thus, many economists argue that implied volatility derived from options should have better forecast ability than models based on historical data. Implied volatility contains market participant's beliefs about future events and incorporates information that is not strictly historical such as forthcoming publication of macroeconomic indicators.

Previous studies on the predictive ability of implied volatility have mainly been conducted on stock options. Even though the results are mixed a major part of the studies conclude that implied volatility is a poor indicator of future volatility. Lamoureux and Lastrapes (1993) investigate the information content of implied volatility from options on individual stocks and find that forecasts based on historical time-series have more predictive power than implied volatility. Canina and Figlewski (1993), who examine implied volatility derived from S&P 100 index options, find that implied volatility has almost no predictive power for future volatility. In a recent study, Christensen and Prabhala (1998) also examine the predictive power of implied volatility on S&P 100 index options. In contrast to previous work they find that implied volatility outperforms historical volatility in forecasting future volatility. Christensen and Prabhala conclude that the difference in their results compared to those of Canina and Figlewski is a consequence of using longer time series and non-overlapping data. They also find evidence that there was a regime shift following the October 1987 stock market crash, with implied volatility being more biased before than after the crash. Christensen and Prabhala's data set covers a period from 1983 to 1995 while Canina and Figlewski only examine a pre-crash period.

Empirical research conducted on currency options, on the other hand, in general come to the conclusion that implied volatility on short maturity contracts performs well in forecasting future volatility and contains information that is not present in historical volatility. Jorion (1995) investigates the information content of implied volatility from currency options traded on the Chicago Mercantile Exchange. Jorion finds that statistical time-series models are outperformed by the volatility implied in short-term options although implied volatility appears to be a biased forecast. Galati and Tsatsaronis (1996) also evaluate the predictive power of volatility implied in currency options, with the difference that they use options traded on the "over-the-counter" market. They also conclude that implied volatility from short-maturity options performs well in forecasting future volatility although it is a biased estimator. For longer horizons they find that neither historical nor implied volatility provides a good forecast of future volatility.²

Most of the work on volatility implied in currency options has been conducted on options written on the major currencies. The purpose of this paper is to evaluate the information content and the predictive power of implied volatility of currency options

² See also Bank of Japan (1995).

on the Swedish krona against the US dollar and the D-mark. It is of interest to study if volatility implied in options on a smaller peripheral currency such as the Swedish krona performs equally well in predicting future volatility as implied volatility from options on the major currency pairings. Options on smaller currencies are less heavily traded, which might result in less efficiently priced options, making implied volatility from these options a less reliable predictor. For comparison purposes, I investigate the predictive power of implied volatility for the DEM/USD exchange rate, which should be a liquid and well traded derivative contract. Furthermore, I compare the predictive power of implied volatility for the Swedish krona with that of the Australian dollar against the US dollar. The Australian dollar is also a relatively small currency and, more importantly, Australia conducts a monetary policy similar to the Swedish in the sense that both countries have an inflation target and freely floating exchange rates.

A further aim of this study is to evaluate implied volatility against two different types of forecast models. The first area of investigation is to evaluate if implied volatility is the preferred option based predictor. Many studies on the implied volatility of exchange rates do not investigate if there are other forecast measures derived from option prices that outperform implied volatility. The implied probability distribution perceived by the market participants might differ from the probability distribution implied by the Black-Scholes model. Taking this possible difference into account might yield better volatility forecasts than the *at-the-money* implied volatility obtained from the Black-Scholes model. I therefore evaluate implied volatility against the standard deviation of implied risk-neutral probability distributions derived from option prices.

I also evaluate implied volatility against predictors of future volatility that are backward looking by nature such as historical volatility and different generalised autoregressive conditional heteroskedasticity (GARCH) models. In contrast to Jorion(1995), the different GARCH volatility forecasts are based on the same data that is available to market participants at the time they make their volatility estimates in order to make the forecast evaluations as realistic as possible. I estimate conditional volatility forecasts based on GARCH(1,1) and EGARCH(1,1) models. The EGARCH model is estimated in order to take into account possible asymmetries in the volatility of the exchange rates. Heynen and Kat (1994) find that the GARCH and EGARCH models perform equally well in forecasting exchange rate volatility. Nevertheless, it is possible that the exchange rate volatility of smaller currencies such as the Swedish krona is asymmetrically affected by positive and negative shocks. It has, in fact, been observed that volatility tends to be higher when the Swedish krona depreciates than when it appreciates.

The remainder of the paper is organised as follows. Section two gives a description of the data used in the analysis. Section three examines the predictive power of option-based volatility forecasts by evaluating implied volatility against the standard deviation of probability density functions implied in option prices. Section four deals with forecasts based on historical data. Implied volatility is evaluated against historical volatility and the different GARCH conditional volatility forecasts. Finally, the study is concluded in section five.

2. Data

The option data employed in this study consists of quotations from the so-called “over-the-counter” (OTC) market. OTC options are traded directly between banks and have non-standardised expiration dates, as opposed to options traded on an exchange. This means that it is possible any day to trade an option, with for instance three month to maturity. Another advantage of using prices from the OTC market is that liquidity is better as most of the currency options are traded in the OTC market.³ Furthermore, smaller currencies such as the Swedish krona are only traded in the OTC market. Yet another and probably the most important advantage is that prices in the OTC market are quoted directly in terms of volatility, which means that, the implied volatility does not have to be calculated from the option prices. One therefore avoids errors that may arise when using data for the spot exchange rates and option prices that are not simultaneously recorded as may be the case when backing out implied volatility using exchange traded options.

The data has been provided by Citibank and consists of daily quotations of implied volatility for currency options on the SEK/DEM, SEK/USD, DEM/USD, and AUD/USD exchange rates with maturities between one week and one year. The quotations are collected at noon (London time) and consist of the average of the bid and ask volatilities. The implied volatility quotations are for “at-the-money forward” options which is the most traded option type in the OTC-market.⁴ Citibank has also provided simultaneous exchange rate data, making it possible to calculate the ex post or realised volatility for exactly the same forecast horizon as implied volatility. This exchange rate series are also used to generate the GARCH and EGARCH conditional volatility forecasts. Daily quotations on two common combinations of “out-of-the-money” options, the so called “strangle” and “risk reversal”, have been provided by Chase Manhattan Bank. These derivative products are used in the derivation of the implied probability distributions.

The sample period contains about 1500 daily observations and runs from January 3, 1993 to December 30, 1998. There are several reasons for choosing this period. After a period of severe speculative attacks against the Swedish krona the Riksbank abandoned the fixed exchange rate in November 1992. Sweden had not experienced a floating exchange rate since the early thirties. Thus, there was clearly a regime shift in November 1992 and there is reason to believe that the period leading up to the abandonment of the peg may have contained serious peso problems.

3. The predictive power of option based volatility measures

This section evaluates the ability of volatility measures extracted from option prices to forecast future volatility. Option prices are by nature forward looking since they

³ Average daily turnover on the OTC-currency options market was \$87 billion on the second quarter of 1998, while turnover on exchanges was only \$1.9 billion. See “International Banking and Financial Market Developments” BIS Quarterly Review, November 1998.

⁴ An option is “at-the-money forward” when its exercise price is equal to the forward price of the underlying asset.

incorporate information about the market participant's expectations regarding the future exchange rate. Hence, this paper evaluates whether the market's estimate of future volatility is actually a good forecast.

The Black-Scholes (1973) model is the most used option pricing model.⁵ The model assumes that the price of the underlying asset follows a geometric Brownian motion (GBM).⁶ This implies that the price of the underlying asset will be lognormally distributed and that its return will be normally distributed with constant variance. Based on these and other assumptions the model relates the price of a currency call option C , to five variables:

$$C=f(S, X, \sigma, t, i-i^*),$$

where S is the underlying currency, X is the strike price, σ is the standard deviation of the underlying currency, t is time to maturity, and $i-i^*$ is the difference between the domestic and foreign interest rate. All variables are known when the option is priced except for the standard deviation σ , which has to be estimated by the market participants. If the option price is taken as given and σ is solved for, we get what is commonly referred to as the *implied volatility*. In other words, implied volatility is the market participant's best "guess" of the future volatility during the life of the option. As previously argued, implied volatility should be a superior forecast since market participants may take into account information of anticipated future events in addition to historical information.

In order to be able to evaluate the predictive power of implied volatility, future or *realised* volatility must be defined. Realised volatility is defined as follows for a given exchange rate S_t :

$$RV_t = \sqrt{\frac{1}{m} \sum_{j=1}^m [\ln(S_{t+j} / S_{t+j-1}) - \mu]^2 \times 250} \quad (1)$$

where m is the number of trading days during the option's life, i.e. during the forecast horizon, and where μ is the estimated average of the returns, and where time is measured in days. That is, realised volatility is the annualised standard deviation of daily continuous returns assuming 250 trading days per year. Realised volatility is then calculated for the entire exchange rate series by moving the m -day window one step forward for each day. One-month implied volatility, for example, may then be compared to realised volatility calculated using a window of 20 trading days. Thus, implied volatility will be compared to ex post volatility during the month for which implied volatility predicted future volatility.

⁵ It was later modified for currency options by Garman-Kohlhagen (1983).

⁶ According to the GBM, the change in the asset price, dS , can be written as $dS = \mu S dt + \sigma S dz$, where μ is the expected return, σ is the volatility, dt denotes the time increment and dz is the increment of a standardised Wiener process.

Before evaluating whether implied volatility is a good predictor of future volatility, it is illustrative to study graphically how these two variables behave. Graphs 1 to 4 show plots of implied and realised volatility for each currency and maturity of the data set. A quick glance at the graphs shows that realised and implied volatility seem to be quite strongly correlated, at least for the shorter maturities.

Another feature of the data is that both implied and realised volatility become less variable as the time horizon gets longer. This is most apparent for realised volatility, which is extremely variable for the shortest maturities while one-year realised volatility is quite stable. For realised volatility this is due to the way it is calculated over a moving window, which results in more stable volatility as this window gets longer. Nevertheless, implied volatility also seems to be more variable for shorter maturities, which is especially apparent for the one-week implied volatility.

As the time horizon gets longer the relationship between realised and implied volatility becomes less clear. For maturities over 6 months the correlation between realised and implied volatility seems to be very weak for all exchange rates. Hence, just by studying the two volatility series graphically the predictive power of implied volatility appears to be very poor for longer horizons.

3.1 Estimations on the predictive power of implied volatility

A simple way of testing the predictive power of the at-the-money implied volatility as a forecast of future volatility is by estimating the following OLS regression:

$$RV_t = \alpha + \beta IV_t + \varepsilon_t \quad (2)$$

where RV_t is realised volatility for the forecasted period starting at time t , IV_t is the at-the-money implied volatility at time t , α is a constant, β is the regression coefficient and ε_t is the error term. In order for implied volatility to have some explanatory power for realised volatility, β has to be statistically significantly different from zero ($\beta \neq 0$), and if implied volatility is an unbiased forecast of realised volatility the hypothesis that $\alpha = 0$ and $\beta = 1$ should not be rejected.

Before presenting the empirical results, an important feature of the data should be pointed out. Since daily data is used, the frequency of the observations is shorter than the period covered by the option contracts i.e. the forecast horizon. As a consequence, forecast errors will be serially dependent since implied volatility will be forecasting actual volatility over overlapping periods. This will result in misleading inference when using standard statistical tests.

One way to solve this problem is to use so called Newey-West (1987) robust standard errors. This method generates asymptotically consistent standard errors in the case of serial correlated regression residuals. The inference for all the regressions presented

in this paper are therefore based on Newey-West standard errors and the F-tests have been calculated taking into account the overlapping data problem by using the Newey-West covariance matrix estimator.⁷

Tables 1 to 4 display the results for the four estimated exchange rates and all maturities. The results show that at-the-money implied volatility seems to be a quite good estimator of realised volatility at least for the shorter forecast horizons. For maturities between one week and three months the coefficient is significantly different from zero for all exchange rates. On the other hand, the results indicate that implied volatility is a poor indicator of realised volatility for longer horizons. It is apparent that the value of the β -coefficient tends to decrease as the time horizon gets longer. The fact that the predictive power of implied volatility tends to deteriorate is also apparent in the way the adjusted R^2 values generally tend to get smaller as the forecast horizon increases. The F-tests show that the joint hypothesis that $\alpha=0$ and $\beta=1$ is almost always rejected. The hypothesis is, however, not rejected for the SEK/USD one-month forecast horizon and for maturities between one- and six month for the AUD/USD exchange rate.

The implied volatility on AUD/USD also seems to have more predictive power than implied volatility on the other three exchange rates up to the three month forecast horizon. The β -coefficients for AUD/USD are closer to one and the adjusted R^2 values are higher compared to other currency pairs. To sum up, at-the-money implied volatility seems to have some explanatory power in predicting future volatility but it does not seem to be an unbiased predictor for most of the estimated exchange rates. These results are in line with earlier work, as for instance Galati and Tsatsaronis (1996).


An alternative to the Newey-West method of handling the overlapping data problem is to use non-overlapping data, i.e. to only use observations that are m periods apart. For the one-month forecast horizon, this means, using observations that are 20 days apart. One obvious problem with this method is that the sample will be greatly reduced. However, the sample used in this paper covers about 1500 daily observations, which should be sufficiently many for the purpose of using non-overlapping data on the shorter option contracts. Estimations using non-overlapping data have been conducted for one week, one-month, and two -months forecast horizons. The results, which can be found in appendix one, are very similar to those obtained using overlapping data. However, one difference in the results is that implied volatility seems to be an unbiased predictor of future volatility when non-overlapping data is used for some of the estimated exchange rates. Table 1a shows that implied volatility appears to be an unbiased predictor for the one-, and two-month horizons for the SEK/DEM exchange rate and for the one- month horizon for the SEK/USD exchange rate. This is however not the case for the DEM/USD exchange rate and the result of the AUD/USD are not altered. These results should be interpreted with caution since the number of observations is low. Nevertheless, the results are in line with Christensen and Prabhala who find that implied volatility is an unbiased predictor using non-overlapping data while it is biased using overlapping data.

⁷ However, as noted by Jorion (1995), since the volatility series are highly correlated processes the slope coefficient may be bias downward.

3.2 Volatility forecasts based on implied risk-neutral probability distributions

The Black-Scholes option pricing model assumes volatility to be constant across all exercise prices. But, in reality implied volatility is usually higher for *out-of-the money* and for *in-the-money* options than for *at-the-money* options creating the so-called volatility smile, which is illustrated in graph a.⁸

Graph a.



The existence of the volatility smile implies that the *at-the-money* volatility could be the lowest volatility quoted in the OTC-market. As a consequence, the *at-the-money* implied volatility may underestimate the “true” expected volatility, since the market usually prices volatility of in-and out-of-the money options higher.

The existence of the volatility smile is not only evidence that the Black-Scholes model does not hold exactly but it is also an indication that market participants make more complex assumptions about the path of the underlying asset price than what is assumed in the model⁹. As mentioned earlier, the Black-Scholes model assumes that the underlying asset price is lognormally distributed and that its return is normally distributed. Empirical research has shown that the distributions of nominal exchange rate returns appear to be leptokurtic, skewed and time-varying. The markets assessment of the probability distribution of the future exchange rate can be extracted from option prices. It is commonly assumed that options can be priced using risk-

⁸ A call option is said to be “out-of-the-money” if its exercise price is greater than the price of the underlying asset. If the exercise price is lower than the price of the underlying asset the call option is said to be “in-the-money”.

⁹ When the assumptions of the Black-Scholes model do not hold, the Black-Scholes formula might still be used, with the purpose of transforming the implied volatility quotations into option prices.

neutral valuation so that the price of e.g. a European call option is given by the present value of the option's expected future payoff:¹⁰

$$c = e^{-r(T-t)} \int_X^{\infty} q(S_T)(S_T - X) dS_T \quad (3)$$

where $q(S_T)$ is the risk-neutral density function (RND). Since expression (3) contains the risk-neutral density function, it is clear that the implied RND can be extracted by differentiating equation (3) twice with respect to the exercise price. It is therefore possible to estimate the implied RND using quotations on *at-the-money* implied volatility and two combinations of *out-of-the-money* options quoted on the OTC-market, the “*strangle*” and the “*riskreversal*”. The implied distributions can be interpreted as the market's estimate of the future probability distribution of the underlying exchange rate, if markets participants were risk neutral. If investors are risk averse, the “true” perception of the distribution will differ from the implied risk neutral distribution. However, Rubinstein (1994) shows that (under certain assumptions) the true equity price distribution for U.S. stocks shifts to the right relative to the risk neutral distribution but the general shape remains the same.


The implied RNDs in this paper are extracted using a method that relies on an assumption of the functional form for the volatility smile that was first proposed by Malz(1997)¹¹. The volatility smile is estimated by fitting a specific function to the three observed option prices. The obtained volatility function is substituted into the Black-Scholes formula, which is then differentiated twice to obtain the implied risk neutral probability density function¹². Unfortunately, this method tends to produce relatively symmetric distributions. This is due to the limited information that can be extracted from the three available option prices. Graph b shows that the extracted RND does not need to be normally distributed and may capture features such as skewness and leptokurtosis.

Graph b.

¹⁰ This assumes absence of arbitrage opportunities, and that the future cash flows of the option can be replicated. However, for some more complicated processes such as jump-diffusion processes options are not priced exclusively based on the principle of absence of arbitrage opportunities.

¹¹ See Campa, Chang and Reider (1998) for a discussion of other methods using OTC options and see Bhara(1997) for methods using exchange traded options.

¹² See appendix 2 for a brief description of the method or see the original article by Malz.



Since the RND based standard deviation takes into account the skewness and excess kurtosis it should, in principle, contain incremental information over the *at-the-money* implied volatility. The standard deviation of the RND should therefore be a better predictor than the *at-the-money* implied volatility. In this section I will evaluate whether the standard deviation derived from implied RND functions outperform the at-the-money implied volatility as a predictor of future volatility. The implied RND:s have been derived from SEK/USD, SEK/DEM, DEM/USD and AUD/USD option contracts with one month to expiration for each day of the exchange rate series. Only one-month implied distributions have been estimated since, unfortunately, this is the only maturity available in the data set for strangles and risk reversals.

Graphs 5a to 5d show that the RND based standard deviation on average is somewhat higher than the at-the-money implied volatility for the four estimated exchange rates. However, note that the difference between the RND based standard deviation and the at-the-money implied volatility is very small. One possible explanation is that the standard deviation might incorporate the skewness and excess kurtosis of the implied distribution, which under certain circumstances will result in a higher standard deviation. The kurtosis for the investigated period is always positive which indicates that the market participants seem to have expected more extreme movements in the future exchange rates than implied by the Black-Scholes option pricing model. The skewness is also always non-zero indicating that market participants have assigned a higher probability to a depreciation (appreciation) than an appreciation (depreciation) of the currencies.

To test the predictive power of the RND based standard deviation, the following OLS equation is estimated for the one-month forecast horizon:

$$RV_t = \alpha + \delta STD_t + \varepsilon_t \quad (4)$$

Where RV_t is the realised volatility at time t , and STD_t is the RND based standard deviation at time t . Table 5 presents the results of the one-month RND based standard deviation, and of the corresponding at-the-money implied volatility to facilitate a comparison. The results show that the regression coefficient δ is significantly different

from zero for all exchange rates. The values of the coefficients are very similar to the values of the coefficient of the at-the-money implied volatility. It seems that the at-the-money implied volatility and the RND based standard deviation have very similar predictive power in forecasting future volatility. However, the standard deviation is a biased estimator of future volatility for all exchange rates, while at-the-money implied volatility is unbiased for the SEK/USD and AUD/USD exchange rates.¹³

Another way to evaluate which of the uncertainty measures is the best predictor of future volatility, is to compare statistical loss functions of the two forecasts. The forecast error is defined as the difference between the forecasted and the realised volatility. The measures used for comparison are the root mean squared error (RMSE) and the heteroskedasticity adjusted mean squared error (HMSE). The reason for including HMSE is simply to take account of the heteroskedasticity in the data¹⁴.

Table 6 shows these measures for the *at-the-money* implied volatility and for the RND based standard deviation. The two measures are in general slightly larger for the forecasts based on the RND standard deviation for all currencies. Hence, the RND based standard deviation does not appear to be a better predictor of future volatility than *at-the-money* implied volatility. The regressions show that the two uncertainty measures seem to have similar predictive power, although the standard deviation is a biased estimator for all exchange rates. This is not surprising considering that graphs 5a to 5d show that the at-the-money implied volatility and the standard deviation series are extremely similar.

Thus, it does not seem to be worthwhile calculating the implied RND when the purpose is to find a good predictor of future volatility since the at-the-money implied volatility does just as well or even better. One possible explanation could be that nominal exchange rate returns might be close to being normally distributed. If this is the case, the Black-Scholes model should not be a bad approximation. Another reason

¹³ It is possible to examine the incremental predictive power of one forecast over the other using a method by Fair and Shiller (1990). Consider the following equation:

$$RV_t = \alpha + \beta IV_t + \delta STD_t + \epsilon_t$$

If neither implied volatility nor the RND based standard deviation contains any useful information for forecasting realised volatility the estimates of the coefficients β and δ should be zero. If the information in implied volatility is a subset of the information in the RND based standard deviation and if STD is an unbiased estimator, then the hypothesis that: $\alpha=0$, $\beta=0$ and $\delta=1$ should not be rejected. Since these two volatility measures are highly correlated the equation is subject to severe multicollinearity. The equation has been estimated but the results are not presented since one coefficient tend to be positive while the other tend to be of the same magnitude but negative. This regression is estimated for all volatility forecasts but are not presented in the paper since all forecast series yield severe multicollinearity.

¹⁴ $RMSE = \sqrt{\frac{1}{T} \sum_{t=1}^T (\sigma_{t+k} - \hat{\sigma}_{t,t+k})^2}$ and $HMSE = \frac{1}{T} \sum_{t=1}^T \left(\frac{\sigma_{t+k}}{\hat{\sigma}_{t,t+k}} - 1 \right)^2$ where σ_{t+k} is ex post volatility from time t to $t+k$

and $\hat{\sigma}_{t,t+k}$ is the forecast made at time t for the period up to $t+k$. Measures such as the root median square error (RMDSE) which are not sensitive to outliers are not used since there is no reason to believe that the data set contains extreme outliers. Furthermore, implied volatility contains information that is not only historical but also forward-looking, meaning that implied volatility in principle should be able to forecast sudden jumps in volatility.

could be due to the way the implied RND:s are extracted from option prices. The Malz (1997) method uses only three prices and tends to result in very symmetric distributions. It is possible that the results might have been different if the implied RND:s had been extracted using prices of exchange traded options. The next section evaluates at-the-money implied volatility (henceforth referred to as implied volatility) against forecasts based on strictly historical information.

4. The predictive power of historical based volatility measures

This section investigates the predictive power of implied volatility in forecasting future volatility in relation to forecasts based on historical information. The section begins by evaluating historical volatility, which is the simplest forecast measure. Then, volatility forecasts based on GARCH and EGARCH models are evaluated (all evaluations are performed out of sample).

4.1 The predictive power of historical volatility

For a given exchange rate series S_t historical volatility will be defined as follows:

$$HV_t = \sqrt{\frac{1}{m} \sum_{j=0}^{m-1} [\ln(S_{t+j} / S_{t+j-1}) - \mu]^2 \times 250} \quad (5)$$

where m is the number of trading days of the forecast horizon and μ is the estimated average. Historical volatility is past realised volatility since the window over which the calculation is performed is backward-looking. The reason for evaluating the predictive power of historical volatility is that it is a widely used and a very simple volatility measure. The evaluation will be performed in the same way as the evaluations in section 2. That is, the following OLS regression is estimated:

$$RV_t = \alpha + \delta HV_t + \varepsilon_t \quad (6)$$

Where α is a constant, HV_t is historical volatility at time t . The δ coefficient has to be statistically significant different from zero for historical volatility to have any predictive power for future volatility. In order for historical volatility to be an unbiased predictor of future volatility the hypothesis that $\alpha=0$ and $\delta=1$ should not be rejected.

Tables 7 to 10 show the results of the estimation. The δ -coefficients are significantly different from zero for all estimated exchange rates, except for the longer horizons. Thus, historical volatility appears to have some explanatory power for future volatility. However, a comparison with table 1 to 4 shows that the coefficient of implied volatility is much closer to one for almost all forecast horizons and exchange rates. Hence in this sense, implied volatility seems to be a better predictor of future volatility than historical volatility. On the other hand, the adjusted R^2 values do not entirely verify this

story. Some of the adjusted R^2 values are actually lower for the regressions with implied volatility as the explanatory variable. The F-test indicates, as expected, that historical volatility is a biased predictor of future volatility for all forecast horizons and exchange rates, with the exception of the AUD/USD exchange rate, for a few horizons.

As was mentioned in section 3, a way of evaluating the forecasts against each other is to compare different statistical loss functions. Table 11a to 11f show the values of the RMSE and HMSE for implied and historical volatility. In most cases, the values of these measures are higher for historical volatility than for implied volatility. One exception is the AUD/USD exchange rate for which the two loss functions are lower for historical volatility for the 6 month and 1 year horizons. Considering that neither implied nor historical volatility appears to be good predictors of future volatility for longer horizons this result should be interpreted cautiously. Hence, implied volatility seems to be a better predictor of future volatility than historical volatility for forecast horizons up to three months. In general, for longer forecast horizons both implied and historical volatility have very little, if any, explanatory power for future volatility.

The next section investigates if implied volatility still is a superior predictor of future volatility when comparing it with more sophisticated historical measures of volatility.

4.2 The predictive power of GARCH models

The Black-Scholes model assumes constant volatility during the life of the option. Graphs 6a to 6d show daily changes in the natural logarithm, i.e. returns, of the SEK/DEM, SEK/USD, DEM/USD, and AUD/USD exchange rates. The return series show that the mean appears to be constant while the variance changes over time. In other words, there seems to be volatility clustering, that is, periods of turbulence tend to be followed by periods of turbulence and calm periods by calm periods. Thus, volatility is apparently not constant over time as is assumed in the Black-Scholes model. Because of the clustering phenomena recent movements should be more useful in predicting future volatility than movements far in the past.

During the last two decades a family of autoregressive conditional heteroscedasticity models, so called ARCH models, have been developed to model time varying variance. The ARCH model was first presented by Engle (1982) and later generalised by Bollerslev(1986). The ARCH model captures the clustering effects of volatility by assigning greater weights to recent data. The model postulates that the conditional variance of the returns is a linear function of past squared innovations. To explain this further assume that the dependent variable y_t is generated by:

$$y_t = \beta X_t + \varepsilon_t \quad (7)$$

where X_t is a vector of exogenous variables and ε_t is assumed to be normally distributed with mean zero and variance h_t conditional on the information set available at time $t-1$:

$$\varepsilon_t | \Psi_{t-1} \sim N(0, h_t) \quad (8)$$

$$\Psi_{t-1} = \{y_{t-1}, x_{t-1}, y_{t-2}, x_{t-2}, \dots\}$$

In an ARCH(q) model the conditional variance is generated by:

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \dots + \alpha_q \varepsilon_{t-q}^2 ; \quad \alpha_i \geq 0 \quad \forall i \quad (9)$$

The conditional variance is an increasing function of the magnitude of the lagged errors, irrespective of their signs. Large (small) errors of either sign tend to be followed by large (small) errors of either sign. Thus, the conditional variance function captures the clustering effects of shocks. The order of the lag q determines the length of time that shocks persist in the conditional variance.

Bollerslev (1986) extended and generalised the ARCH model by including past conditional variances in the conditional variance function, thereby introducing long memory into the volatility process. He proposed that the conditional variance in a GARCH(p,q) should be specified as:

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \dots + \alpha_q \varepsilon_{t-q}^2 + \beta_1 h_{t-1} + \dots + \beta_p h_{t-p} \quad (10)$$

It has been empirically demonstrated that the GARCH(1,1), (i.e. p=1 and q=1), often is sufficient to capture the dynamics of the volatility process in most financial series¹⁵. Nevertheless, I test this by estimating a GARCH(1,1) model for the four exchange rates. Table 12 contains the results of the GARCH(1,1) estimations and the corresponding Ljung-Box tests which show that the hypothesis of no autocorrelation in the squared standardised residuals cannot be rejected, at the 5%-level, for any of the estimated exchange rates¹⁶. Thus, a GARCH(1,1) model seems to be enough for capturing the dynamics of the volatility of the exchange rates. Hence, a GARCH (1,1) model is used henceforth to generate volatility forecasts. The GARCH(1,1) conditional variance is given by:

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta h_{t-1} \quad (11)$$

In order to compare the empirical properties of the GARCH(1,1) forecasts the model is applied to daily returns on the SEK/DEM, SEK/USD, DEM/USD and AUD/USD exchange rate. The full sample runs from January 1993 to October 1998. For the first set of out-of-sample forecasts an estimation period of two years is used, starting 1993-01-03. The model is estimated with data starting in 1993 since the Riksbank let go of the currency peg in November 1992. It would be inappropriate to use an estimation period that includes observations of the fixed exchange rate regime to generate volatility forecasts in a floating exchange rate environment. Once the GARCH(1,1) parameters have been estimated using data up to time τ the variance forecasts at time τ are generated for $\tau+k$ by:

¹⁵ See Bollerslev et al (1992).

¹⁶ Table 10 also shows that there appears to be no autocorrelation in the level either.

$$E_{\tau}[\sigma_{\tau+k}^2] = \alpha_0 \frac{1 - (\alpha_1 + \beta)^{(k-1)}}{1 - (\alpha_1 + \beta)} + (\alpha_1 + \beta)^{(k-1)} \sigma_{\tau+1}^2 \quad k \geq 2 \text{ and } (\alpha_1 + \beta) < 1 \quad (12)$$

where

$$\sigma_{\tau+k}^2 = \alpha_0 + \alpha_1 \varepsilon_{\tau}^2 + \beta \sigma_{\tau}^2$$

In this way, a volatility forecast is generated for every day one year ahead. The one week volatility forecast is then given by calculating the average of the first five daily volatility forecasts, the one-month forecast by calculating the average of the first 20 daily volatility forecasts and so on. In the next step, the estimation window for the GARCH parameters is increased by one observation and new volatility forecasts are generated for the forecast period starting on this day. This procedure is repeated until the entire series of volatility forecasts has been estimated¹⁷. Thus, the GARCH model is estimated with the historical information that is available to market participants at the time the implied volatility is observed. Graphs 7a to 7f plots the GARCH(1,1) volatility forecast and the corresponding implied volatility for the SEK/DEM exchange rate. The series are quite well correlated and no systematic pattern is apparent in the sense that there are periods when the GARCH(1,1) forecast is above implied volatility and periods when it is below.

As has been mentioned earlier GARCH volatility does only depend on the magnitude and not on the sign of the past error. However, it is possible that exchange rate volatility could also be asymmetrically affected by positive and negative shocks. Nelson (1991) developed an exponential GARCH model that allows for such asymmetric effects on the conditional volatility. The so called EGARCH model works with the logarithm of the conditional variance and an additional term allowing the conditional variance to respond asymmetrically to positive and negative unexpected changes in the exchange rate. The logarithm of the conditional variance in this model is specified as:

$$\ln(h_t) = \phi_0 + \phi_1 \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \phi_2 \left[\frac{|\varepsilon_{t-1}|}{\sqrt{h_{t-1}}} - \sqrt{\frac{2}{\pi}} \right] + \rho \ln(h_{t-1}) \quad (13)$$

No restrictions are necessary on the model's parameters to ensure non-negativity since the EGARCH model is specified for the logarithm of the conditional variance. The second term of equation (13) allows the conditional variance to respond asymmetrically to positive and negative changes in the exchange rate. Empirical evidence demonstrates that there often is a negative relation between stock returns and volatility. It is not that clear that this negative relation is also apparent in the case of exchange rate volatility. Heynen and Kat (1994) conclude that there are no apparent asymmetrical relation between unanticipated returns and exchange rate volatility. Nevertheless, it has been observed that volatility tends to be higher when the Swedish krona is depreciating than when it is appreciating, and a positive shock (i.e. an

¹⁷ Note that the volatility forecasts are not independent since the estimation windows overlap.

unexpected depreciation of the krona) should therefore have a larger effect on the conditional variance than a negative shock. When a normal distribution is assumed for the error term the volatility forecast expressed as the variance on a daily basis can be specified as follows¹⁸:

$$E_{\tau}[\sigma_{\tau+k}^2] = (\sigma_{\tau+k}^2)^{\hat{\theta}^{(k-1)}} \exp \left[\left(\hat{\phi}_0 - \hat{\phi}_2 \sqrt{\frac{2}{\pi}} \right) \frac{(1 - \hat{\theta}^{k-1})}{(1 - \hat{\theta})} + \frac{1}{2} (\hat{\phi}_1^2 + \hat{\phi}_2^2) \frac{(1 - \hat{\theta}^{2(k-1)})}{(1 - \hat{\theta}^2)} \right] \quad (14)$$

$$\times \prod_{i=0}^{k-2} \{ N[\hat{\theta}^i (\hat{\phi}_2 + \hat{\phi}_1)] \exp(\hat{\theta}^{2i} \hat{\phi}_1 \hat{\phi}_2) + N[\hat{\theta}^i (\hat{\phi}_2 - \hat{\phi}_1)] \exp(-\hat{\theta}^{2i} \hat{\phi}_1 \hat{\phi}_2) \}$$

where

$$\sigma_{\tau+1}^2 = \exp \left(\hat{\phi}_0 + \hat{\phi}_1 \frac{\varepsilon_t}{\sqrt{h_t}} + \hat{\phi}_2 \left[\frac{\varepsilon_t}{\sqrt{h_t}} - E \left[\left| \frac{\varepsilon_t}{\sqrt{h_t}} \right| \right] \right] + \hat{\theta} \ln(\sigma_{\tau}^2) \right)$$

and

$$N[a] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^a \exp \left(-\frac{z^2}{2} \right) dz$$

In the same way as for the GARCH(1,1) forecast, the parameters are first estimated for the 1993-1995 period. Then, a volatility forecast is generated for every day one year ahead. The different forecasts are then estimated by calculating the average of the number of days in each forecast horizon. Thereafter, the estimation window is increased by one observation and the procedure is repeated until the entire series is estimated. Graphs 8a to 8d show that the one-month GARCH(1,1) and the EGARCH(1,1) volatility forecasts are very similar.

GARCH models are based on backward looking information and tend to forecast tomorrow's volatility very similar to today's volatility. Implied volatility on the other hand is forward looking since it should reflect the market participants' best guess of the average volatility during the life of the option given that the Black-Scholes model holds. Implied volatility should therefore contain the expected impact of future events whereas GARCH forecasts can only respond to changes in market conditions after the fact and are not able to forecast changes in these conditions. In this sense, implied volatility should be able to provide a better forecast of future volatility than GARCH based volatility.

To evaluate the GARCH(1,1) and the EGARCH(1,1) conditional volatility forecasts I use the same method used in the previous sections by estimating the following OLS regressions:

$$RV_t = \alpha + \delta \text{GARCH}_t + \varepsilon_t \quad (15)$$

and

¹⁸ See Heynen, Kemna and Vorst (1994) and Nelson (1991).

$$RV_t = \alpha + \gamma EGARCH_t + \varepsilon_t \quad (16)$$

Where α is a constant, RV_t is realised volatility, $GARCH_t$ is the GARCH(1,1), and $EGARCH_t$ is the EGARCH(1,1) volatility forecast at time t . Since the GARCH(1,1) and EGARCH(1,1) volatility forecast are estimated from 1995-01-03, I have re-estimated regression (2) for this period in order to be able to compare the results of the GARCH(1,1) and EGARCH(1,1) forecasts with those of implied volatility. Tables 13 to 16 show the estimation results for the three regressions.

The results show that both the GARCH(1,1) and the EGARCH(1,1) conditional volatility forecasts seem to be relatively good predictors of future volatility. Both the δ and γ coefficients are quite close to one and statistically significantly different from zero at the one percent level, at least for the short forecast horizons for all the estimated exchange rates. The prediction ability of the GARCH(1,1) and the EGARCH(1,1) forecast deteriorates as the prediction horizon increases. This is especially evident in the case of the SEK/USD exchange rate. Note that the adjusted R^2 values are considerably higher for the SEK/DEM exchange rate than for the SEK/USD exchange rate.

Comparing the results of the GARCH(1,1) and the EGARCH(1,1) forecasts with those of the implied volatility it is apparent that the coefficients for the GARCH and EGARCH forecasts are closer to one than the coefficient of implied volatility for forecasting horizons up to three months in the case of the SEK/DEM and SEK/USD exchange rate. Hence, surprisingly enough, the GARCH based conditional volatility forecasts appear to be better predictors of future volatility than implied volatility for these exchange rates. The F-tests show that all three volatility forecasts appear to be unbiased predictors of future volatility for the 1 to 6-months maturities of the SEK/DEM, while they are found to be biased predictors for the SEK/USD exchange rate. These results are not in line with previous studies both in the sense that the GARCH based forecasts outperforms implied volatility as better predictors of future volatility but also because all predictors seem to be unbiased for most of the maturities of the SEK/DEM exchange rate.

The results for the other two exchange rates do not show the same pattern. In the case of the DEM/USD exchange rate, the coefficient for the implied volatility seems on average to be closer to one compared to the corresponding GARCH and EGARCH coefficients. Furthermore, the adjusted R^2 values tend to be lower for the regressions including the GARCH(1,1) and the EGARCH(1,1) forecast as explanatory variables. Thus, in this case, implied volatility seems to be a better predictor of future volatility. For the AUD/USD exchange rate on the other hand it is difficult to draw any conclusions since the results are mixed. The F-tests show that for both the DEM/USD and the AUD/USD exchange rates all forecasts are found to be unbiased predictors of future volatility for most forecast horizons.

As in the previous section I evaluate the forecast accuracy of the two methods by comparing the statistical loss functions which are shown in Tables 17a to 17f. The tables basically confirm the results of the regressions. The loss functions of the EGARCH(1,1) forecast are lower than those of both implied volatility and of the GARCH(1,1) forecasts for the SEK/DEM exchange rate. Hence, positive and negative

shocks appear to affect the SEK/DEM volatility asymmetrically and it is worthwhile taking this effect into consideration when making volatility forecasts. For the SEK/USD exchange rate, on the other hand, the loss functions of the GARCH(1,1) forecast have the lowest value. Finally, the loss functions for the implied volatility are in general lower than for the GARCH based forecasts both for the DEM/USD and AUD/USD exchange rate. The results for the DEM/USD and AUD/USD exchange rate are therefore more in line with previous studies in that implied volatility seems to outperform the GARCH based forecasts.

Concluding this section, it seems that for the DEM/USD and AUD/USD exchange rate implied volatility is the best predictor of future volatility and in some cases it is also an unbiased estimator. For the exchange rates involving the Swedish krona on the other hand, forecast models based on historical data appear to perform better in predicting future volatility than implied volatility. Further, the EGARCH volatility forecast appears to outperform the GARCH forecast for the SEK/DEM exchange rate implying that there are asymmetries in the way positive and negative shocks affect volatility. One explanation for these somewhat conflicting results could be that the options market for DEM/USD and AUD/USD might be larger, more liquid and therefore more well functioning. In this case, new information of relevance to the development of the exchange rate should be incorporated more quickly in the option prices, which could result in better volatility forecasts. However, it should be pointed out that the differences in the predictive power of these three forecasts are relatively small, hence, the results should be interpreted cautiously.

5. Conclusions

The purpose of this study is to investigate the predictive power of implied volatility for future volatility. Further, I evaluate implied volatility both against another forecast based on option prices and against forecasts based on purely historical information. Implied volatility barely outperforms the standard deviation of implied probability distributions, although these two forecasts are extremely similar. Implied volatility does also outperform forecasts based on historical data for the DEM/USD and the AUD/USD exchange rates. For the exchange rates involving the Swedish krona, on the other hand, implied volatility is outperformed by forecasts based on GARCH models. More specifically, an EGARCH(1,1) conditional volatility forecasts seems to be the best predictor of future volatility for the SEK/DEM exchange rate while a GARCH(1,1) forecast appears to be the best predictor for the SEK/USD exchange rate. Hence, in contrast to the results of Heynen and Kat (1994) there appears to be an asymmetrical relation between unexpected return and volatility at least for the SEK/DEM exchange rate. Most studies find implied volatility to be a biased predictor of future volatility. However, I find implied volatility and the GARCH based forecasts to be unbiased predictors for most horizons of the SEK/DEM, DEM/USD and AUD/USD exchange rates when the estimation period covers 1995-1998.

References

Bank of Japan, (1995), "Empirical Analysis of the Information Content of Implied Volatility", *Quarterly Bulletin*, February 1995.

Bank of International Settlements (1995), "Central Bank Survey of Foreign Exchange and Derivatives Market Activity".

Bera A. K., and M. L. Higgings, (1993), " ARCH Models: Properties, Estimation and Testing", *Journal of Economic Surveys*, Vol 7, No 4.

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, 307-327.

Bollerslev, T., R.Y. Chou, and K.F. Kroner, (1992), "ARCH Modeling in Finance", *Journal of Econometrics*, 52, 5-59.

Bahra, B., (1997), "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application", *Bank of England Working Paper* No 66.

Black, F. and M. Scholes, (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 81, 637-659.

Breeden, D.T, and R.H. Litzenberger, (1978), "Prices of State-Contingent Claims Implicit in Option Prices", *Journal of Business*, 51, 621-651.

Campa, J.M., Chang, P.H.K., and R. L. Reider, (1998), "Implied Exchange Rate Distributions: Evidence from OTC Option Market", *Journal of International Money and Finance*, 17, 117-160.

Canina, L. and S. Figlewski, (1993), "The Informational Content of Implied Volatility", *Review of Financial Studies*, Vol.6, No.3.

Christensen, B.J., and N.R. Prabhala, (1998), "The Relation Between Implied and Realized Volatility", *Journal of Financial Economics*, 50, 125-150.

Engle, R.F., (1982), "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of the United Kingdom", *Econometrica*, 50, 987-1007.

Fair, R.C. and R. J. Shiller, (1990), "Comparing Information in Forecasts from Econometric Models", *American Economic Review*, 80, 375-389.

Galati, G. and K. Tsatsaronis, (1996), "The information content of implied Volatility from Currency Options", *Bank of International Settlement*.

Garman, M.B., and S.W. Kohlhagen, (1983), "Foreign Currency Option Values", *Journal of International Money and Finance*, 2, 231-237.

Heynen, R.C., and H.M. Kat, (1994), "Volatility Prediction: A Comparison of the Stochastic Volatility, GARCH(1,1), and EGARCH(1,1) Models", *Journal of Derivatives*, 2, 50-65.

Heynen, R.C., A. Kemna, and T. Vorst, (1994), "Analysis of the Term Structure of Implied Volatility", *Journal of Financial and Quantitative Analysis*, 29, 31-56.

Jorion, P., (1995), "Predicting Volatility in the Foreign Exchange Market", *Journal of Finance*, 50, 507-528.

Lamoureux, G.G. and W.D. Lastrapes, (1993) "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities", *Review of Financial Studies*, Vol. 6, No. 2.

Malz, A.M., (1997), "Option-Implied Probability Distributions and Currency Excess Returns", *Federal Reserve Bank of New York Staff Report* # 32.

Nelson, D.D., (1991), "Conditional Heteroscedasticity in Asset Returns: A New Approach", *Econometrica*, 59, 347-370.

Newey, W.K., and K. D. West, (1987), "A Simple Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", *Econometrica*, Vol. 55, No. 3.

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