

ARBITRAGE-FREE PREDICTION OF THE IMPLIED VOLATILITY SMILE

PETROS DELLAPORTAS AND ALEKSANDAR MIJATOVIĆ

ABSTRACT. This paper gives an arbitrage-free prediction for future prices of an arbitrary co-terminal set of options with a given maturity, based on the observed time series of these option prices. The statistical analysis of such a multi-dimensional time series of option prices corresponding to n strikes (with n large, e.g. $n \geq 40$) and the same maturity, is a difficult task due to the fact that option prices at any moment in time satisfy non-linear and non-explicit no-arbitrage restrictions. Hence any n -dimensional time series model also has to satisfy these implicit restrictions at each time step, a condition that is impossible to meet since the model innovations can take arbitrary values. We solve this problem for any $n \in \mathbb{N}$ in the context of Foreign Exchange (FX) by first encoding the option prices at each time step in terms of the parameters of the corresponding risk-neutral measure and then performing the time series analysis in the parameter space. The option price predictions are obtained from the predicted risk-neutral measure by effectively integrating it against the corresponding option payoffs. The non-linear transformation between option prices and the risk-neutral parameters applied here is *not* arbitrary: it is the standard mapping used by market makers in the FX option markets (the SABR parameterisation) and is given explicitly in closed form. Our method is not restricted to the FX asset class nor does it depend on the type of parameterisation used. Statistical analysis of FX market data illustrates that our arbitrage-free predictions outperform the naive random walk forecasts, suggesting a potential for building management strategies for portfolios of derivative products, akin to the ones widely used in the underlying equity and futures markets.

Key words and phrases. Prediction of option prices in FX, Risk-neutral measure, Implied Volatility, trading strategy for options.

We thank the RBS FX options desk in London for making available the USDJPY options data used in this paper. We are grateful to Dilip Madan for suggesting this problem to us. This work is financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF), project ARISTEIA-LIKEJUMPS..

1. INTRODUCTION

One might expect that the time series of past option prices or, equivalently, their implied volatilities, could be used to predict their future behaviour. However, there is a major theoretical and practical impediment to constructing a statistical model that can be estimated on a time series of option prices,¹ namely that the state space of the model has to consist of arbitrage-free option prices only *and* the expected value of the model at any time also needs to lie within this state space. Otherwise the predictions are not arbitrage-free and are hence of little use in practice since they can certainly not be fulfilled (the contemporaneously quoted option prices in the derivatives markets are always free of arbitrage). Since the arbitrage-free requirement is both non-linear and non-explicit, it is impossible to satisfy it in a time series model for option prices, as in such a framework the model innovations are allowed to take arbitrary values thus making the model necessarily violate the no-arbitrage constraints with positive probability.

This paper provides a framework that overcomes the difficulty arising from the arbitrage-free requirement in forecasting option prices and shows that the predictions obtained are not only internally consistent, but also outperform, for the data used in our empirical study, the ones based on a model that assumes a random walk process for the option prices. The random walk (or white noise) model is a standard benchmark in the related literature for forecasting equities and it is typically very hard to beat; see, for example, Welch and Goyal [2008]. It simply forecasts the option prices at time $t + 1$ by the observed price at the time t and it is by default arbitrage-free. We consider prediction of the *volatility smile*, which is equivalent to the prediction of option prices with any strikes and a fixed maturity (see e.g. Breeden and Litzenberger [1978], Renault [2010] and Section 3). Our methodology consists of the following key idea. First, for a given maturity, encode the time series of the historical implied volatilities across strikes in terms of the parameters of a corresponding risk-neutral measure. Second, perform a time series analysis on the implied historical parameter observations. We provide empirical evidence that, at least in the Foreign Exchange (FX) derivatives market from which we have available data, the time series based on the observed parameter values exhibits predictive ability

¹Throughout the paper we assume that a manager has access to a good quality implied volatility time series data. Our data set is described in Section 4.1 below.

for implied volatilities (and hence option prices) through ARMA-GARCH based statistical models for the risk-neutral parameters. The key benefit of this approach (i.e. first performing a non-linear transformation of implied volatilities and then applying a model to the time series of the implied risk-neutral parameters) lies in the fact that, in the parameter space, the no-arbitrage constraints are reduced to ensuring that the state space of the time series model is defined by the natural parameter restrictions. As we shall see in Section 4, this is typically easy to satisfy and, in particular, implies that the forecasts also lie within the state space. Equivalently put, the predicted option prices will be free of arbitrage. In Section 4.3 a simple trading strategy for strangles and risk-reversals, based on our approach and out-of-sample options data, is presented.

2. THE PROBLEM

2.1. Prediction of implied volatilities. Our aim is to investigate the problem of the arbitrage-free forecasting of option prices in the context of the co-terminal (i.e. with the same expiry) call options on the USDJPY exchange rate. As mentioned in the introduction, call and put contracts on the exchange rates of the currencies of major economies are among the most liquid derivatives contracts in the financial markets. The FX derivatives market, where such contracts trade, is very deep, both in terms of the volume of options traded and the size of the underlying notional of the contracts. We will develop our approach using call options but, due to the put-call parity (see e.g. [Hull, 2002, Sec. 8.4]), our results are directly applicable to put options.

Recall that a call option struck at K and expiring at some future time T is a contract that gives the owner the right, but not the obligation, to purchase N \$ at expiry T for the price of K ¥ per \$. Furthermore, the price $C(K, T)$ of a vanilla call option described above is well-defined for practical purposes, because the bid-offer spread is tight due to the volume, frequency and size of the trades. The market makers in FX options keep a close eye on and record throughout a trading day the prices of all the options on a given currency pair across strikes and maturities. In practice, option traders do not refer to option prices *per se* but rather express them through the Implied Volatility (IV) metric, which allows one to compare directly option prices across different strikes, maturities, expiration times and underlyings in the same units (IV is derived using the Black-Scholes formula, see Section 3.1 for definition and Gatheral [2006] for more details). For this reason, in what

follows we may (and shall) consider the prediction of the implied volatilities $IV(K, T)$ instead of the option prices $C(K, T)$, without loss of generality and with the benefits mentioned above.

The main obstacle to a direct application of econometric tools to the multi-dimensional time series of option prices can be described as follows. To make the discussion simpler let us fix a maturity T (e.g. 1 month). Note first that there are uncountably many possible option contracts with maturity T at any moment in time as there are uncountably many possible strikes K . Assume further that we have chosen strikes K_i , $i = 1, \dots, n$ (with e.g. $n = 40$), and would like to model the evolution of the corresponding implied volatility prices $IV_t(K_i, T)$, $i = 1, \dots, n$, over time (indexed by $t \in \mathbb{N}$). Let the n -dimensional vector \mathbf{Y}_t , with co-ordinates given by

$$\mathbf{Y}_t(i) := IV_t(K_i, T), \quad i = 1, \dots, n,$$

describe the midday implied volatilities on USDJPY on day t for all the strikes K_i and assume that we are given such implied volatilities over the time period $t = 1, \dots, t_0$. A direct approach based on a general multi-dimensional time series model predicts \mathbf{Y}_{t+1} with a predictor $G(\mathbf{Y}_{1:t})$ based on a general function G , by minimizing the squared error loss $E(\mathbf{Y}_{t+1} - G(\mathbf{Y}_{1:t}))^2$, or, equivalently, by setting $G(\mathbf{Y}_{1:t})$ as the conditional mean $E(\mathbf{Y}_{t+1} | \mathbf{Y}_{1:t})$. Since the above minimum mean-squared error predictor of \mathbf{Y}_{t+1} given $\mathbf{Y}_{1:t}$ presupposes knowledge of the joint distribution of $\mathbf{Y}_{1:t+1}$, which is not available, we assume that it belongs to a parametric family of distributions indexed by a parameter vector θ . Then, the optimality of the above predictor cannot be obtained since for every θ we shall have a different optimal predictor $E(\mathbf{Y}_{t+1} | \mathbf{Y}_{1:t}, \theta)$. However, a simple way to proceed that gives reasonable predictions is to replace the unknown parameter θ with a suitable estimate based on past history $\mathbf{Y}_{1:t}$. What makes this approach infeasible to our problem of predicting implied volatilities is the fact that at the time step $t + 1$, the components of the vector \mathbf{Y}_{t+1} , i.e. the modelled implied volatilities $IV_{t+1}(K_i, T)$, $i = 1, \dots, n$, have to satisfy non-linear and non-explicit no-arbitrage restrictions. It is an extremely hard problem to identify an estimate of θ that will result in a state vector \mathbf{Y}_{t+1} with the desired no-arbitrage restrictions.²

²It follows from Breeden and Litzenberger [1978] that, when translated back to the call option prices $C(K, T)$, the no-arbitrage restrictions amount to the function $K \mapsto C(K, T)$ being non-increasing and


Before describing the solution to this problem, we should stress here that, due to the liquidity in the FX options market, the recorded implied volatilities always satisfy the no-arbitrage requirements, and therefore it is hard to perform the statistical analysis of the options data set. This fact by no means implies that in a non-efficient market the statistical analysis of implied volatilities would be easy. On the contrary, if the options market is not efficient enough to have contemporaneous implied volatility quotes across a family of strikes for a given maturity, the problem is not well-defined because there is no internally consistent time series of implied volatilities and a more sophisticated approach for the estimation of the pricing kernel (see e.g. Gagliardini, Gourieroux, and Renault [2011]) is required before any kind of statistical analysis can begin.

2.2. A solution. The above discussion indicates that the arbitrage-free prediction of implied volatilities is a difficult problem, a solution to which has, to our knowledge, not yet been proposed. We propose a solution below, which will be shown consistently to outperform random walk predictions in a large out-of-sample exercise (see Figure 1 in Section 4.2 for the numerical results).

In order to predict the implied volatility vector \mathbf{Y}_{t+1} , we first apply a non-linear transformation to the implied volatilities data up to and including time t . This one-to-one transformation maps the implied volatilities into the parameters of a risk-neutral measure \mathbb{Q} (see Section 3 for the definition and properties of the parametric form of the risk-neutral measure used in FX options markets). The core of our approach rests on the fact that every element in the space of \mathbb{Q} parameters yields an arbitrage-free set of implied volatilities across all strikes. This allows us to perform the statistical analysis in the space of \mathbb{Q} -parameters. We use the daily recorded parameters that characterise the choice of the risk-neutral measure given by the quoted implied volatilities on that day, to forecast, via a time series model, tomorrow's \mathbb{Q} -parameter values (see Section 4.1 for the precise description of our data set). The predicted implied volatilities $\hat{\mathbf{Y}}_{t+1}$ are then computed via a non-linear pricing function that guarantees that the no-arbitrage restrictions are satisfied.

convex. However, imposing such restrictions on the stochastic dynamics of $C_t(K_i, T)$, $i = 1, \dots, n$, at each moment of the running time t constitutes a problem of equivalent complexity.

A schematic description of our approach to this problem is given in the following diagram:



In the mathematical finance literature, arbitrage-free evolution of option prices under a risk-neutral measure has been studied extensively (see e.g. Carmona and Nadtochiy [2009], Kallsen and Krühner [2010], Schweizer and Wissel [2008] and the references therein). While these approaches are both mathematically involved and very interesting, we should point out that the problem of the arbitrage-free evolution of the implied volatility smile/surface under the risk-neutral measure is distinct from the main aim of the present paper. In this paper we model the evolution of the implied volatility smile under the real-world measure and are solely concerned with the arbitrage-free constraint in a static sense, i.e. at each moment of the running time,³ in order to obtain viable predictions for tomorrow's option prices. Furthermore, note that the stochastic dynamics under the continuous-time models, like the ones in Carmona and Nadtochiy [2009], Kallsen and Krühner [2010], Schweizer and Wissel [2008], is typically very complicated making it unclear how and whether such models can be efficiently estimated and applied for the statistical prediction of future option prices.

3. ARBITRAGE-FREE OPTION PRICES IN FOREIGN EXCHANGE

Intuitively, a set of implied volatilities $IV_t(K_i, T)$, $i = 1, \dots, n$, allows arbitrage if and only if it is possible to trade (i.e. buy and sell) the corresponding contracts, at prices given by $IV_t(K_i, T)$, $i = 1, \dots, n$, and follow a self-financing trading strategy on the FX spot rate (i.e. buy and sell the underlying currencies) in such a way that at some future time the portfolio of options and the gains (i.e. P&L) from the trading strategy in agregat have a non-negative value with certainty and a strictly positive value with positive probability. In other words, if a market maker (i.e. a large trading desk whose business it is to buy or sell options of any strike and maturity) quotes option prices that allow arbitrage, a counterparty could enter into a contract with the market maker *without any risk of loss*

³This is completely analogous to the way pricing theory is applied in the derivatives markets.

(here we are allowed to ignore the fact that prices are quoted in terms of bids and offers, because, as mentioned in Section 2, the spread in USDJPY options is typically only a few basis points). It is evident that, when the number of option prices is large (it is in the thousands or even tens of thousands, in the case of market makers) it is very hard to check *a priori* from the quoted option prices whether they are arbitrage-free. In practice this is achieved by using the theoretical concept of the *risk-neutral* measure, which arises in the fundamental theorem of asset pricing: the implied volatility prices $IV_t(K_i, T)$, $i = 1, \dots, n$, satisfy a no-arbitrage restriction if they are obtained as discounted expectations of their respective payoffs under a risk-neutral measure \mathbb{Q} . In other words, if they are of the form

$$e^{-(r_d - r_f)T} \mathbb{E}^{\mathbb{Q}} [(S_T - K_i)^+],$$

where S_T denotes the FX rate (e.g. USDJPY) at the future time T , r_d (resp. r_f) represent the dollar (resp. yen) interest rate over the time interval of length T and, $x^+ = \max\{0, x\}$ for any $x \in \mathbb{R}$. The defining feature of a risk-neutral measure \mathbb{Q} is that the mean of the random FX rate S_T at time T is equal to the exponential of the interest rate differentials of the two currencies over the time period from now until expiry T :

$$(1) \quad \mathbb{E}^{\mathbb{Q}}[S_T] = F_T, \quad \text{where} \quad F_T = e^{(r_d - r_f)T} S_0.$$

The condition in (1) is clearly satisfied by many probability measures \mathbb{Q} , as it only fixes the mean of the random variable S_T . This gives the market makers (and other market participants) a large amount of freedom in choosing the option prices in a no-arbitrage way as there clearly exists a vast family of probability measures, and hence risk-neutral models for the FX rate S_T , such that (1) holds.

The following two observations play an important role in how the concept of the risk-neutral measure is applied in the financial markets: (a) The choice of a risk-neutral measure does not necessarily give the correct (i.e. market) price for a given derivative contract. All such a choice does is to provide arbitrage-free prices for all traded options with the same maturity. Note that, since the market maker is interested in the consistent pricing of derivatives, this is precisely what they are after. (b) The correct price of an option, or equivalently the correct choice of the risk-neutral measure, is down to supply and demand in the market and crucially depends on the market makers view of where the specific contracts should trade. This is why flexible parametric forms of the risk-neutral measure

are useful: they allow the market maker to express their view of where the liquid derivatives should trade, while providing an arbitrage-free pricing mechanism for all derivatives.

As it turns out, in the FX option markets there is a standard parametric form for the choice of the risk-neutral measure, which we will describe in detail in Sections 3.1 and 3.2. This not only allows us to circumvent the difficulty arising from the no-arbitrage restriction for option prices, but it expresses the vanilla option for any strike and a given maturity in terms of the values of three parameters, the current level of the FX spot rate and the two interest rates in the respective currencies over the period from current time until maturity. There are two reasons why the particular parameterisation of option prices, based on the SABR formula (see Section 3.2) has been adopted as market standard:

- (i) the parsimonious description of all option prices with a given maturity is very robust and easy to use for the traders;
- (ii) each of the three parameter values is closely related to the three fundamental features of the risk-neutral distribution (which determines option prices): the over-all level of option prices and the skewness and kurtosis of the risk-neutral law of the FX spot rate S_T .

It should be stressed here that the analytical SABR formula (see Equation (4) below) in general yields arbitrage-free option prices except in the extreme wings (i.e. for the strikes that are many standard deviations away from the at-the-money value). However, the points made above and the fact that for such extreme strikes the notion of the market defined price is questionable due to the lack of liquidity,⁴ suggest that it is as reasonable to apply the SABR parametrisation of the risk-neutral measure for purposes of our problem as it is to use it for the pricing, hedging and risk management of the portfolios of options, a very common day-to-day practice in the FX options markets. A further argument in favour of using the SABR formula, besides the lack of liquidity in the wings, is that the price of options in the extreme wings is more uncertain than in the “near” wings due to the increase in the bid-offer spread for such derivatives.

3.1. Implied volatility in the FX markets. The value $BS(F_T, K, T, \sigma)$ of the European call option with strike K and expiry T in a Black-Scholes model with constant volatility

⁴We thank the referee for this observation.

$\sigma > 0$ is given by the Black-Scholes formula

$$(2) \quad \text{BS}(F_T, K, T, \sigma) := e^{-r_d T} [F_T N(d_+) - K N(d_-)],$$

where

$$d_{\pm} = \frac{\log(F_T/K) \pm \sigma^2 T/2}{\sigma \sqrt{T}},$$

and $N(\cdot)$ is the standard normal cumulative distribution function. The implied volatility that corresponds to the market price $C(K, T)$ for the strike $K > 0$ and maturity $T > 0$ is the unique positive number $IV(K, T)$ that satisfies the following equation in the variable σ :

$$(3) \quad \text{BS}(F_T, K, T, \sigma) = C(K, T).$$

Implied volatility is well-defined since the function $\sigma \mapsto \text{BS}(F_T, K, T, \sigma)$ is strictly increasing for positive σ (the *vega* of a call option $\frac{\partial \text{BS}}{\partial \sigma}(F_T, K, T, \sigma) = e^{-r_d T} F_T N'(d_+) \sqrt{T}$ is clearly strictly positive) and the right-hand side of (3) lies in the image of the Black-Scholes formula if and only if the call option price $C(K, T)$ satisfies a no-arbitrage restriction. It is clear from the definition, albeit suppressed in the notation, that the implied volatility $IV(K, T)$ also depends on the current level of the FX spot rate S_0 and the interest rate differential between the two currencies.

As noted earlier, the implied volatility $IV(K, T)$ is nothing more than a convenient number (in %) to express the price of a call option with expiry T and strike K via the Black-Scholes formula in (2). The convenience, from the perspective of market participants, lies in the fact that $IV(K, T)$ contains information about the price of a call option, which is independent of the FX rate and can therefore be easily compared across currency pairs. Furthermore, implied volatility is being used in the FX option markets to specify parametrically the call option prices for all strikes and a given maturity. It is well-known (see e.g. Breeden and Litzenberger [1978]) that the information contained in knowing the prices of the implied volatilities for all strikes and a given maturity T is equivalent to specifying the risk-neutral law of the FX rate S_T . Therefore, the SABR formula for the implied volatility $IV(K, T)$ given in the next section together with the Black-Scholes formula in (2) specify an easy parameterisation of the risk-neutral law of the spot rate S_T .

3.2. The parametric form of the risk-neutral measure in the FX markets. The version of the SABR formula for the implied volatility, derived in Hagan, Kumar, Lesniewski, and Woodward [2002], which will be used in this paper takes the form

$$(4) \quad IV(K, T) = \alpha \frac{1 + T \left[\frac{1}{4 \cdot 24} \frac{\alpha^2}{(F_T K)^{1/2}} + \frac{1}{4} \frac{\rho \nu \alpha / 2}{(F_T K)^{1/4}} + \frac{2 - 3\rho^2}{24} \nu^2 \right]}{(F_T K)^{1/4} \left\{ 1 + \frac{1}{4 \cdot 24} \log^2 \left(\frac{F_T}{K} \right) + \frac{1}{16 \cdot 1920} \log^4 \left(\frac{F_T}{K} \right) \right\}} \cdot \frac{z}{x(z)}$$

where

$$z = \frac{\nu}{\alpha} (F_T K)^{1/4} \log \left(\frac{F_T}{K} \right) \quad \text{and} \quad x(z) = \log \left\{ \frac{\sqrt{1 - 2\rho z + z^2} + z - \rho}{1 - \rho} \right\}.$$

The market data in this formula consist of the current FX spot rate S_0 and the interest rate differential for the maturity T , which features in F_T (see formula in (1)). The parameters that allow the market participant to express their view on the price of the call option struck at K with maturity T are given by

α	instantaneous (current) level of volatility,
ν	volatility of volatility,
ρ	instantaneous correlation.

A general version of the SABR formula contains a further parameter β the value of which is taken to be 1/2 in (4). In practice $\beta = 1$ is also used, marginally changing the analytical expression in (4). However, this makes no material difference for pricing purposes as a change in β can be compensated by a change in the value of ρ since both parameters mainly effect the skew of the smile (cf. the final paragraph of this subsection).⁵

The formula in (4) yields a natural parameterisation of the risk-neutral law of S_T . It was developed in Hagan et al. [2002] based on an assumption that the FX spot rate process $(S_t)_{t \geq 0}$ evolves under a risk-neutral measure as a stochastic volatility process: α denotes the value of the volatility process in this model at time zero (i.e. the current value), ρ is the correlation between the two Brownian motions driving the spot and the volatility processes, and ν is the volatility of the stochastic volatility process; see Hagan et al. [2002] for more details.

⁵We thank one of the referees for this remark.

A simple and yet important fact is that for each value of the state-vector (α, ν, ρ) , where α and ν are positive and ρ is between -1 and 1 , the function $K \mapsto \text{BS}(F_T, K, T, IV(K, T))$, given by the formulae (2) and (4), represents an arbitrage-free collection of option prices for any positive strike K . The parameters (α, ν, ρ) are used to control the risk-neutral distribution in the following way: a change in the parameter α has the effect of changing the overall level of option prices, the parameter ν and ρ control the kurtosis and skewness of the risk-neutral distribution of S_T . This natural interpretation of the parameters has made (4) a standard parameterisation of a risk-neutral law for the spot in the FX option markets.

It should be stressed, however, that our approach does not depend on the specific parametric form of a risk-neutral law. We use the time series of parameters of formula (4), because they describe the market implied risk-neutral law of the FX rate S_T extremely accurately: since the formula has emerged as the market standard, the values of the parameters (α, ν, ρ) are used by the traders to express the market consensus on what shape the risk-neutral law of S_T takes. In particular, the parameter β mentioned in the first paragraph of this subsection could have taken the value $\beta = 1$ (instead of $\beta = 1/2$) without ramifications for our approach. More generally, in a different market the method outlined here can be applied to the parameters of the risk-neutral law given by the standard of that particular market. For example in the equity option markets one can use Heston's parameterisation Heston [1993] of the risk-neutral law.

4. ARBITRAGE-FREE MODELLING OF CO-TERMINAL OPTIONS IN FOREIGN EXCHANGE

4.1. Description of the data set and FX option market conventions. In foreign exchange the liquid options of a given market-defined maturity T , e.g. $T = 1$ month, have a rolling expiry with respect to the current time t . In other words, at time t a liquid call option expiring in one month would cover the time interval $[t, t + T]$. The next day, i.e. at time $t + 1$, a liquid one-month option will be a different derivative contract covering the time interval $[t + 1, t + 1 + T]$.

Our data set consists of the time series for the parameter values (α, ν, ρ) implied by the USDJPY liquid one-month options, the FX spot rate level and the interest rate differentials at London midday over the period 29/9/2006 to 16/12/2011. The data set has

been provided by the RBS FX options desk, one of the largest market makers in FX options. The option quotes are taken at London midday to ensure that all the option prices are temporally consistent. In other words, if the option prices that are used to obtain the parameters (α, ν, ρ) are not recorded simultaneously, there is no guarantee that the parameter triplet reflects the market view on the risk-neutral law of the FX spot rate. Since London midday is the only time at which the trading desk records all the liquid option prices simultaneously, the parameters (α, ν, ρ) obtained in this way describe all the option prices with expiry T without the usual difficulty with options data arising from the illiquidity in the market. Furthermore, this also gives us a very clear interpretation for the predicted parameters: our prediction yields a risk-neutral law of S_T that can be converted to a vector of option prices and compared (out-of-sample) to the option prices that are observed at future London midday fixing times.

4.2. Time series modelling. We have obtained the data set $(\alpha_t, \nu_t, \rho_t)$ for $t = 1, \dots, 1360$ for the USDJPY exchange rate. Our modelling construction is executed as follows: we use the first $n_0 = 1000$ values to fit a model and the next $n_1 = 360$ values to test its predictive ability. We first transform our data using

$$\begin{aligned} A_t &:= \log(\alpha_t) \\ N_t &:= \log(\nu_t) \\ R_t &:= \log\left(\frac{\rho_t + 1}{1 - \rho_t}\right) \end{aligned}$$

so that all transformed parameters lie in the real line. This will guarantee that the predictions achieved through the time series models obey the restrictions $\alpha > 0, \nu > 0$ and $-1 < \rho < 1$. Model building has been guided through an ARMA-GARCH time-series methodology which includes autocorrelation and partial autocorrelation plots, AIC and BIC information criteria and parameter significance tests; see, for example, Brockwell [2005]. In particular, in all series there was overwhelming evidence that the hypothesis of unit root cannot be rejected (with augmented Dickey-Fuller and Phillips-Peron tests) so the difference operator was applied to all series and in the resulting series the unit root hypothesis was rejected. All (differenced) series exhibited strong GARCH effects that were evident by both inspecting the autocorrelation plots of their squares and by Engle's ARCH test. After removing the GARCH effects (in which student-t errors seemed to

perform much better than Normal errors when BIC and AIC values were compared and when residuals were inspected) the autocorrelation and partial autocorrelation plots of the residual series indicated that forms of ARMA-GARCH models might be appropriate. Our subsequent modelling choice procedure consisted of fitting a series of such models and finally the ‘best’ models have been chosen with respect to the best AIC/BIC values attained. In some ARMA-GARCH models the problem of near root cancellation that results in misleading inference (see, for example, Ansley and Newbold [1980]) was observed so we chose to remove the moving average component. Parameter estimates have been obtained through MATLAB software. A mathematical formulation of our models is as follows.

By first defining, for $t = 2, \dots, n_0$, $\Delta\alpha_t := A_t - A_{t-1}$, $\Delta\nu_t := N_t - N_{t-1}$, and $\Delta\rho_t := R_t - R_{t-1}$, we obtained as best models the following general specifications:

$$\begin{aligned}\Delta\alpha_t &= \mu^\alpha + \sum_{i=1}^{p^\alpha} \phi_i^\alpha \Delta\alpha_{t-i} + \sum_{j=1}^{q^\alpha} \theta_j^\alpha \epsilon_{t-j}^\alpha + \epsilon_t^\alpha \\ \Delta\nu_t &= \mu^\nu + \sum_{i=1}^{p^\nu} \phi_i^\nu \Delta\nu_{t-i} + \sum_{j=1}^{q^\nu} \theta_j^\nu \epsilon_{t-j}^\nu + \epsilon_t^\nu \\ \Delta\rho_t &= \mu^\rho + \sum_{i=1}^{p^\rho} \phi_i^\rho \Delta\rho_{t-i} + \sum_{j=1}^{q^\rho} \theta_j^\rho \epsilon_{t-j}^\rho + \epsilon_t^\rho\end{aligned}$$

where the error terms $\epsilon_t^\alpha, \epsilon_t^\nu, \epsilon_t^\rho$ follow a GARCH(1,1) model with Student-t errors. The resulting estimated parameters, based on the initial time series of 1000 data points, are given in Table 1.

The out-of-sample prediction exercise was performed by fitting the model of Table 1 as each new data point $t = 1001, 1002, \dots, 1360$ arrives and then predicting the triplet of the parameters through the fitted model for one, two and three days ahead. These predicted triplets were then transformed back to predicted implied volatilities via (4). The three ingredients that were unknown in (4), namely the future spot rate S_0 and the two interest rates r_d, r_f (yielding the future forward rate F_T , see (1)), have been predicted with a simple random walk model, so were taken to be equal to current values. Strikes K were selected to be forty equally spaced values between 90% and 110% of the current spot rate. For every future day and every strike we have calculated the predicted implied volatility and recorded the absolute error against the implied volatility derived using the true values

Parameter estimates for the USDJPY series					
μ^α	-0.0002 (0.0001)				
ϕ^α	0.9104 (0.0160)				
θ^α	-0.9789 (0.0053)				
ϵ^σ	0.0002 (0.0001)	0.1801 (0.0444)	0.7807 (0.0413)	3.8903 (0.5788)	
ϕ^ν	-0.1844 (0.1082)	-0.2279 (0.0372)	-0.2269 (0.0441)	-0.1096 (0.0428)	0.2753 (0.0341)
θ^ν	0.0042 (0.1105)				
ϵ^ν	0.0000 (0.0000)	0.0317 (0.0145)	0.9397 (0.0311)	6.1987 (1.3818)	
μ^ρ	0.0013 (0.0022)				
ϕ^ρ	-0.0520 (0.0273)				
ϵ^ρ	0.0004 (0.0003)	0.0445 (0.0249)	0.9262 (0.0444)	2.9694 (0.4177)	

TABLE 1. Estimated parameters (standard errors in brackets) for the time series models. The four ϵ parameters correspond to the GARCH(1,1) model estimates: the constant, the ARCH parameter, the GARCH parameter and the degrees of freedom of the t-density respectively.

of the parameters α, ν, ρ and (of course) the realised values of the spot rate S_0 and the interest rates r_d, r_f . Figure 1 illustrates that the mean error varies between 28 basis points for predictions one day ahead and may even reach 65 basis points for predictions three days ahead.

In Figure 1 we have depicted (with blue dots) the corresponding errors produced by the naive random walk predictor (i.e. tomorrow's option prices are the same as today's). Our methodology clearly outperforms the random walk forecasts for all strikes with the improvement reaching up to four basis points for out-of-the-money options. Beyond the three day time horizon, this improvement vanishes, indicating that our ARMA-GARCH model does not have further predictive ability.




FIGURE 1. USDJPY. Red cross: our predictions; blue dots: random walk predictions.

An important observation is that the predicted risk-neutral measure in our framework depends, not only on the SABR parameters, which our methodology predicts, but also on the future level of the FX forward rate. Since the latter is notoriously difficult to predict, in our approach we use the random walk prediction for it. Recall that the level of the forward rate is given by the appropriately scaled mean of the FX rate under the risk-neutral measure, see (1). Figure 1 clearly shows that our forecasts outperform the random walk by a larger amount for options that are struck further out-of-the-money (i.e. their strike is far from the FX forward rate). The shape of the payoff of such options dictates that their price depends to a greater extent on the tail and less so on the mean of the predicted risk-neutral distribution. Furthermore, it is interesting to note that, as clearly shown by Figure 1, historical option price data contain sufficient information to improve the prediction of future option prices compared to a random walk forecast even though the future option prices depend on the future FX forward rate, which is very hard to forecast. Perhaps this offers an avenue for the improved prediction of the FX rate itself.

4.3. Simple trading strategy for strangles and risk-reversals. Based on the results of the previous section, we perform an illustrative trading exercise to demonstrate the potential of our methodology. We choose to buy or sell strangles ($P(K_-, T) + C(K_+, T)$)

and risk-reversals ($C(K_+, T) - P(K_-, T)$) at strikes $K_- = 0.9S$ and $K_+ = 1.1S$, where S is the current level of spot. We emphasize here that this choice of K_{\pm} is arbitrary and that structures consisting of any finite combinations of options with maturity T may be traded, since we have arbitrage-free forecasts for the entire implied volatility smile at T . A realistic strategy requires a trading rule in which trades are executed only when the return forecast of a strangle/risk-reversal exceeds a given value of $\delta > 0$. Put differently we go long a strangle if the current $(P(K_-, T), C(K_+, T))$ and predicted $(\tilde{P}(K_-, T), \tilde{C}(K_+, T))$ prices satisfy

$$\min \left\{ \frac{\tilde{C}(K_+, T) - C(K_+, T)}{C(K_+, T)}, \frac{\tilde{P}(K_-, T) - P(K_-, T)}{P(K_-, T)} \right\} > \delta,$$

and act analogously in the case of a risk-reversal. The realised P&L of the trade is the next day's price of the structure minus $P(K_-, T) + C(K_+, T)$ (if we were long a strangle on the previous day).

Figure 2 depicts the out-of-sample performance of this strategy for the data used in Section 4.2. The results indicate that there is consistent positive return expressed by both the average daily standardized return and the average daily return of the strategy. Notwithstanding that the P&L reported in Figure 2 is gross of transaction costs, we believe that there is potential to adopt the arbitrage-free option price forecasting methodology described in this paper both for risk management purposes and as a profitable quantitative trading strategy. It is worth emphasising in this connection that both the trading strategy and the forecasting model we have used were chosen for their simplicity, rather than their ability to capture the full potential of the approach.

REFERENCES

- C.F. Ansley and P. Newbold. Finite sample properties of estimators for autoregressive moving average models. *Journal of Econometrics*, 13(2):159–183, 1980.
- D.T. Breeden and R.H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978.
- P.J. Brockwell. *Time Series Analysis*. Wiley Online Library, 2005.
- R. Carmona and S. Nadtochiy. Local volatility dynamic models. *Finance & Stochastics*, 13:1–48, 2009.




FIGURE 2. Performance of an out-of-sample trading strategy. Dashed green line: average daily return; Solid blue line: average daily standardized return, i.e. average daily return divided by the standard deviation of the return. It should be noted that for $\delta \sim 10^{-5}$, the strategy trades on 82% of days (i.e. 281 out of 342) and for $\delta \sim 5 * 10^{-4}$, it trades on 3.2% of days (i.e. 11 out of 342). We could of course increase the trading frequency for any fixed threshold δ by exploiting all possible structures of vanilla options (expiring at T) to obtain a stronger signal. This is possible because we have at our disposal all arbitrage-free option prices for maturity T .

- P. Gagliardini, C. Gouriéroux, and E. Renault. Efficient derivative pricing by the extended method of moments. *Econometrica*, 79(4):1181–1232, July 2011.
- J. Gatheral. *The volatility surface: a practitioner's guide*. John Wiley & Sons, Inc., 2006.
- P. Hagan, D. Kumar, A. Lesniewski, and D. Woodward. Managing smile risk. *Willmott Magazine*, pages 84–108, 2002.
- S.L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.
- J.C. Hull. *Options, Futures and Other Derivatives*. Prentice Hall, New Jersey, 5th edition, 2002.

- J. Kallsen and P. Krühner. On a Heath-Jarrow-Morton approach for stock options. *Preprint*, 2010. URL <http://www.math.uni-kiel.de/numerik/kallsen/personen/kallsen/pub/HJM.pdf>.
- E. Renault. Econometrics of option pricing. *Encyclopedia of Quantitative Finance*, 2010.
- M. Schweizer and J. Wissel. Arbitrage-free market models for option prices: The multi-strike case. *Finance & Stochastics*, 12:469–505, 2008.
- I. Welch and A. Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.

DEPARTMENT OF STATISTICS, ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS, ATHENS, GREECE
E-mail address: `petros@aueb.gr`

DEPARTMENT OF MATHEMATICS, IMPERIAL COLLEGE LONDON, UK
E-mail address: `a.mijatovic@imperial.ac.uk`