Predicting the dynamics of implied volatility surfaces: A new approach with evidence from OTC currency options*

Georgios Chalamandaris† Andrianos E. Tsekrekos‡

This version: January, 2008

Abstract

Volatility implied from observed option contracts systematically varies with the contracts strike price and time to expiration, giving rise to an instantaneously non–flat implied volatility surface (IVS) that exhibits substantial time variation. We propose a new approach that jointly models the cross–sectional characteristics and the time–series dynamics of the IVS. First, instead of imposing a parametric specification of moneyness and time–to–maturity to explain the IVS cross–sectionally, we derive directly from the data a number of orthogonal statistical factors that are shown to accurately reproduce the IVS observed on any given day. These statistical factors are shown to have a natural interpretation in the law of motion of the IVS. At a second stage, we attempt to exploit the factors identified for forecasting purposes, by modeling their evolution with simple, parsimonious econometric specifications. We demonstrate that our approach achieves a high–quality fit of the surface and of its evolution over time, using OTC currency options. The out–of–sample forecasting accuracy of the approach up to three days in the future is found to be significantly higher than that of hard–to–beat widely–used benchmarks.

*This paper is an updated version of a draft, entitled "*Common factors and causality in the dynamics of implied volatility surfaces: Evidence from OTC currency options*", and has benefited from the comments and suggestions of participants at the 6$^{\text{th}}$ Hellenic Finance & Accounting Association meeting in Patras and at the Accounting & Finance seminars series in AUEB. Any errors are our own responsibility.

†Corresponding author, Department of Accounting & Finance, Athens University of Economics & Business (AUEB), 76 Patision Str., 104 43, Athens, Greece. Tel:+30-210-8203758, email: gchalamandar@aueb.gr.

‡Tel/Fax:+30-210-8203762, email: tsekrekos@aueb.gr.

1 Introduction

Observed option prices implicitly contain information about the volatility expectations of market participants. Using an option pricing model, these volatility expectations can be extracted, and if market participants are rational, then these implied volatilities should contain all the information that is relevant for the pricing, hedging and management of option contracts and portfolios.

Conventionally, implied volatilities are found by first equating observed option prices to the Black–Scholes–Merton (Black and Scholes (1973), Merton (1973)) theoretical prices, and then by inverting for the unknown and unobservable volatility parameter, given the characteristics of the option contract (strike level K, option maturity T), the characteristics of the underlying asset (price level S, cash flows distributed by the asset until T, expressed either as a constant proportion δ of the asset price S or as a present value D of discrete flows at known times) and the rate of return of a riskless bond, i.e. given option prices

$$\sigma^{BS} : (S, K, T, r, \delta) \rightarrow \sigma^{BS} (S, K, T, r, \delta) \tag{1}$$

Contrary to the Black–Scholes–Merton assumption of constant (or deterministically time–dependent) volatility, the mapping in (1) has two features that have attracted the interest of researchers and practitioners in financial modeling: First, the volatilities implied from observed contracts systematically vary with the options strike price K and date to expiration T, giving rise to an instantaneously non–flat *implied volatility surface*, $\sigma^{BS} (K, T)$ (hereafter IVS). Canina and Figlewski (1993) and Rubinstein (1994) provide evidence that when plotted against moneyness, $m = K/S$ (the ratio of the strike price to the underlying spot price), implied volatilities exhibit either a 'smile' or a 'skew', while Heynen et. al. (1994), Xu and Taylor (1994) and Campa and Chang (1995) show that implied volatilities are a function of time to expiration and thus exhibit a 'term structure'. The second feature is that the IVS changes dynamically over time, i.e. $\sigma_t^{BS} (K, T)$ with a time subscript, as prices in the options market respond to new information that affects investors' beliefs and expectations.

In order to exploit this empirically observed profile of the IVS for forecasting purposes, researchers and practitioners usually fit parametric specifications of time–to–maturity and moneyness to the available IVS at a point in time. A well–cited example of this approach is offered by Dumas, Fleming and Whaley (1998) that treat implied volatilities as linear–in–parameters functions of the strike price and time–to–maturity (what they refer to as the ad hoc implied volatility model). Estimating the model on weekly cross sections of S&P 500 option prices, they observe that the parameter estimates produced are highly unstable over time, a result that is also confirmed by Christoffersen and Jacobs (2004).

This work is motivated by this instability of estimates of such parametric specifications, which is evidence of substantial time variation in the IVS. In this paper we propose an approach that jointly models the cross–sectional characteristics and the time–series dynamics of the IVS. Thus our approach augments the existing literature that either focuses on the cross–sectional fitting of the IVS, ignoring the time dimension, or models the time–series dynamics of an arbitrarily chosen segment on the IVS.

The approach is straightforward to set up and delivers accurate forecasts of implied volatilities for any moneyness or maturity level. It is based on simple time–series models of the evolution of few orthogonal statistical factors that are identified in the daily dynamics of the IVS. These statistical factors have a natural interpretation in the law of motion of the IVS that is intuitively appealing to practitioners.

To assess the performance of our IVS modeling approach, we study its ability to accurately predict the level and the direction of change of 1 to 5–day ahead implied volatilities across the surface. The forecasting accuracy is found to be very good, at least up to 3 days ahead, both in absolute terms and relatively to natural benchmarks such as a random walk for implied volatilities. The superior forecasting ability of the approach relatively to benchmarks is more profound for medium to long option maturities and is equally good across the moneyness dimension, which is quite interesting.

We demonstrate the approach using daily time–series of implied volatilities for a cross–section of currency options on 30 different currencies quoted against the Euro from the OTC market, and find that both in–sample performance and out–of–sample forecasting are uniformly good, across currencies. Moreover, there is nothing in the approach that prevents it from being applied to exchange–traded contracts.

A few existing papers are related to our approach, with Diebold and Li (2006) and Gonçalves and Guidolin (2006) closest in spirit. Both papers first apply parametric specifications at the cross–sectional level, and then fit time series models on the coefficients estimated from the first step. Moreover, both papers are concerned with forecasting: the yield curve and the IVS of S&P 500 index options respectively.

Diebold and Li (2006) apply a Nelson and Siegel (1987) type of specification to the yield curve derived from the cross–section of U.S. government bond prices, and the estimated coefficients of the specification are fitted to an autoregressive integrated moving average model. Gonçalves and Guidolin

(2006) first estimate a five–parameter version of the ad hoc implied volatility model of Dumas et. al. (1998) on the volatility surface implied by S&P 500 index options. They then model the time evolution of estimated parameters with a vector autoregressive model.

In contrast, our approach does not impose any "ad hoc" parametric specification on the cross–sectional description of the IVS: we simply decompose the surface into *orthogonal* unobservable *linear statistical factors* that are derived directly from the data and completely characterise the IVS. These factors capture systematic movements in the evolution of the IVS and turn out to have an intuitive interpretation. A parsimonious version of the model, with only few significant factors (decided via a number of criteria) used, turns out to produce both excellent in–sample fit and accurate out–of–sample forecasts.

Our paper is also related to Skiadopoulos, Hodges and Clewlow (1999), Derman and Kamal (1997), Mixon (2002) and Cont and da Fonseca (2002) that apply similar statistical–factor decompositions of the volatility surface implied by equity index options and attempt to relate the factors to observable economic variables; however none of the aforementioned contributions is concerned with using the factors for forecasting purposes as we do here.

Furthermore–and to the best of our knowledge–our paper is the first that examines the dynamics of the IVS from currency options; several authors have reported results from index options markets, however applications to the currency options market are only concerned with the time–to–maturity dimension of the IVS (e.g. Xu and Taylor (1994) or Krylova et. al. (2005)) and not the dynamics of the whole surface.

The rest of the paper is organised as follows: Section 2 describes the data, presents our methodology for describing the implied volatility surface and discusses the estimation results of this methodology. In section 3 we propose two alternative model based on our approach, that can capture the time–series dynamics of the cross–sectional estimates obtained in the previous section. Section 4 is devoted to the assessment of the out–of–sample forecasting performance of our approach, while Section 5 concludes the paper.

2 The volatility surface implied by currency options

2.1 The data

The data used in this study consist of daily time–series of implied volatilities for a cross–section of OTC currency options on 30 different currencies quoted against the Euro, kindly supplied by a major market participant. The time series are from 1/1/1999 to 21/5/2007, a total of 2,184 weekdays. The currencies examined and some exchange rate statistics are reported in Table 1.

In comparison to exchange–traded currency options, the OTC market is far more liquid. According to a Bank of International Settlements survey (2007), the outstanding notional amount of OTC currency options on the Euro in December 2006 was approximately 3.65 trillion US$ (2.54 trillion Euros). The corresponding amount of exchange–traded currency options was 78.64 million US$ (54.77 million Euros), far less than 1% of the notional amount outstanding in the OTC market.

As is typical in the OTC market, currency options are quoted in terms of the implied volatilities, which are conventionally converted into prices using the Garman and Kohlhagen (1983) version of the Black and Scholes (1973) option pricing formula. Our data–set consists of implied volatility quotes for the following fourteen expirations: 1 week, 1 month, 2 months, 3 months, 6 months, 9 months, 12 months, 18 months, 2–5 years, 7 years and 10 years. For each of these maturities, the implied volatility is observed for options with five different Black–Scholes deltas: OTM puts with $\Delta^{BS} = -0.10$ and $\Delta^{BS} = -0.25$, ATM calls and OTM calls with $\Delta^{BS} = 0.10$ and $\Delta^{BS} = 0.25$, where

$$\Delta^{BS} = \frac{\partial O}{\partial S} = \begin{cases} e^{-r_f T}\mathcal{N}(d_1), & \text{if the option } O \text{ is a call} \\ e^{-r_f T}\left[\mathcal{N}(d_1) - 1\right], & \text{if the option } O \text{ is a put} \end{cases} \tag{2}$$

$$d_1 = \frac{\ln(S/K) + (r_d - r_f + \sigma^2/2)\,T}{\sigma\sqrt{T}} \tag{3}$$

with r_d, r_f the risk–free interest rate in the domestic and the foreign country respectively, S the spot exchange rate, K the strike price of the option, T the time to option maturity in years and σ the exchange rate's volatility. Hence, for each exchange rate and on each observation date, a vector of 70×1 implied volatilities is observed.

Code	Currency	Average	Min-Max
AUD	Australian $	1.679	1.504-1.915
BRL	Brazilian Real	2.678	1.395-3.977
CAD	Canadian $	1.496	1.256-1.804
CHF	Swiss Franc	1.544	1.444-1.656
CLP	Chilean Peso	648.5	459.6-849.3
CZK	Czech Koruny	32.17	27.48-38.68
DKK	Danish Kroner	7.445	7.416-7.512
GBP	British	0.658	0.571-0.724
HKD	Hong Kong $	8.564	6.463-10.68
HUF	Hungarian Forint	253.5	234.5-284.6
IDR	Indonesian Rupiahs	9916.9	6726.9-13220.4
INR	Indian Rupees	49.90	38.65-60.03
ISK	Icelandic Kronur	83.18	68.06-97.44
JPY	Japanese ¥	126.2	89.34-164.1
KRW	South Korean Won	1236.9	943.4-1517.8
MXN	Mexican Peso	11.41	7.576-15.31
MYR	Malaysian Ringgits	4.222	3.149-6.407
NOK	Norwegian Kroner	8.058	7.228-8.947
NZD	New Zealand $	1.961	1.638-2.302
PHP	Philippine Peso	55.54	36.80-76.59
PLN	Polish Zlotych	4.064	3.351-4.900
RUB	Russian Ruble	31.21	23.13-37.85
SEK	Swedish Kronor	9.064	8.070-9.937
SGD	Singapore $	1.863	1.453-2.233
SKK	Slovakian Koruny	40.89	32.83-48.30
THB	Thailand Baht	44.15	34.30-53.20
TRY	Turkish (New) Lira	1.341	0.370-2.139
TWD	Taiwanese (New) $	36.32	26.48-45.47
USD	United States $	1.100	0.829-1.366
ZAR	South African Rand	8.013	6.099-12.09

Table 1: Average, minimum and maximum middle exchange rates of 30 different currencies against the Euro from January 1999 to May 2007. Source: European Central Bank.

Three necessary exclusionary criteria are applied to all currency options: [a] Days with at least one of the 70 implied volatilities missing are excluded, [b] days with flat implied volatility profiles (i.e. no "smile" or "skew") *for all 14 maturities* are excluded as misrecordings, and [c] maturities for which the implied volatility does not change from day to day in more than half the weekdays in our sample are excluded, as thinly–traded. Table 2 report the starting date and the number of days remaining in our sample after the above criteria have been applied.

Several different profiles of implied volatility surfaces are observed in our sample period. As an indication, in Figures 1–4 the average IVS profile and the daily standard deviation of the IVS from EUR/USD and HUF/EUR options are plotted. In the USD/EUR case, the implied volatility surface exhibits a clear symmetric "smile" with an increasing term structure on average, and a fair amount of variability around this average profile (ranging from a fourth to a tenth of its typical value). In contrast, the HUF/EUR implied volatility surface exhibits a "skew", with either an increasing or a humped–shaped term structure, and a significantly asymmetric variability for short maturities. Similar patterns emerge in all currencies examined; to conserve space the corresponding figures for the remaining 28 currencies are relegated to Appendix D (available from the authors upon request).

Given the origin of the data, one possible criticism is that idiosyncratic effects, specific to the market participant supplying the quotes, could influence the analysis. There are however reasons to believe that such effects (if any) are not strongly affecting our analysis. First, our focus here is on systematic factors in the volatility surface, not on specific events or outliers of the surface. Secondly, given the liquidity of the market and the size of the market participant supplying the data, it should be fairly unlikely that our data are substantially away from typical values.

However using OTC data has many advantages in comparison to exchange–traded data. Besides superior liquidity, OTC currency options are available for longer maturities than the currency options traded in exchanges. Moreover, OTC options have a constant time–to–maturity, unlike exchange–traded options whose maturity varies from day to day. In practical terms, this alleviates the need for grouping options into maturity bins (see for example Skiadopoulos et. al. (1999)) or for creating synthetic *fixed–maturity* series via interpolation (as in Alexander (2001)). This should translate to less noisy IVSs and more precision in the identification of factors affecting their dynamics. Similar OTC currency options data have been used in previous studies by Campa and Chang (1995), (1998) and Christoffersen and Mazzotta (2005); the latter study actually concludes that OTC currency options data are of superior quality for volatility forecasting purposes.

Currency	Start Date	# of days	Currency	Start Date	# of days
AUD	08-Sep-2000	1729	MXN	02-Jan-2006	361
BRL	08-Apr-2003	1064	MYR	24-May-2006	259
CAD	03-Nov-2003	883	NOK	04-Sep-2000	1745
CHF	11-Jul-2000	1783	NZD	05-Dec-2005	381
CLP	08-Dec-2004	638	PHP	08-Sep-2003	965
CZK	05-Sep-2000	1741	PLN	05-Dec-2005	381
DKK	02-Jun-2004	773	RUB	03-Jan-2006	359
GBP	05-Sep-2000	1744	SEK	05-Sep-2000	1741
HKD	05-Dec-2005	381	SGD	05-Dec-2005	381
HUF	05-Dec-2005	381	SKK	08-Apr-2003	1062
IDR	05-Apr-2005	555	THB	05-Jan-1999	2172
INR	05-Dec-2005	381	TRY	14-Nov-2000	1688
ISK	03-Jan-2006	357	TWD	05-Dec-2005	381
JPY	04-Sep-2000	1745	USD	04-Sep-2000	1745
KRW	27-Apr-1999	2082	ZAR	05-Dec-2005	381

Table 2: For each of the thirty different currency options in our sample, the table reports the starting date and the number of trading days in the time series. The end date in all time series is 21/5/2007.

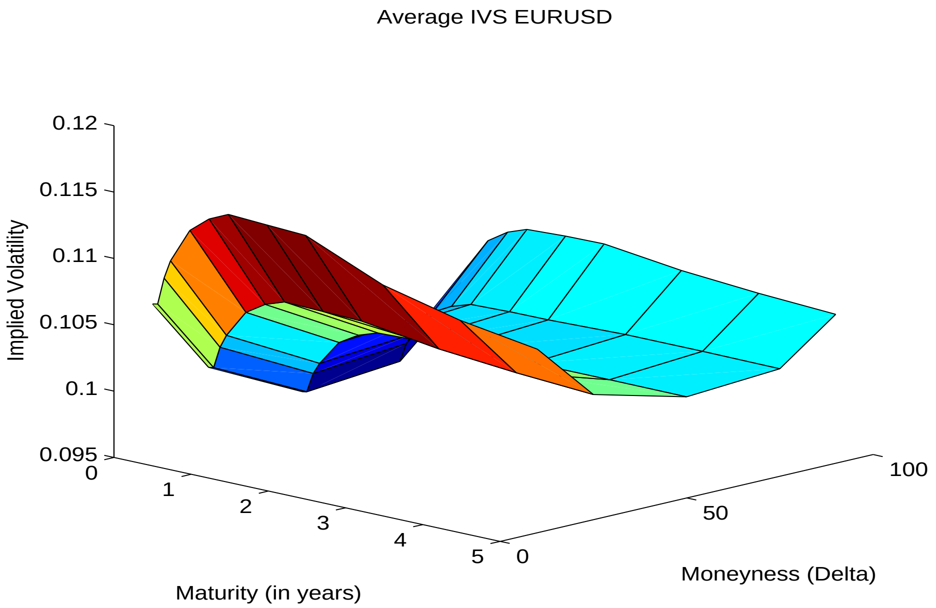

Figure 1: Average implied volatility surface from EUR/USD options, for the period 4/9/2000–21/5/2007.

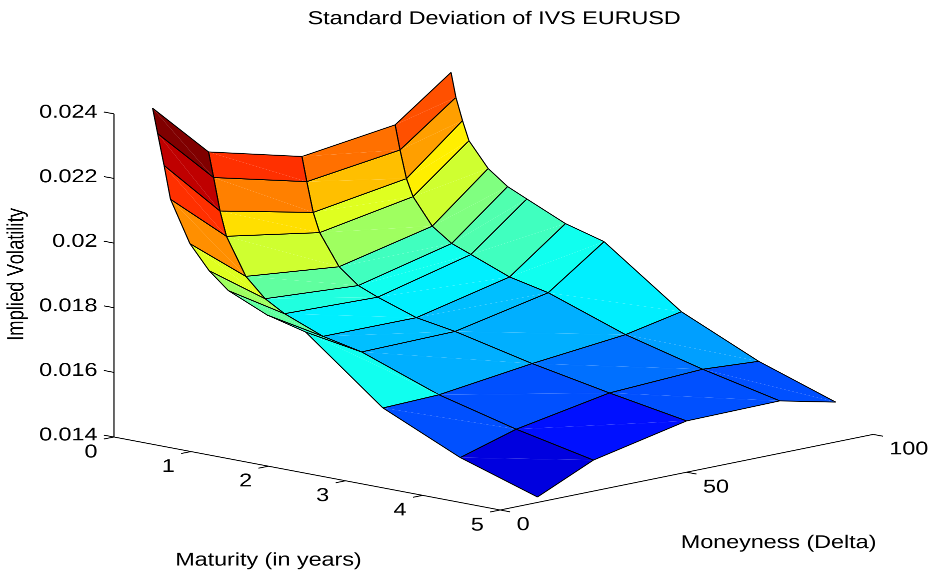

Figure 2: Daily standard deviation of EUR/USD implied volatilities as a function of moneyness and time to maturity for the period 4/9/2000–21/5/2007.

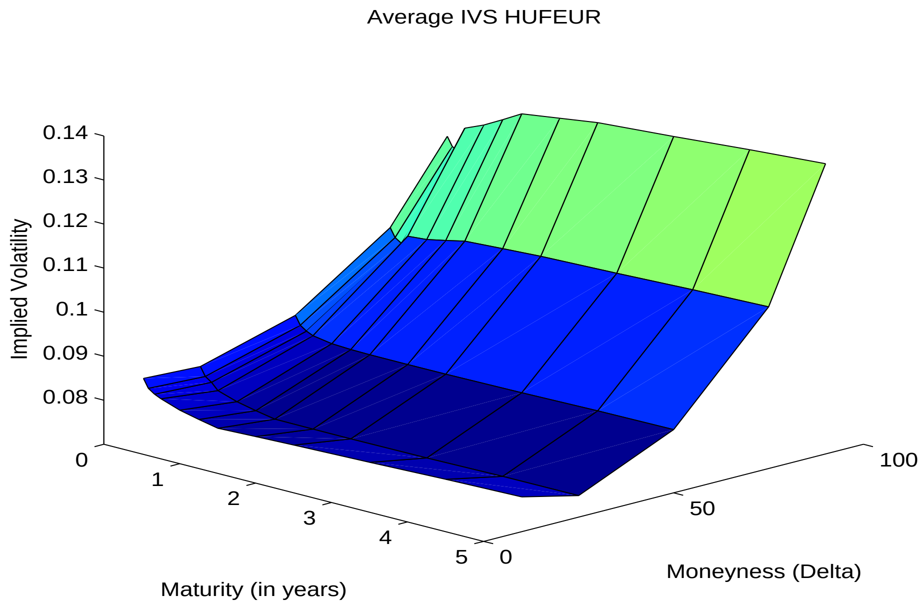

Figure 3: Average implied volatility surface from HUF/EUR options, for the period 5/12/2003–21/5/2007.

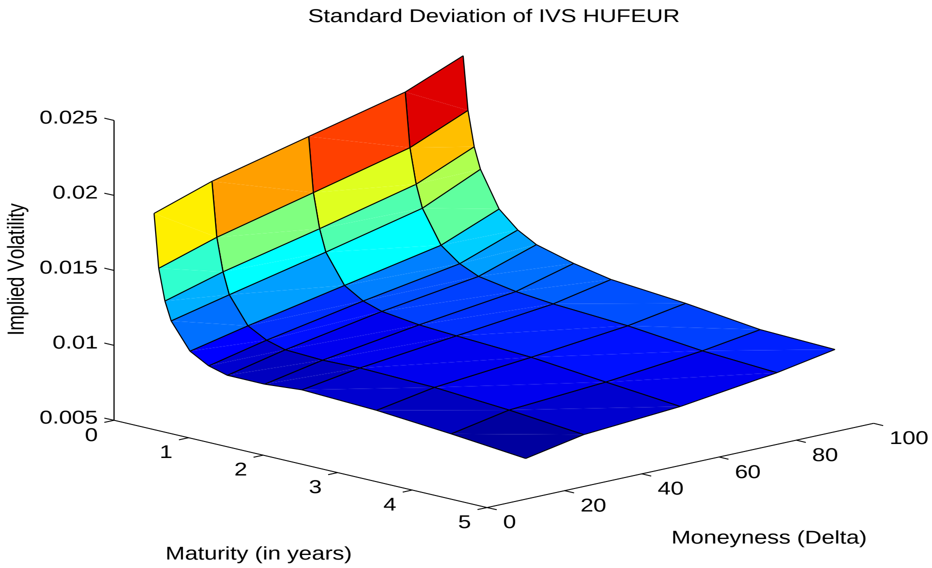

Figure 4: Daily standard deviation of HUF/EUR implied volatilities as a function of moneyness and time to maturity for the period 5/12/2005–21/5/2007.

2.2 Decomposition of the implied volatility surface

A common practice in describing the implied volatility surface on any given day, is to fit cross–sectionally a parametric specification of moneyness and time–to–maturity. For example, following Dumas et. al. (1998), Gonçalves and Guidolin (2006) use the following specification in their investigation of the IVS of S&P 500 index options:

$$\ln \sigma_i = \beta_0 + \beta_1 M_i + \beta_2 M_i^2 + \beta_3 T_i + \beta_4 \left(M_i \times T_i\right) + \epsilon_i, \tag{4}$$

where ϵ_i is the random error term, $i = 1, \ldots, p$, with p the number of options available in each daily cross section, σ_i, M_i and T_i the implied volatility, moneyness and time to maturity of option i respectively. Each day, a vector $\beta = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)'$ is stored, and a VAR model is fitted to the time–series of such vectors, with a view towards forecasting future implied volatility. Equation (4) can be rewritten, in matrix notation, as

$$\sigma_{p\times 1} = V_{p\times 5}\beta + \epsilon_{p\times 1}, \tag{5}$$

with $\epsilon_{p\times 1}$ a vector of random errors and $V_{p\times 5}$ a matrix of option characteristics with typical rows $(M_i, M_i^2, T_i, M_i \times T_i)$, $i = 1, \ldots, p$.

Instead of imposing a parametric specification of moneyness and time–to–maturity as matrix V in (5), we simply assume that the implied volatilities of the p observed options at a given date are coming from a linear model of the form

$$\sigma_{p\times 1} = \mu_{p\times 1} + L_{p\times m} F_{m\times 1} + \varepsilon_{p\times 1} \tag{6}$$

where $\mathbb{E}\left[\sigma\right] = \mu$, $\mathbb{E}\left[F\right] = \mathbb{E}\left[\varepsilon\right] = 0$, $\mathbb{E}\left[\varepsilon\varepsilon'\right] = \mathbb{I}$, and $\mathbb{E}\left[FF'\right] = \Psi$, with $\mathbb{I}$ the $p \times p$ identity matrix and Ψ a diagonal matrix.

Equation (6) assumes there are m common factors, the vector $F_{m\times 1}$, that are associated with the observed implied volatility levels. The $p\times m$ matrix L, the factor loadings matrix, represents the sensitivity of the implied volatility to each of the factors. The covariance structure that equation (6) implies is $\Sigma = LL' + \Psi$.

Since theory does not direct the choice for the number of factors m that should be used, we let the data decide on a currency–by–currency basis. Specifically, we estimate the factor loadings by the principal components method.[1] In order to decide how many factors are required to adequately

[1] The Principal Components method is the empirical (sample) version of what is known as the Karhunen–Loève decomposition in the theory of stochastic processes. It is essentially a statistical method of extracting the most important uncorrelated sources of variation in a multivariate system that is characterised by a high degree of collinearity. It has been used extensively in finance, especially in the modeling of yield curves (see Rebonato (2000), Wilson (1994) and the references therein).

explain the variation of each IVS, we use the Guttman–Kaiser and Velicer criteria, both reviewed in Appendix A.

Using the estimated dimension $\widehat{m}$, the estimated factor loadings $\widehat{L}$, and estimated vectors and variances $\widehat{\varepsilon}, \widehat{\Psi}$ as true values, then from (6) the factor realisations or scores can be recovered. Specifically, for each day, the following weighted least–squares estimation is performed

$$\min_{F_m} \sum_{i=1}^{p} \frac{\varepsilon_i^2}{\psi_i} = \min_{F_m} \varepsilon' \Psi^{-1} \varepsilon = \min_{F_m} (\sigma - \mu - LF)' \, \Psi^{-1} (\sigma - \mu - LF) \tag{7}$$

The approach accounts for the fact that the diagonal elements of Ψ are most likely not equal (see Bartlett (1937), Knez, Litterman and Scheinkman (1994) and Mixon (2002)).

Replacing population with sample analogs in (7) and minimising at each day, the estimated factor scores are recovered,

$$\widehat{F}_m = \left(\widehat{L}' \widehat{\Psi}^{-1} \widehat{L}\right)^{-1} \widehat{L}' \widehat{\Psi}^{-1} (\sigma - \widehat{\mu}) \tag{8}$$

These are by construction orthogonal.

Table 3 reports the variance decomposition results for the volatility surface implied by the each of the 30 currency options in our sample. The number of retained principal components reported in the table is the most conservative (greatest) from the Guttman–Kaiser and Velicer criteria.

On average, the retained factors can explain 99.09% of the variance in the daily volatility surface implied by currency options. The proportion of variance explained ranges from a minimum of 97.20% for options on CZK/EUR to a maximum of 99.91% in the TWD/EUR case.

On average, the first principal component accounts for 88.09% of the IVS variation across currencies; it can range from 60.67% (ISK/EUR) to 97.16% (AUD/EUR). In all but two of the currency options examined it can explain more than 75% of the IVS.

The second and third principal components contribute, on average, an additional 8.29% and 2.11% respectively, while PCs 4–9 collectively explain approximately 1.5% of the IVS variation, wherever retained. Closer inspection of the results in Table 3 suggests that the three principal components usually retained in empirical applications explain collectively 98.49% on average; moreover, in all but one case (CZK/EUR) they can explain 95% of the IVS or more.

To get a better feeling of the results, the estimated three–factor loadings are plotted against moneyness (Δ^{BS}) and time–to–maturity (T) in Figures 5

Currency Code	Velicer f_n	G–K $\bar{\lambda}$	Total Variance Explained (%)	Proportion of Total Variance Explained by PC 1	PC 2	PC 3	PC 4	PC 5	PC 6	PC 7	PC 8	PC 9
AUD	5	1	99.83	97.16	1.64	0.66	0.21	0.16				
BRL	7	2	99.76	85.15	11.51	1.77	0.76	0.29	0.19	0.09		
CAD	2	3	97.95	91.35	4.66	1.94						
CHF	4	4	98.31	79.31	11.33	5.15	2.52					
CLP	7	2	99.86	91.46	6.36	1.14	0.46	0.24	0.13	0.07		
CZK	4	4	97.20	79.94	10.83	4.16	2.27					
DKK	1	3	99.17	86.90	10.92	1.35						
GBP	3	2	99.31	95.75	2.53	1.03						
HKD	2	3	99.44	95.91	3.10	0.43						
HUF	7	3	99.82	91.58	5.31	1.98	0.45	0.24	0.17	0.09		
IDR	5	3	99.57	87.70	7.89	2.54	1.16	0.28				
INR	3	2	98.74	83.72	14.01	1.01						
ISK	3	3	97.81	60.67	30.47	6.67						
JPY	5	2	99.79	94.98	3.38	1.02	0.30	0.11				
KRW	5	2	99.41	93.13	4.11	1.36	0.50	0.31				
MXN	9	2	99.71	90.41	5.46	1.60	0.95	0.51	0.28	0.23	0.16	0.11
MYR	3	3	98.84	64.46	30.55	3.83						
NOK	4	3	99.37	91.19	4.93	2.51	0.74					
NZD	7	2	99.88	88.33	9.18	1.19	0.64	0.23	0.20	0.11		
PHP	1	3	98.14	82.75	11.96	3.43						
PLN	3	2	99.17	91.09	7.45	0.63						
RUB	7	3	99.58	89.42	6.17	1.87	1.11	0.53	0.33	0.15		
SEK	1	3	97.66	87.88	5.63	4.15						
SGD	3	3	99.05	92.85	4.50	1.70						
SKK	3	3	97.76	76.69	17.34	3.73						
THB	1	3	99.44	95.09	2.65	1.70						
TRY	5	2	99.64	94.36	2.85	1.73	0.47	0.23				
TWD	6	2	99.91	96.08	2.41	0.88	0.31	0.19	0.04			
USD	3	2	99.41	96.20	2.03	1.18						
ZAR	3	2	99.17	91.22	6.94	1.01						
Average	4.1	2.6	99.09	88.09	8.29	2.11	0.86	0.28	0.19	0.12	0.16	0.11

Table 3: For each of the thirty different currency options in our sample, the table presents the decomposition of the exchange rate implied volatility surface per retained principal component. The number of retained principal components reported in the table is the greatest from the Guttman–Kaiser ($\bar{\lambda}$) and Velicer (f_n) criteria.

and 6 for two representative cases: options on the EUR/USD and HUF/EUR exchange rates.[2,3]

The character of the factors seems intuitive. Factor 1 in Panels (a) represents a shock that affects all maturities and deltas (i.e. moneyness) in the same direction (same sign). The effect is strongest at the short horizons and it dampens over time. It can be interpreted as a *level effect*, and it is consistent with a mean–reverting model of stochastic volatility. Factor 1 affects OTM and ATM volatility differently. This is consistent with the notion that a change in volatility alters the steepness of the "smile" and correspondingly the skewness of the implied risk neutral density.

Factor 2 affects short–term and long–term impled volatility with different signs (it appears to change sign around the 3–month/6–month option maturity). Thus, this factor separates between different ends of the volatility term structure, i.e. it is a *term–structure effect*. The effect is almost uniform across the moneyness dimension.

Finally, the third factor appears to change sign ATM. It separates the effect between OTM puts and calls and it is present in all maturities. However its effect is more pronounced for short–dated options. Changes along this factor alter the steepness of the implied volatility smile; it can be interpreted as a *jump–fear effect*. Similar factors emerge in all currency options examined. Factors 4–9 that contribute less than 2% whenever retained, are more difficult to interpret.

Turning our attention to the factor realisations or scores, $\widehat{F}_m$, Figures 7 and 8 (from EUR/USD and HUF/EUR options again) suggest that the IVS is fluctuating significantly over time. The autocorrelations and partial autocorrelations plotted suggest that some structure may exist in the factor dynamics. Thus, in the next section, we turn our attention to the in–sample modeling of the time–series dynamics of the factors, with a view towards forecasting future implied volatility out–of–sample.

[2] It should be noted that the estimated factor loadings are unique only up to an orthogonal rotation of the factors. An orthogonal rotation corresponds to altering the directions of the vectors, thus changing the interpretation of the factors, however it does not alter the space that the factors span. In the exposition and discussion that follows, the unrotated factor loadings are used.

[3] Again, the corresponding figures for all other currency options are available in Appendix D.

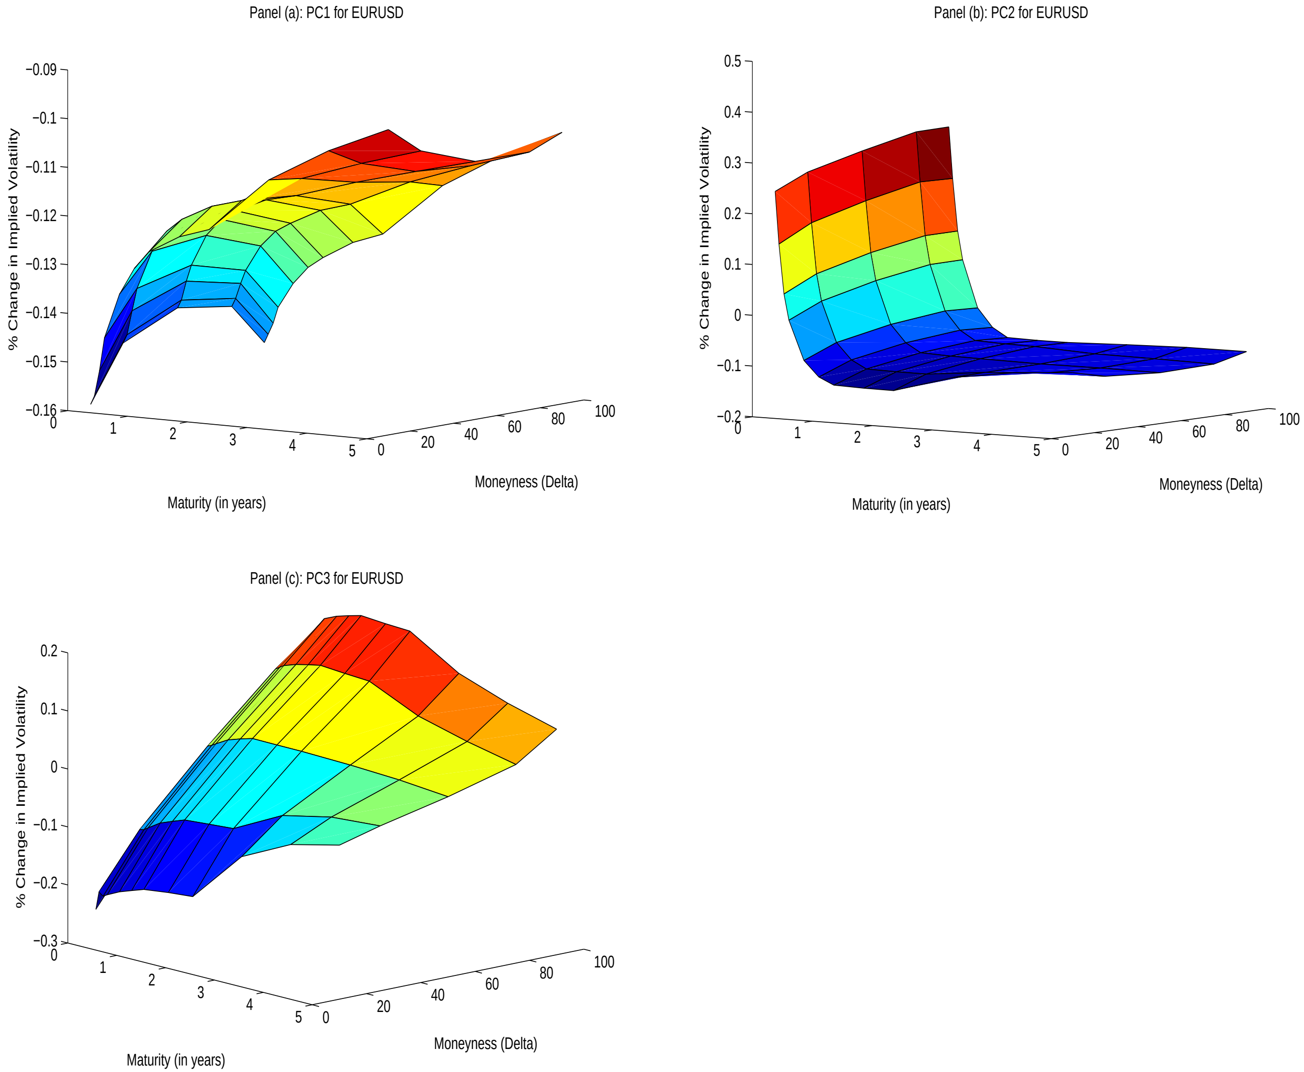

Figure 5: Factor loadings of the three retained principal components that are identified in the daily time series of the volatility surface implied by options on the EUR/USD exchange rate from 4/9/2000–21/5/2007.

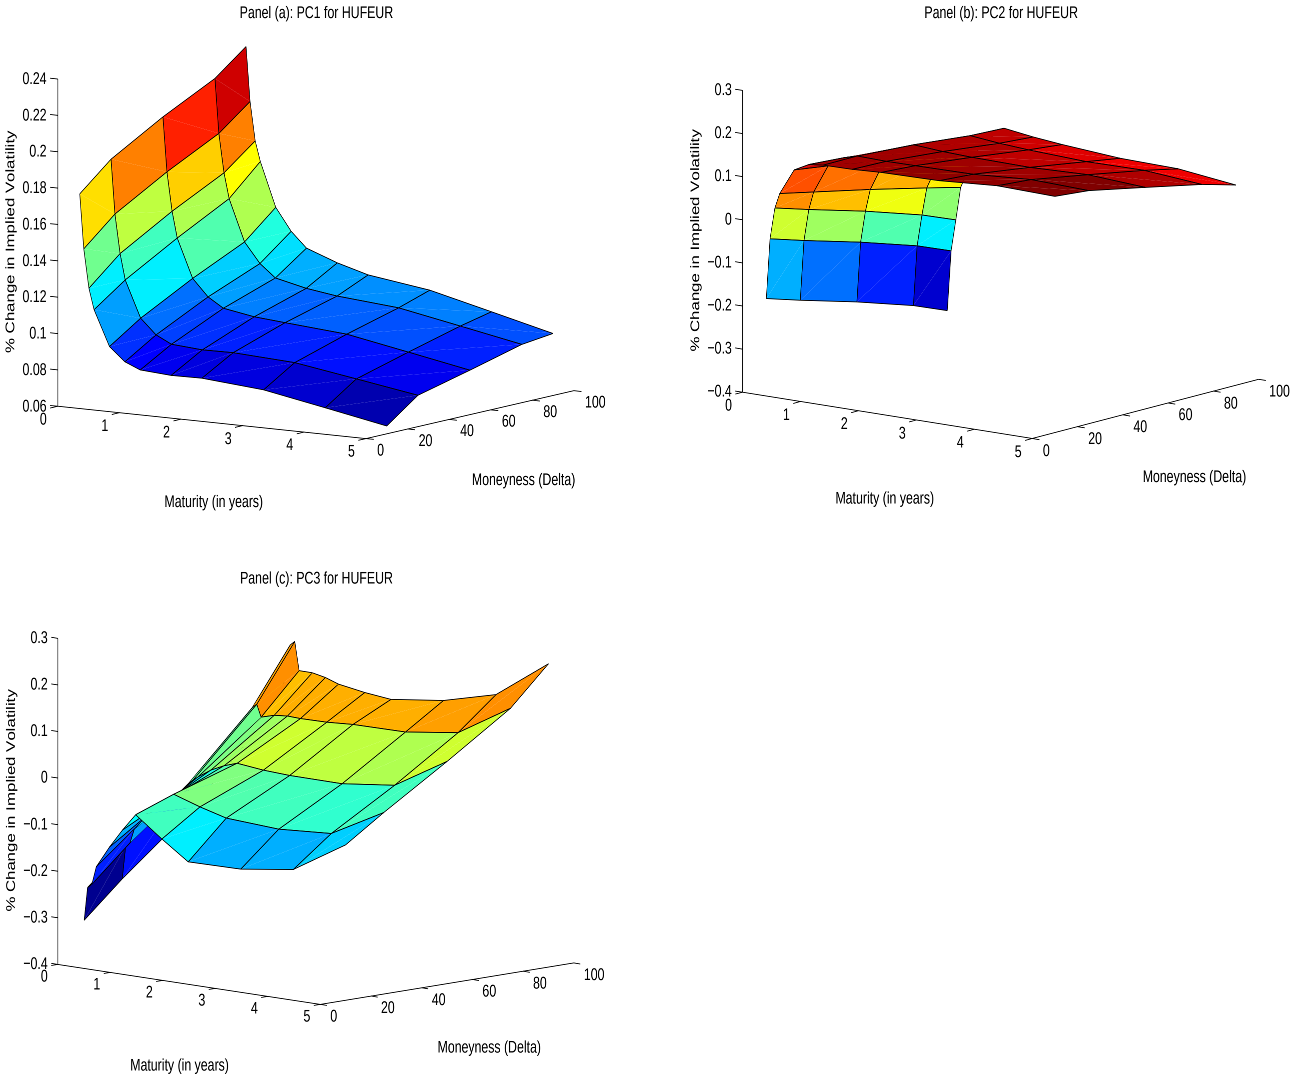

Figure 6: Factor loadings of the first three retained principal components that are identified in the daily time series of the volatility surface implied by options on the HUF/EUR exchange rate from 5/12/2005–21/5/2007.

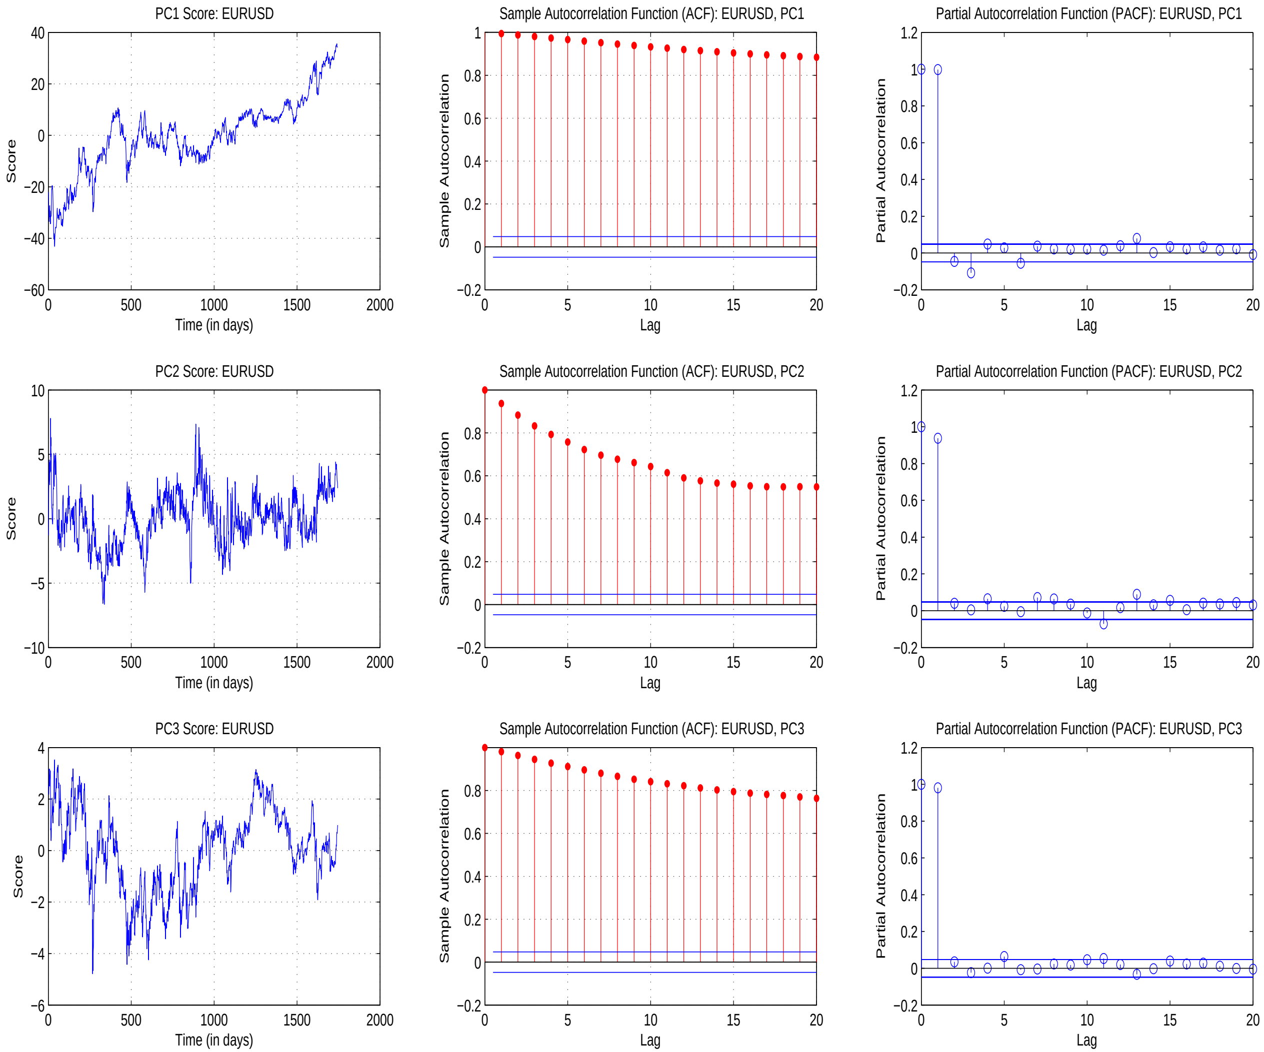

Figure 7: Factor realisations or scores of the three principal components that are identified in the daily time series of the volatility surface implied by options on the EUR/USD exchange rate from 4/9/2000–21/5/2007, and their autocorrelation (ACF) and partial autocorrelation (PACF) for up to 20 lags. The blue horizontal lines in the ACF and PACF graphs correspond to the $\alpha = 5\%$ significance level.

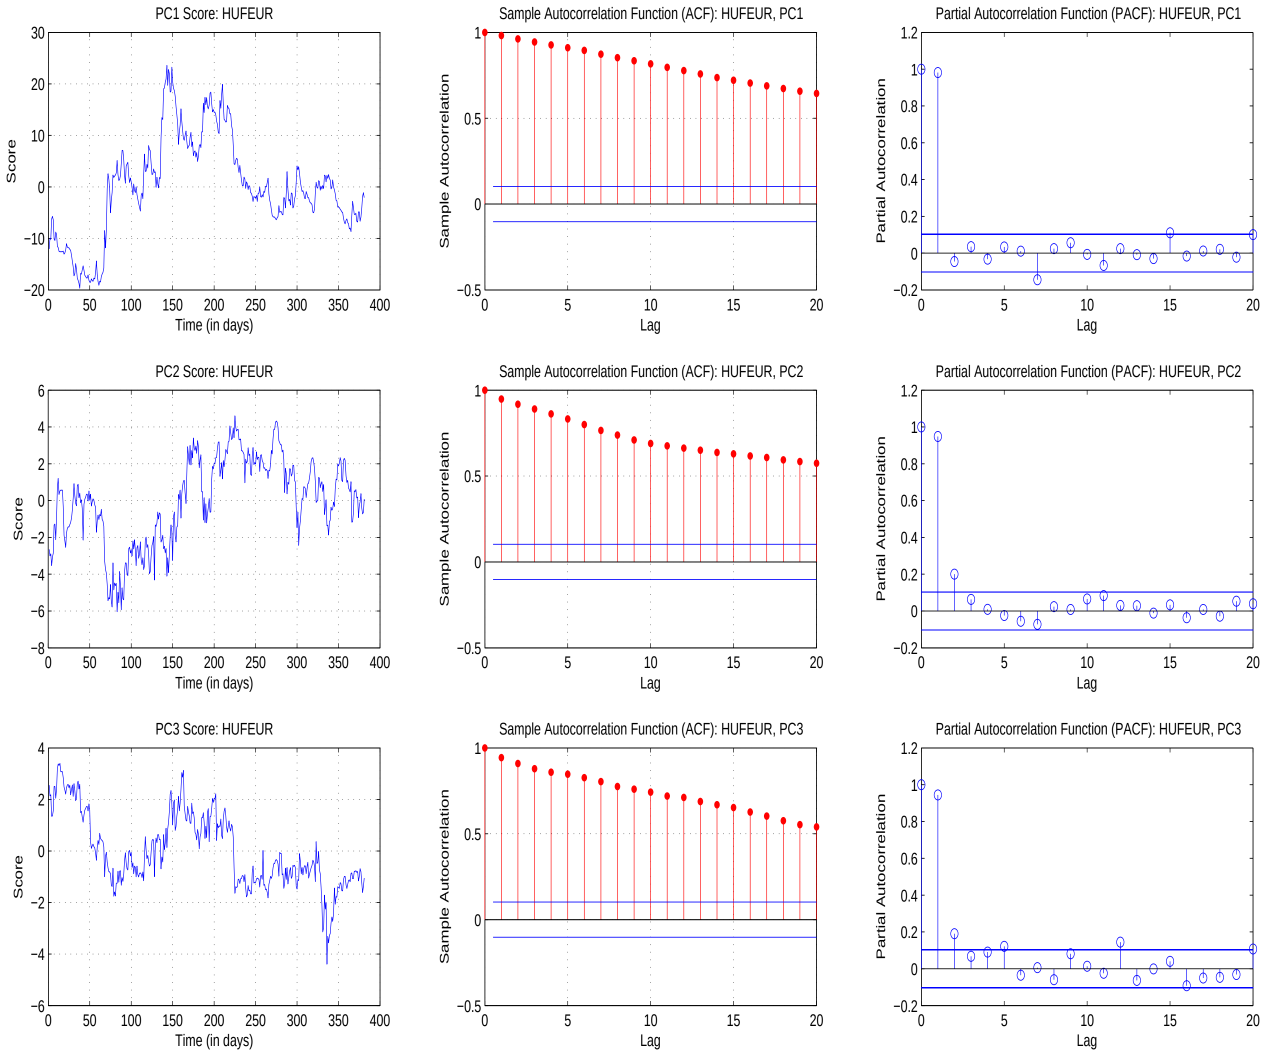

Figure 8: Factor realisations or scores of the first three principal components that are identified in the daily time series of the volatility surface implied by options on the HUF/EUR exchange rate from 8/4/2003–21/5/2007, and their autocorrelation (ACF) and partial autocorrelation (PACF) for up to 20 lags. The blue horizontal lines in the ACF and PACF graphs correspond to the $\alpha = 5\%$ significance level.

3 Modeling the dynamics of the implied volatility surface

In this section we model the time variation in the IVS as captured by the dynamics of factor realisations identified in the cross–sectional analysis that preceded.

In order to determine whether the implied volatility factor scores are stationary, the augmented Dickey–Fuller and Phillips–Perron unit root tests are conducted. The lag length of the tests for each series is decided by the Bayesian Information Criterion of Schwarz (1978). Results (relegated in Appendix B) indicate that the null of a unit root is soundly rejected in all cases at the 10% significance level, and in almost all cases at the 5%.

In section 2.2 we have identified 3–9 factors that explain the volatility surface implied by options on 30 different exchange rates. Instead of using all the factors identified in each IVS series in our sample, we have decided to retain and model only the first 3 factors for all currencies. One should expect that, if anything, this choice will worsen our results in the cases where several more factors are present in the data. However, our objective is to investigate whether a parsimonious model, with only the most important (and intuitive) factors, can achieve good in–sample fit and accurate out–of–sample prediction.

Following the lead of Gonçalves and Guidolin (2006), we first consider a vector autoregressive (VAR) model for the time series of estimated factors $\widehat{\mathbf{F}}_t = \left(\widehat{F}_{1,t}, \widehat{F}_{2,t}, \widehat{F}_{3,t}\right)'$ of the form:

$$\widehat{\mathbf{F}}_t = \mathbf{c} + \sum_{j=1}^{d} \mathbf{\Phi}\widehat{\mathbf{F}}_{t-d} + \mathbf{e}_t \tag{9}$$

where $\mathbf{e}_t \sim \mathcal{N}_{\text{i.i.d.}}(\mathbf{0}, \mathbf{\Omega})$. Although factors are virtually *contemporaneously orthogonal*, the intuition behind (9) is that *lagged realisations* of one factor (e.g. the "*jump–fear*") might influence the current realisation of another factor (e.g. the "*level*"). This is consistent with empirical observations regarding the IVS, such as that when the level increases, the steepness of the smile decreases, etc. We label (9) as model VAR, estimate it by OLS equation by equation, and select d by successive applications of the log–likelihood ratio test at the $\alpha = 1\%$ significance level, starting with a maximum value of $d = 10$.[4] Estimation results are summarised in Tables 4 and 5.

[4]The degrees of freedom correction factor proposed by Sims (1980) is used in the log–likelihood ratio tests.

PC	Code	d	$R^2_{adj.}$	$LB\,(d)$	$A\,(d)$	Code	d	$R^2_{adj.}$	$LB\,(d)$	$A\,(d)$	Code	m	$R^2_{adj.}$	$LB\,(m)$	$A\,(m)$
1	AUD	9	0.9580	0.6094	0.0000*	IDR	4	0.9459	0.9996	0.7228	PLN	7	0.9342	0.9998	0.0605
2		9	0.8870	0.5871	0.0000*		4	0.9367	0.6945	0.0000*		7	0.9013	0.8667	0.4837
3		9	0.9173	0.7350	0.0000*		4	0.9675	0.9999	0.6620		7	0.7353	0.2579	0.7148
1	BRL	6	0.9611	0.9991	0.0000*	INR	7	0.9343	0.8046	1.0000	RUB	7	0.8488	0.9639	0.9389
2		6	0.9043	0.1474	0.0000*		7	0.8598	0.8495	0.7361		7	0.6848	0.9990	0.9913
3		6	0.7956	0.0167*	0.0000*		7	0.8623	0.9922	0.1541		7	0.8763	0.9261	0.7014
1	CAD	6	0.9248	0.8298	0.0000*	ISK	5	0.9736	0.8177	1.0000	SEK	10	0.9701	0.1681	0.0000*
2		6	0.8045	0.9067	0.0000*		5	0.9620	0.0930	1.0000		10	0.9381	0.1469	0.0000*
3		6	0.7292	0.0011*	0.0000*		5	0.9377	0.7850	0.9962		10	0.9150	0.4525	0.0000*
1	CHF	8	0.9437	0.0255*	0.0000*	JPY	10	0.9779	0.9266	0.0000*	SGD	10	0.8398	0.9939	0.9998
2		8	0.9011	0.8653	0.0000*		10	0.9591	0.9031	0.0000*		10	0.8505	0.9777	0.8441
3		8	0.9402	0.9963	0.0000*		10	0.9103	0.5907	0.0000*		10	0.8799	1.0000	0.4011
1	CLP	9	0.9070	0.9998	0.0280*	KRW	6	0.9782	0.1701	0.0000*	SKK	2	0.9514	0.1953	0.0038*
2		9	0.8789	0.9764	0.0433*		6	0.9607	0.0000*	0.0000*		2	0.8581	0.0920	0.3107
3		9	0.9561	0.6061	0.7253		6	0.9115	0.0665	0.0000*		2	0.9081	0.0023*	0.0000*
1	CZK	9	0.9582	0.9981	0.0000*	MXN	4	0.9639	0.3297	0.7855	THB	10	0.9632	0.3900	1.0000
2		9	0.9107	0.4170	0.0000*		4	0.7903	0.3036	0.0000*		10	0.9484	0.0001*	0.0000*
3		9	0.9333	0.5220	0.0000*		4	0.7270	0.6976	0.0012*		10	0.9654	0.1428	0.0000*
1	DKK	1	0.9603	1.0000	1.0000	MYR	1	0.9411	0.1501	1.0000	TRY	8	0.9749	0.8547	0.3926
2		1	0.9726	0.9999	0.0199*		1	0.8986	0.1170	0.0220*		8	0.9047	0.9986	0.0000*
3		1	0.9463	0.0000*	0.9878		1	0.8673	0.8610	0.9999		8	0.6970	0.6156	0.0000*
1	GBP	10	0.9869	0.5964	0.0000*	NOK	8	0.9700	0.9130	0.0057*	TWD	3	0.7308	0.5829	0.0165*
2		10	0.9454	0.8447	0.0526		8	0.8250	0.1419	0.0000*		3	0.7542	0.7212	0.0006*
3		10	0.9335	0.9152	0.5774		8	0.8714	0.9932	0.0000*		3	0.7975	0.9959	0.9753
1	HKD	2	0.9056	0.8258	0.0382*	NZD	3	0.9352	0.4849	0.9030	USD	10	0.9764	0.6766	0.0000*
2		2	0.8724	0.9451	0.9296		3	0.8129	0.0934	0.2829		10	0.8722	0.0000*	0.0000*
3		2	0.6678	0.9995	0.0126*		3	0.5481	0.0000*	0.0008*		10	0.9574	0.55522	0.0000*
1	HUF	7	0.9277	0.7347	0.0186*	PHP	6	0.8740	0.5240	0.7748	ZAR	7	0.9661	0.6095	0.0010*
2		7	0.8757	0.9870	0.4454		6	0.4542	0.0037*	0.9999		7	0.9332	0.4348	0.7377
3		7	0.8439	0.2398	0.1955		6	0.8674	0.9966	0.7745		7	0.7873	0.4616	0.1263

Table 4: For the volatility surface implied by the thirty different currency options in our sample, the table reports the results from the estimation of the VAR model (equation (9)) on the first three factor scores. The lag length, d, is selected by successive applications of the log–likelihood ratio test at the 1% significance level, starting with a maximum value of $d = 10$. Under $LB\,(d)$ and $A\,(d)$, p–values for the Ljung–Box statistic (H_0 : absence of autocorrelation up to lag d in the residuals) and the Engle test (H_0 : i.i.d. Gaussian residuals) are reported. The length of the time series for each currency is reported in Table 2.
An * denotes that the null is rejected at the $\alpha = 5\%$ significance level.

	PC1		PC2		PC3	
Currency	is Granger–caused at $\alpha = 1\%$ by					
Code	PC2	PC3	PC1	PC3	PC1	PC2
AUD					✓	
BRL		✓	✓	✓	✓	✓
CAD		✓		✓		✓
CHF			✓	✓		
CLP	✓					
CZK	✓			✓	✓	✓
DKK					✓	
GBP					✓	
HKD					✓	
HUF			✓			
IDR	✓				✓	✓
INR						
ISK	✓	✓				
JPY	✓					✓
KRW			✓			
MXN	✓	✓	✓			
MYR	✓		✓			
NOK	✓	✓	✓			
NZD				✓		
PHP	✓		✓	✓		✓
PLN				✓	✓	
RUB						
SEK	✓			✓	✓	✓
SGD					✓	
SKK	✓	✓	✓	✓	✓	
THB					✓	✓
TRY	✓	✓		✓		✓
TWD		✓			✓	✓
USD	✓					
ZAR					✓	

Table 5: For the volatility surface implied by the thirty different currency options in our sample, the table reports the results of Granger–causality tests from the estimation of the VAR model (equation (9)) on the first three factor scores. A ✓ denotes evidence of Granger–causality at the $\alpha = 1\%$ significance level.

The adjusted R^2's in Table 4 suggest that the in–sample fit is extremely good. Moreover, our suspicion that the estimated factors might be interrelated is confirmed by Table 5, that reports Granger–causality tests as implied by the VAR estimation. The table reports significant (at the 1%) off–diagonal elements of $\mathbf{\Phi}$ in (9): for example, the tickmarks in JPY suggest that both the "level" and the "jump–fear" factor is Granger–caused by the "term–structure" of the IVS.

The Ljung–Box $LB\,(d)$ lack–of–fit statistic in Table 4 suggests that in the vast majority of cases the in–sample fit is fairly good, with uncorrelated residuals. However, in most currencies, several lags are required to achieve this fit, which suggests that there might be something that our VAR model is missing. Indeed, the results of Engle's (1982) ARCH/GARCH test in Table 4 suggest that there is substantial conditional heteroscedasticity in the residuals of the VAR fitted specification.

To account for this, we estimate a second model, labeled RMPQ, of the following form:

$$\begin{aligned}
\widehat{F}_{i,t} &= k_i + \sum_{j=1}^{r} \phi_{ij} \widehat{F}_{i,t-k} + \epsilon_{i,t} + \sum_{\kappa=1}^{m} \theta_{i\kappa} \epsilon_{i,t-\kappa} \\
\epsilon_{i,t} &\sim \mathcal{N}\left(0, h_{i,t}^2\right) \\
h_{i,t}^2 &= \omega_i + \sum_{n=1}^{p} a_{in} h_{i,t-n}^2 + \sum_{w=1}^{q} b_{iw} \epsilon_{i,t-w}^2 \\
i &= 1, 2, 3
\end{aligned} \tag{10}$$

This is a univariate autoregressive moving–average of order (r, m) model for the estimated factor scores separately, with Gaussian GARCH (p, q) innovations. In contrast to the model in (9), the system of equations in (10) accounts for conditional heteroscedasticity, but ignores relationships between estimated factors. We estimate RMPQ via maximum likelihood and select r, m, p, q by the Bayesian Information Criterion of Schwarz (1978), starting with maximum values $r, m = 10$ and $p = q = 4$.

The estimation results are reported in Table 6. Fairly parsimonious specifications can adequately fit the factors scores in sample, across currencies. In the majority of cases, both the structure of factor realisations and the heteroscedasticity in the residuals seem adequately modeled.

Table 6: For the volatility surface implied by the thirty different currency options in our sample. the table reports the results from estimating an $ARMA\,(r,m)$ process with Gaussian $GARCH\,(p,q)$ innovations in the first three retained principal components scores. The specification (r,m,p,q) is selected by the Bayesian Information Criterion of Schwarz (1978). starting with a maximum of $r=m=10$ and $p=q=4$. Under LB and A. p–values for the Ljung–Box statistic (H_0 : absence of autocorrelation in the residuals) and the Engle test (H_0 : i.i.d. Gaussian residuals) are reported. The length of the time series for each currency is reported in Table 2.
An * denotes that the null is rejected at the $\alpha = 5\%$ significance level.

Retained	Currency	Specification				Log	Residuals	
PC	Code	r	m	p	q	Likelihood	LB	A
1	AUD	1	0	1	1	-2518.90	0.0881	0.9760
2		4	2	1	1	-1680.60	0.5906	0.7760
3		2	0	1	1	-668.78	0.5554	0.9999
1	BRL	1	0	1	1	-2633.29	0.7831	0.9961
2		1	1	1	1	-2213.28	0.1975	0.7732
3		1	1	1	1	-1555.42	0.8345	0.7337
1	CAD	2	0	2	1	-1124.89	0.1068	0.8994
2		2	1	1	1	-840.58	0.8354	0.4787
3		3	3	2	1	-409.48	0.1356	0.4093
1	CHF	1	0	1	1	-1723.10	0.2978	0.5988
2		1	1	1	1	-1186.14	0.1010	0.7816
3		1	1	1	2	-585.61	0.3326	1.0000
1	CLP	3	4	1	2	-1202.61	0.0401*	1.0000
2		3	4	1	1	-783.29	0.0502	1.0000
3		1	0	1	1	-170.90	0.5444	0.9580
1	CZK	1	1	1	1	-2670.94	0.9979	0.8185
2		1	1	2	1	-1879.08	0.1916	0.0084*
3		2	1	2	1	-1172.28	0.1241	0.2224

Retained PC	Currency Code	Specification				Log Likelihood	Residuals	
		r	m	p	q		LB	A
1	DKK	4	0	2	2	465.12	0.0000*	1.0000
2		3	4	1	1	327.01	0.0000*	1.0000
3		3	2	2	1	805.53	0.0000*	0.8807
1	GBP	1	0	1	1	-1981.12	0.3515	0.9804
2		2	1	1	1	-1147.56	0.0742	0.8320
3		2	1	1	1	-356.63	0.0879	1.0000
1	HKD	2	3	1	1	-491.41	0.0001*	0.8360
2		1	0	1	1	-278.01	0.6221	0.9674
3		2	2	2	1	-11.23	0.2283	0.8749
1	HUF	1	0	1	1	-745.91	0.6568	0.9626
2		1	1	1	1	-412.07	0.9834	0.1547
3		1	1	2	1	-224.55	0.0168*	0.7627
1	IDR	1	0	1	2	-945.98	0.9942	1.0000
2		1	0	1	1	-543.16	0.8948	0.7072
3		4	2	1	1	-90.63	0.4347	1.0000
1	INR	4	1	1	1	-484.55	0.0000*	1.0000
2		3	2	2	1	-485.74	0.0000*	0.9523
3		1	4	1	2	-44.50	0.4122	0.9990
1	ISK	4	1	1	2	-898.02	0.0093*	1.0000
2		3	3	1	2	-831.68	0.0002*	1.0000
3		1	0	2	1	-641.69	0.3470	0.9993
1	JPY	1	0	1	1	-2804.62	0.0421*	0.7922
2		1	0	1	1	-1753.73	0.4914	0.9904
3		2	1	1	1	-1500.36	0.0406*	0.9560
1	KRW	2	0	1	1	-3661.59	0.3843	0.9769
2		2	1	2	2	-2283.15	0.0000*	0.9997
3		1	1	2	1	-2271.26	0.1897	0.9962
1	MXN	1	0	1	1	-627.81	0.7177	0.9892
2		1	1	1	1	-458.56	0.8403	0.8877
3		1	0	1	2	-321.64	0.1539	0.0025*

Retained PC	Currency Code	Specification				Log Likelihood	Residuals	
		r	m	p	q		LB	A
1	MYR	1	4	2	2	-434.02	0.0977	1.0000
2		3	4	1	2	-289.50	0.0001*	1.0000
3		1	0	1	2	-186.59	0.8860	1.0000
1	NOK	1	0	1	1	-2133.90	0.9016	0.9983
2		1	1	2	1	-1145.53	0.0853	0.9916
3		1	1	1	1	-317.69	0.0000*	1.0000
1	NZD	1	0	1	1	-623.61	0.7127	0.9577
2		2	1	1	1	-369.07	0.8456	0.6431
3		4	3	2	1	-162.70	0.1811	0.3949
1	PHP	1	1	2	1	-1555.82	0.0000*	0.9997
2		2	2	2	2	-1198.31	0.0000*	1.0000
3		1	1	2	2	-892.22	0.9982	1.0000
1	PLN	1	0	1	1	-578.68	0.9361	0.9962
2		1	0	1	1	-387.54	0.8244	0.8822
3		1	0	1	1	-96.29	0.1109	0.9634
1	RUB	1	0	1	1	-417.81	0.5697	0.8774
2		3	2	1	1	-247.36	0.0263*	0.9690
3		2	1	1	1	-85.63	0.8849	0.8018
1	SEK	1	0	2	1	-2211.00	0.0574	0.2681
2		2	1	2	1	-1231.17	0.1265	0.0000*
3		1	1	1	1	-1116.21	0.0000*	1.0000
1	SGD	1	0	1	1	-438.14	0.1127	0.9993
2		1	0	1	1	-308.92	0.4048	0.8743
3		1	0	1	1	40.03	0.6078	0.9761
1	SKK	3	4	1	1	-1574.24	0.4971	1.0000
2		1	0	1	1	-947.00	0.2155	0.9999
3		1	0	2	1	-395.57	0.2161	0.7405
1	THB	4	1	2	2	-4547.51	0.0000*	1.0000
2		1	0	2	1	-2170.60	0.4861	1.0000
3		2	1	2	2	-1573.14	0.9652	1.0000

Retained PC	Currency Code	Specification				Log Likelihood	Residuals	
		r	m	p	q		LB	A
1	TRY	1	2	2	1	-5231.82	0.6885	0.4584
2		1	1	1	2	-3893.00	0.6431	0.9195
3		4	2	1	2	-3576.51	0.0000*	1.0000
1	TWD	1	1	1	1	-489.33	0.0000*	0.7517
2		1	0	1	1	-168.74	0.1489	0.5548
3		1	0	1	1	-1.18	0.2013	0.7387
1	USD	1	0	1	1	-2620.68	0.0025*	0.4813
2		2	1	1	1	-1839.53	0.0000*	0.1601
3		1	0	1	1	-223.47	0.1696	0.9669
1	ZAR	2	1	1	1	-891.89	0.0182*	0.9996
2		1	1	1	1	-560.22	0.2777	1.0000
3		1	0	2	1	-357.71	0.4498	1.0000

As Gonçalves and Guidolin (2006) point out, both models estimated in this section, VAR and RMPQ, can be considered reduced–form analogs of more structural models, such as that proposed by Garcia, Luger and Renault (2003). There, predictability in the IVS dynamics arises as a consequence of investors' learning (from option prices) about the processes of fundamentals that are driven by persistent latent variables. In such a setting, simple reduced–form models like equations (9) and (10) can pick up such predictability. The results of this section suggest that in–sample such reduced–form models can perform very well.

4 Forecasting the implied volatility surface

Our modeling approach appears successful in–sample across different implied volatility surfaces. Nevertheless, good out–of–sample predictions should also be provided by an IVS model in order to be considered successful. The objective of this section is to investigate the out–of–sample performance of models VAR and RMPQ at forecasting future daily implied volatilities.

As a benchmark we include the so–called "random walk" (RW) model, according to which today's implied volatility for a given option contract is the best forecast of tomorrow's implied volatility, i.e. our best estimate of the future IVS is the IVS observed today. Although naive in spirit, Harvey and Whaley (1992) point out that practitioners widely use this model in practice.

We set up the forecasting comparison as follows: Using the first 100 implied volatility surfaces we observe for each of the 30 currency options, factor realisations $\widehat{F}_{1,2,3}$ are extracted by estimating equations (6) and (8). These are used as inputs to obtain estimates for the parameter of equations (9) and (10). Using the estimated parameters, forecasts of factor realisations for $\tau = 1, 2, \ldots, 5$ days ahead are produced every day, for the remaining sample period. Since the IVS on day $t + \tau$ can be decomposed into factors realisations τ days ahead, we can produce the whole predicted IVS via the forecasted $\widehat{F}_{t+\tau}$'s. We update the estimates of our approach every day for each new observation of the IVS; however, we keep the lag length d and the specification (r, m, p, q) in (9)–(10) constant and equal to the values estimated using the first 100 IVSs throughout.

To assess out–of–sample forecasting performance, the following three measures are computed each day, for τ–ahead forecasts, for each model:
[a] Mean squared error (MSE), the average (over the number of options) squared deviations of observed implied volatilities from the model's implied volatilities,
[b] Coefficient of determination (R^2) from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and
[c] Mean correct prediction of the direction of change (MCP), the average percentage of observations for which the model–predicted and the observed change in implied volatilities have the same sign

We should stress that the above comparisons are performed across the whole IVS and that we do not make full use of the information provided by our approach, since only the first three identified factor are modeled and employed in forecasting.

Table 7 reports the average values of the out–of–sample performance measures [a]–[c] for one–day ahead predictions. Both the RMPQ and the VAR models perform notably better than the "naive" random walk model: their

Currency Code	One–day ahead forecasts MSE			R^2			MCP		
	RMPQ	VAR	RW	RMPQ	VAR	RW	RMPQ	VAR	RW
AUD	0.0533	0.0569	0.1343	0.9854	0.9846	0.9632	52.65	52.33	47.66
BRL	0.6735	0.8158	1.2341	0.9409	0.9321	0.8999	52.07	52.21	51.19
CAD	0.0368	0.0412	0.2923	0.9377	0.9307	0.6998	53.18	52.86	48.12
CHF	0.0489	0.0502	0.0442	0.9384	0.9386	0.9470	52.34	52.04	47.74
CLP	0.1384	0.1521	0.6249	0.9240	0.9168	0.7774	50.08	49.34	62.71
CZK	0.1220	0.1241	0.0915	0.9193	0.9199	0.9486	50.65	50.70	53.44
DKK	0.0409	0.0471	0.1982	0.9856	0.9714	0.9488	52.47	52.12	49.63
GBP	0.0346	0.0365	0.0719	0.9801	0.9790	0.9608	51.90	51.96	48.50
HKD	0.0418	0.0440	0.5527	0.9715	0.9708	0.7661	50.69	52.22	56.96
HUF	0.0877	0.0907	1.3004	0.9040	0.9059	0.4556	52.26	52.20	45.96
IDR	0.2201	0.2605	0.7358	0.9501	0.9442	0.8734	53.21	53.04	53.23
INR	0.0710	0.0772	0.8744	0.9442	0.9418	0.5795	50.33	50.19	57.52
ISK	0.6458	0.8593	3.3637	0.7445	0.7401	0.6526	51.44	50.30	53.39
JPY	0.0758	0.0800	0.1826	0.9843	0.9836	0.9619	52.38	52.27	46.24
KRW	0.3119	0.3169	0.1996	0.9772	0.9771	0.9862	51.77	51.48	49.70
MXN	0.1326	0.1398	1.4033	0.9414	0.9401	0.6619	53.15	52.13	55.23
MYR	0.1147	0.1165	0.9798	0.6939	0.7258	0.3990	51.13	51.22	51.92
NOK	0.0458	0.0471	0.0856	0.9508	0.9516	0.9092	53.78	53.86	47.36
NZD	0.0469	0.0515	1.1626	0.9408	0.9395	0.4296	56.02	54.24	50.19
PHP	2.3100	33.083	1.4331	0.9037	0.7227	0.8967	50.09	51.34	54.10
PLN	0.0416	0.0427	0.7436	0.9742	0.9745	0.6721	52.80	50.91	46.68
RUB	0.0350	0.0366	0.7531	0.8933	0.8931	0.5598	52.59	52.45	61.05
SEK	0.0696	0.0733	0.0892	0.9583	0.9567	0.9413	51.96	51.76	49.40
SGD	0.0297	0.0304	0.4070	0.9654	0.9655	0.7295	54.92	55.90	55.52
SKK	0.0705	0.0743	0.2045	0.9213	0.9175	0.7950	53.98	52.55	48.55
THB	0.3830	0.4147	0.5107	0.9732	0.9702	0.9643	43.93	53.09	59.95
TRY	2.3048	2.2142	2.0386	0.9694	0.9712	0.9752	52.62	52.56	52.20
TWD	0.0551	0.0713	0.5135	0.9073	0.8785	0.6511	51.74	53.93	55.31
USD	0.0603	0.0632	0.1257	0.9789	0.9782	0.9569	52.10	52.45	46.27
ZAR	0.3119	0.3893	3.1386	0.8861	0.8830	0.4957	52.60	50.89	48.63

Table 7: For the volatility surface implied by the thirty different currency options in our sample, the table reports out–of–sample average prediction errors across models. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). VAR corresponds to equation (9), with d selected via successive applications of the log–likelihood ratio test, starting with a maximum value of 10. RW stands for the random walk model that sets one–day's ahead implied volatility forecast equal to today's level. MSE is the mean squared error, R^2 is the coefficient of determination from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and MCP the mean correct prediction of the direction of change.

average mean squared errors are most of the times one half–and in some occasions one tenth or less–those of the RW model. Concentrating on the most liquid currencies (USD, JPY, GBP, CAD, CHF), the average MSEs of RMPQ (VAR) are 0.0603, 0.0758, 0.0346, 0.0368, 0.0489 (0.0632, 0.0800, 0.0365, 0.0412, 0.0502 respectively), a marked improvement over the performance of the random walk model (0.1257, 0.1826, 0.0719, 0.2923, 0.0442), except the CHF case where the latter outperforms only marginally. In all but 4 cases (CZK, KRW, PHP and TRY) our approach outperforms, in terms of MSEs, the random walk benchmark.

A similar conclusion is reached when the second measure, the R^2's from univariate regressions of observed implied volatilities on the model–predicted implied volatilities, are examined. There is a distinct improvement in explanatory power, ranging from -2.93% (CZK) to 51.11% (NZD), with an average of 15% across all currencies in our sample. In terms of the MCP measure, both the RMPQ and the VAR models perform better, by 2%–3% most of the times, than the 50% benchmark in correctly predicting the direction of change in the future IVS. Taking into account that the MCP (as all performance measures examined) is averaged across (in excess of) 1500 days for some currencies (AUD, CHF, JPY, GBP, USD, NOK, etc.), this constitutes a remarkable improvement in terms of correctly–predicted days.[5]

The prediction accuracy of our approach is also investigated in terms of multistep–ahead forecasts. Tables C.1–C.4 in Appendix C report the results of repeating the comparison for $2, 3, \ldots, 5$ days ahead. Our approach continues to outperform the random walk benchmark for up to 3–4 day ahead forecasts, although its absolute performance declines as expected. It is not until 5–day ahead forecasts are attempted that our approach is not better than a naive random walk in MSEs; however, even at this forecast horizon our approach is better at predicting the sign of change across the IVS than a toss of a coin.

Turning our attention to the comparison between our two proposed models, the results of Table 7 suggest that the RMPQ model outperforms the VAR model across all prediction measures for one–day ahead forecasts. This ranking persists at multistep–ahead forecasting performance as well, as the tables in Appendix C indicate. Although the VAR model is able to capture the reported dependence between factors across time (see Table 5), it seems

[5]Readers should note that MCP, the mean correct prediction of the direction of change in the IVS, cannot be defined for the random walk, RW model. What is reported in Table 7 and others that follow under MCP for the random walk, is the actual average percentage of out–of–sample days for which the implied volatility across the surface did not decrease. That is, in 46.27% of the 1645 days of the EUR/USD out–of–sample period, the average change in the IVS was greater or equal to zero.

that for forecasting purposes, adequately modeling the heteroscedasticity in the factor innovations is more important. By capturing the heteroscedasticity, the RMPQ model picks up autoregressive/moving–average dependence in factors more sharply than the VAR, and improves implied volatility predictions considerably.

In order to gain a better understanding of the RMPQ model's forecasting performance, Table 8 decomposes the out–of–sample prediction errors into moneyness and time–to–maturity categories for the USD/EUR case. Evidently, our approach performs better than the random walk model across all moneyness and maturity categories. Looking at the averages, the improvement of our approach is equally substantial across the moneyness dimension, and more pronounced for medium and long–term contracts.

Within categories, our approach performs the worst for contracts with 1 week time–to–maturity. This was expected in the light of evidence (see Hentschel (2003)) suggesting that there are measurement and liquidity–related biases with the implied volatilities of such short–term contracts. For all option types (calls and puts), our model has lower errors for ATM and OTM contracts than DOTM ones; however the relative ranking of the two appears to be related to maturity. Within each moneyness category, puts appear to be associated with lower MSEs than calls, except in the extremes of the time–to–maturity dimension. Thus it seems that the forecasting strength of our model comes mainly from medium and long–term ATM and OTM segments of the surface.

5 Conclusions

No single empirically observed deviation from the Black–Scholes–Merton option pricing framework has attracted more research effort than the nonconstant pattern of implied volatility versus the moneyness and time to maturity dimensions.

Recently, general equilibrium structural models have proposed economic justifications for the existence of an IVS; under such models, if latent factors entering the pricing kernel are persistent, then several patterns of an IVS can be observed, with substantial time–variation from period to period. This conclusion is consistent with the empirical observation that the coefficients of parametric specifications devised to fit an observed IVS change dramatically over time.

In this paper we propose an approach that jointly models the cross–sectional characteristics and the time–series dynamics of the IVS, which is based on simple time–series models of the evolution of few orthogonal statis-

Forecasting horizon: 1-day ahead. Prediction Measure: MSE														
	ATM		OTM				DOTM							
			Calls		Puts		Calls		Puts		Averages	
Maturity	RMPQ	RW	RMPQ	RW	RMPQ	RW	RMPQ	RW	RMPQ	RW	RMPQ	RW
1wk	0.2039	0.2539	0.2237	0.2727	0.2016	0.2557	0.2682	0.3011	0.2156	0.2714	0.2226	0.2710
1m	0.1111	0.1417	0.1185	0.1559	0.1109	0.1421	0.1421	0.1764	0.1158	0.1524	0.1197	0.1537
2m	0.0718	0.1129	0.0755	0.1238	0.0675	0.1143	0.0878	0.1405	0.0727	0.1247	0.0751	0.1232
3m	0.0529	0.1027	0.0552	0.1123	0.0482	0.1048	0.0644	0.1286	0.0517	0.1158	0.0545	0.1128
6m	0.0290	0.0932	0.0333	0.1018	0.0242	0.0959	0.0436	0.1195	0.0315	0.1081	0.0323	0.1037
9m	0.0240	0.0933	0.0290	0.1026	0.0200	0.0962	0.0403	0.1227	0.0302	0.1100	0.0287	0.1050
12m	0.0225	0.0933	0.0276	0.1033	0.0195	0.0963	0.0388	0.1257	0.0318	0.1115	0.0280	0.1060
18m	0.0210	0.0935	0.0236	0.1037	0.0183	0.0966	0.0327	0.1261	0.0287	0.1119	0.0249	0.1064
2y	0.0217	0.0938	0.0236	0.1040	0.0192	0.0969	0.0321	0.1267	0.0282	0.1124	0.0250	0.1068
3y	0.0262	0.0940	0.0259	0.1038	0.0254	0.0974	0.0279	0.1260	0.0235	0.1127	0.0258	0.1068
4y	0.0359	0.0941	0.0331	0.1036	0.0381	0.0978	0.0388	0.1257	0.0363	0.1132	0.0365	0.1069
5y	0.0495	0.0940	0.0439	0.1032	0.0548	0.0978	0.0518	0.1251	0.0555	0.1133	0.0511	0.1067
Averages	0.0558	0.1133	0.0594	0.1242	0.0540	0.1160	0.0724	0.1454	0.0601	0.1298	0.0603	0.1257

Table 8: For the volatility surface implied by options on the USD/EUR exchange rate from 4/9/2000 to 21/5/2007, the table provides a decomposition of the mean squared error (MSE) of Table 7 across the moneyness and time–to–maturity dimensions. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). RW stands for the random walk model that sets one–day's ahead implied volatility forecast equal to today's level. ATM, OTM and DOTM stand for at–the–money, out–of–the–money, and deep out–of–the–money respectively.

tical factors that are identified in the IVS.

At a first stage, instead of imposing a parametric specification of moneyness and time–to–maturity to explain the IVS cross–sectionally, as the popular practitioners' approach dictates, we derive directly from the data a number of orthogonal statistical factors that are shown to accurately reproduce the IVS observed on any given day. These statistical factors exhibit substantial time–variability, and are shown to have loadings with natural interpretations in the law of motion of the IVS.

At a second stage, we attempt to exploit the factors identified for forecasting purposes by modeling their evolution with simple, parsimonious econometric specifications. We demonstrate that our approach achieves a high–quality fit of the surface and of its evolution over time.

We examine the forecasting ability of our proposed approach out–of–sample in terms of standard prediction measures and in comparison to benchmarks used in practice. We find that our approach clearly outperforms in one–day ahead predictions of the IVS, and continues to outperform for forecasting horizons up to 3 days in the future. Careful examination of our approach's performance suggests that its forecasting power is at the medium and long–term ATM and OTM segments of the surface.

A Appendix: Determining the number of retained principal components

A variety of methods, from simple rules of thumb to proper statistical tests, have been suggested in the literature in order to determine the number of retained principal components in practical applications (see Jackson (1991) for a good review).

One popular rule of thumb is the Guttman–Kaiser criterion, also known as the mean eigenvalue rule of thumb, which retains only the principal components that correspond to eigenvalues larger that the mean of all eigenvalues. Another simple and practical rule is to keep the components which explain at least 95% of the total variance (as in the application of Litterman and Scheinkman (1988)).

We use the Guttman–Kaiser criterion in our application, as well as a testing procedure proposed by Velicer (1976) that examines the partial correlations of the residuals once m principal components are retained.

When $m = 0, 1, 2, \ldots, p$ principal components are retained, the partial correlation matrix of the residuals is given by

$$\mathbf{R}_m = \mathbf{D}^{-\frac{1}{2}} \varepsilon_m \varepsilon_m' \mathbf{D}^{-\frac{1}{2}} \tag{A.1}$$

where $\mathbf{D}$ is a diagonal matrix made up of the diagonal elements in $\epsilon_m \epsilon_m'$. By definition, $\mathbf{R}_0$ is the original correlation matrix of $\sigma_{p\times 1}$. If r_{ij} represents the i th row, j th column element of $\mathbf{R}_n$, then the Velicer (1976) statistic is given by

$$f_m = \sum_i \sum_{j \neq i} \frac{r_{ij}^2}{p\,(p-1)}. \tag{A.2}$$

The f_m statistic always lies in $[0, 1]$ and has a minimum in the range $0 < m < p - 1$, and this should be the number of principal components retained. The logic behind the test is that as long as f_m is declining, the partial covariances are declining faster than residual variances. Thus principal components will be retained until the ones left out represent more variance than covariance.

B Appendix: Unit root tests of identified factors

To determine whether the factor identified in the cross-sectional analysis are stationary, the augmented Dickey–Fuller and Phillips–Perron tests are conducted.

Currency Code	Augmented Dickey–Fuller applied to the scores of PC 1	PC 2	PC 3	PC 4	PC 5	PC 6	PC 7	PC 8	PC 9
AUD	0.0000	0.0000	0.0000	0.0000	0.0000				
BRL	0.0000	0.0000	0.0000	0.0000	0.0000	0.0041	0.0000		
CAD	0.0000	0.0000	0.0000						
CHF	0.0000	0.0000	0.0000	0.0000					
CLP	0.0000	0.0000	0.0150	0.0395	0.0042	0.0000	0.0000		
CZK	0.0000	0.0000	0.0000	0.0000					
DKK	0.0045	0.0294	0.0093						
GBP	0.0691*	0.0000	0.0000						
HKD	0.0128	0.0000	0.0000						
HUF	0.0116	0.0000	0.0000	0.0000	0.0000	0.0010	0.0026		
IDR	0.0045	0.0000	0.0544*	0.0000	0.0000				
INR	0.0029	0.0000	0.0000						
ISK	0.0230	0.0391	0.0626*						
JPY	0.0191	0.0000	0.0000	0.0000	0.0000				
KRW	0.0021	0.0000	0.0000	0.0000	0.0000				
MXN	0.0078	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0267	0.0000
MYR	0.0178	0.0135	0.0041						
NOK	0.0000	0.0000	0.0000	0.0000					
NZD	0.0290	0.0000	0.0000	0.0000	0.0000	0.0047	0.0000		
PHP	0.0000	0.0000	0.0000						
PLN	0.0119	0.0249	0.0000						
RUB	0.0000	0.0000	0.0000	0.0116	0.0316	0.0000	0.0000		
SEK	0.0000	0.0000	0.0000						
SGD	0.0000	0.0000	0.0000						
SKK	0.0000	0.0000	0.0000						
THB	0.0000	0.0000	0.0000						
TRY	0.0000	0.0000	0.0000	0.0000	0.0000				
TWD	0.0000	0.0000	0.0000	0.0095	0.0012	0.0239			
USD	0.0105	0.0000	0.0000						
ZAR	0.0239	0.0237	0.0000						

Table B.1: For the volatility surface implied by the thirty different currency options in our sample, the table reports p-values from the augmented Dickey–Fuller test under the null of a unit root process without drift, applied to the time–series of retained principal component scores. The number of retained principal components reported in the table is the greatest from the Guttman–Kaiser ($\bar{\lambda}$) and Velicer (f_n) criteria. The lag length for the test is selected by the Bayesian Information Criterion of Schwarz (1978) and the length of the time series for each currency is reported in Table 2.
An * denotes that the null cannot be rejected at the $\alpha = 5\%$ level.

Currency Code	Phillips–Perron applied to the scores of								
	PC 1	PC 2	PC 3	PC 4	PC 5	PC 6	PC 7	PC 8	PC 9
AUD	0.0000	0.0000	0.0000	0.0000	0.0000				
BRL	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		
CAD	0.0000	0.0000	0.0000						
CHF	0.0000	0.0000	0.0000	0.0000					
CLP	0.0000	0.0000	0.0049	0.0000	0.0000	0.0000	0.0000		
CZK	0.0000	0.0000	0.0000	0.0000					
DKK	0.0039	0.0296	0.0031						
GBP	0.0232	0.0000	0.0000						
HKD	0.0034	0.0000	0.0000						
HUF	0.0049	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		
IDR	0.0026	0.0000	0.1044*	0.0000	0.0000				
INR	0.0076	0.0000	0.0000						
ISK	0.0256	0.0381	0.0229						
JPY	0.0000	0.0000	0.0000	0.0000	0.0000				
KRW	0.0018	0.0000	0.0000	0.0000	0.0000				
MXN	0.0311	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
MYR	0.0236	0.0058	0.0010						
NOK	0.0000	0.0000	0.0000	0.0000					
NZD	0.0027	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		
PHP	0.0000	0.0000	0.0000						
PLN	0.0074	0.0033	0.0000						
RUB	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000		
SEK	0.0000	0.0000	0.0000						
SGD	0.0000	0.0000	0.0000						
SKK	0.0000	0.0000	0.0000						
THB	0.0000	0.0000	0.0000						
TRY	0.0000	0.0000	0.0000	0.0000	0.0000				
TWD	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000			
USD	0.0032	0.0000	0.0000						
ZAR	0.0083	0.0067	0.0000						

Table B.2: For the volatility surface implied by the thirty different currency options in our sample, the table reports p-values from the Phillips–Perron test under the null of a unit root process without drift, applied to the time–series of retained principal component scores. The number of retained principal components reported in the table is the greatest from the Guttman–Kaiser ($\bar{\lambda}$) and Velicer (f_n) criteria. The lag length for the test is selected by the Bayesian Information Criterion of Schwarz (1978) and the length of the time series for each currency is reported in Table 2. An * denotes that the null cannot be rejected at the $\alpha = 5\%$ level.

In the augmented Dickey–Fuller, under the null hypothesis the true underlying process of factor realisations $\widehat{F}_{m,t}$ is

$$\widehat{F}_{m,t} = \widehat{F}_{m,t-1} + \zeta_1 \Delta \widehat{F}_{m,t-1} + \zeta_2 \Delta \widehat{F}_{m,t-2} + \ldots + \zeta_l \Delta \widehat{F}_{m,t-l} + u_t \qquad \text{(B.1)}$$

with Δ the first difference operator, while under the alternative

$$\widehat{F}_{m,t} = \phi \widehat{F}_{m,t-1} + \zeta_1 \Delta \widehat{F}_{m,t-1} + \zeta_2 \Delta \widehat{F}_{m,t-2} + \ldots + \zeta_l \Delta \widehat{F}_{m,t-l} + u_t \qquad \text{(B.2)}$$

with $\phi < 1$. In Phillips–Perron, the null and alternative hypothesis are

$$\widehat{F}_{m,t} = \widehat{F}_{m,t-1} + \nu_t \qquad \text{(B.3)}$$

and

$$\widehat{F}_{m,t} = \phi \widehat{F}_{m,t-1} + \nu_t \qquad \text{(B.4)}$$

respectively.

C Appendix: Prediction accuracy of multi–step ahead forecasts

In the tables that follow, the out–of–sample prediction accuracy of our RMPQ and VAR models is assessed, for $\tau = 2, 3, 4, 5$ day–ahead implied volatility forecasts.

Currency	Two–day ahead forecasts								
	MSE			R^2			MCP		
Code	RMPQ	VAR	RW	RMPQ	VAR	RW	RMPQ	VAR	RW
AUD	0.0791	0.0849	0.1384	0.9788	0.9776	0.9631	53.68	52.97	45.95
BRL	0.9439	1.1486	1.3584	0.9155	0.9049	0.8906	52.94	52.24	47.74
CAD	0.0520	0.0603	0.2452	0.9147	0.9013	0.7228	54.64	53.41	47.23
CHF	0.0582	0.0605	0.0499	0.9243	0.9255	0.9401	53.32	52.49	46.86
CLP	0.1894	0.1923	0.5314	0.8986	0.8943	0.7913	50.36	49.29	60.55
CZK	0.1500	0.1519	0.1205	0.9004	0.9041	0.9350	51.41	51.00	50.65
DKK	0.0523	0.0549	0.2008	0.9802	0.9683	0.9453	53.08	52.88	48.61
GBP	0.0474	0.0506	0.0735	0.9736	0.9717	0.9606	52.75	52.61	47.77
HKD	0.0636	0.0650	0.4468	0.9564	0.9576	0.7907	51.31	52.54	54.63
HUF	0.1408	0.1458	1.0704	0.8604	0.8675	0.4769	53.18	52.18	45.97
IDR	0.3114	0.3924	0.6912	0.9298	0.9193	0.8748	54.28	52.84	48.95
INR	0.0997	0.1099	0.7039	0.9150	0.9163	0.6159	50.98	50.04	56.91
ISK	0.8834	1.1433	2.7850	0.6999	0.7250	0.6781	52.59	51.13	50.33
JPY	0.1124	0.1210	0.1903	0.9772	0.9761	0.9615	52.99	53.09	43.77
KRW	0.4080	0.4124	0.3012	0.9706	0.9709	0.9796	52.85	52.29	47.27
MXN	0.1718	0.1859	1.1304	0.9268	0.9262	0.6927	56.05	53.48	51.16
MYR	0.1893	0.1909	0.8293	0.5112	0.5870	0.3850	49.24	52.14	51.38
NOK	0.0681	0.0689	0.0981	0.9276	0.9328	0.9013	54.59	54.07	46.40
NZD	0.0766	0.0866	0.9175	0.9003	0.9017	0.4569	56.67	55.87	46.37
PHP	3.0487	43.6199	2.4852	0.8649	0.5828	0.8579	51.77	52.11	51.97
PLN	0.0678	0.0687	0.5946	0.9599	0.9617	0.7082	53.87	53.22	46.25
RUB	0.0408	0.0421	0.5757	0.8827	0.8849	0.5916	54.13	52.90	54.75
SEK	0.0950	0.0968	0.1045	0.9419	0.9436	0.9356	52.52	52.15	47.15
SGD	0.0427	0.0437	0.3282	0.9484	0.9504	0.7559	55.71	56.62	53.27
SKK	0.1017	0.1047	0.2049	0.8904	0.8922	0.7976	54.55	53.33	47.17
THB	0.5808	0.5996	0.6074	0.9596	0.9580	0.9577	43.78	53.28	57.26
TRY	3.1716	2.8775	2.7887	0.9585	0.9627	0.9661	53.24	53.10	50.50
TWD	0.0651	0.0694	0.4105	0.8873	0.8806	0.6759	51.38	53.39	53.62
USD	0.0894	0.0934	0.1402	0.9690	0.9682	0.9531	52.31	53.47	46.13
ZAR	0.5758	0.7634	2.7097	0.8099	0.8129	0.5091	54.60	49.97	47.66

Table C.1: For the volatility surface implied by the thirty different currency options in our sample, the table reports out–of–sample average prediction errors across models. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). VAR corresponds to equation (9), with d selected via successive applications of the log–likelihood ratio test, starting with a maximum value of 10. RW stands for the random walk model that sets two–day's ahead implied volatility forecast equal to today's level. MSE is the mean squared error, R^2 is the coefficient of determination from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and MCP the mean correct prediction of the direction of change.

Currency Code	MSE			R^2			MCP		
	Three–day ahead forecasts								
	RMPQ	VAR	RW	RMPQ	VAR	RW	RMPQ	VAR	RW
AUD	0.1045	0.1139	0.1422	0.9723	0.9701	0.9630	54.80	53.42	44.87
BRL	1.1966	1.4428	1.4634	0.8908	0.8807	0.8823	53.32	52.28	46.91
CAD	0.0675	0.0734	0.1950	0.8903	0.8817	0.7561	54.70	53.72	45.76
CHF	0.0662	0.0690	0.0542	0.9115	0.9152	0.9354	53.35	52.13	46.36
CLP	0.2354	0.2266	0.4375	0.8777	0.8740	0.8087	51.25	48.67	60.14
CZK	0.1790	0.1809	0.1481	0.8805	0.8881	0.9222	52.13	51.54	48.58
DKK	0.0631	0.0677	0.2026	0.9745	0.9618	0.9528	53.11	53.01	48.02
GBP	0.0596	0.0632	0.0732	0.9673	0.9649	0.9612	52.77	52.83	46.96
HKD	0.0843	0.0832	0.3351	0.9423	0.9463	0.8245	51.11	52.47	53.59
HUF	0.1945	0.1985	0.8184	0.8165	0.8318	0.5210	53.86	52.45	44.87
IDR	0.4258	0.5748	0.6711	0.9036	0.8876	0.8719	54.96	53.15	46.04
INR	0.1294	0.1435	0.5343	0.8832	0.8910	0.6645	53.45	50.56	55.93
ISK	1.0507	1.4320	2.2412	0.6614	0.6998	0.7058	52.59	51.13	49.66
JPY	0.1507	0.1639	0.1986	0.9698	0.9681	0.9609	53.11	53.35	43.25
KRW	0.4910	0.4902	0.3811	0.9646	0.9655	0.9740	53.49	52.61	46.19
MXN	0.2114	0.2210	0.8434	0.9098	0.9139	0.7365	58.04	55.15	49.64
MYR	0.2611	0.2468	0.6644	0.3504	0.4892	0.3734	48.86	51.91	49.68
NOK	0.0903	0.0913	0.1091	0.9026	0.9130	0.8942	55.59	53.32	45.96
NZD	0.1082	0.1253	0.6700	0.8543	0.8608	0.5021	58.22	54.82	46.05
PHP	2.1226	401.4787	2.4353	0.8494	0.3848	0.8623	52.64	52.33	51.09
PLN	0.0918	0.0931	0.4367	0.9455	0.9495	0.7594	55.96	52.38	43.68
RUB	0.0455	0.0479	0.4024	0.8744	0.8771	0.6333	55.04	53.18	51.62
SEK	0.1225	0.1206	0.1178	0.9231	0.9295	0.9305	52.46	52.71	47.49
SGD	0.0532	0.0540	0.2453	0.9340	0.9387	0.7932	55.69	56.82	51.24
SKK	0.1299	0.1322	0.1986	0.8599	0.8678	0.8046	55.27	53.89	46.53
THB	0.7907	0.7746	0.6957	0.9452	0.9440	0.9516	43.97	53.25	56.33
TRY	3.8264	3.3033	3.1154	0.9505	0.9571	0.9615	53.30	53.15	49.90
TWD	0.0749	0.0801	0.3060	0.8678	0.8605	0.7108	50.42	51.64	52.89
USD	0.1208	0.1272	0.1556	0.9581	0.9569	0.9487	52.70	53.98	45.33
ZAR	0.8549	1.2075	2.2300	0.7285	0.7434	0.5401	53.84	49.13	46.56

Table C.2: For the volatility surface implied by the thirty different currency options in our sample, the table reports out–of–sample average prediction errors across models. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). VAR corresponds to equation (9), with d selected via successive applications of the log–likelihood ratio test, starting with a maximum value of 10. RW stands for the random walk model that sets three–day's ahead implied volatility forecast equal to today's level. MSE is the mean squared error, R^2 is the coefficient of determination from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and MCP the mean correct prediction of the direction of change.

	Four–day ahead forecasts								
Currency	MSE			R^2			MCP		
Code	RMPQ	VAR	RW	RMPQ	VAR	RW	RMPQ	VAR	RW
AUD	0.1263	0.1378	0.1413	0.9666	0.9639	0.9641	55.57	53.67	44.22
BRL	1.4353	1.6877	1.5528	0.8658	0.8594	0.8748	53.45	51.76	46.99
CAD	0.0806	0.0845	0.1405	0.8694	0.8647	0.8063	53.54	53.32	45.36
CHF	0.0743	0.0780	0.0576	0.8977	0.9041	0.9313	53.96	52.37	45.96
CLP	0.2932	0.2739	0.3459	0.8526	0.8477	0.8294	50.95	47.50	60.71
CZK	0.2073	0.2082	0.1737	0.8605	0.8732	0.9102	52.41	51.67	47.51
DKK	0.0829	0.0900	0.1755	0.9715	0.9573	0.9566	53.23	53.05	47.98
GBP	0.0712	0.0750	0.0720	0.9613	0.9584	0.9622	53.09	52.99	46.74
HKD	0.1057	0.1009	0.2241	0.9277	0.9352	0.8678	50.63	50.46	55.74
HUF	0.2591	0.2638	0.5664	0.7603	0.7852	0.5934	53.66	50.98	44.90
IDR	0.5649	0.8077	0.6690	0.8715	0.8514	0.8658	55.11	52.68	43.94
INR	0.1581	0.1754	0.3663	0.8482	0.8679	0.7335	53.97	52.30	59.07
ISK	1.1620	1.6274	1.6539	0.6389	0.6874	0.7498	52.28	49.79	46.88
JPY	0.1886	0.2068	0.2046	0.9622	0.9601	0.9609	52.99	52.88	43.53
KRW	0.5683	0.5626	0.4528	0.9589	0.9604	0.9689	53.26	52.41	45.59
MXN	0.2540	0.2628	0.5621	0.8920	0.9010	0.7984	59.75	55.57	48.27
MYR	0.3404	0.3138	0.4966	0.1957	0.3830	0.3502	46.99	51.25	51.40
NOK	0.1122	0.1122	0.1188	0.8762	0.8949	0.8882	55.85	53.14	46.19
NZD	0.1356	0.1589	0.4122	0.8078	0.8224	0.5918	59.04	52.18	44.00
PHP	3.3095	588.3826	2.3952	0.8494	0.2505	0.8682	52.21	52.06	51.83
PLN	0.1196	0.1188	0.2865	0.9279	0.9357	0.8258	56.67	51.96	42.05
RUB	0.0490	0.0519	0.2264	0.8669	0.8717	0.7045	55.70	53.46	49.31
SEK	0.1504	0.1397	0.1271	0.9028	0.9188	0.9274	52.66	52.53	48.08
SGD	0.0621	0.0632	0.1588	0.9209	0.9282	0.8456	55.76	56.46	50.85
SKK	0.1561	0.1556	0.1879	0.8298	0.8475	0.8166	56.13	54.29	45.06
THB	1.0059	1.1303	0.7757	0.9307	0.9258	0.9459	44.26	53.05	56.20
TRY	4.4857	3.8264	3.5851	0.9431	0.9503	0.9551	53.75	53.22	49.20
TWD	0.0837	0.0878	0.1986	0.8498	0.8464	0.7651	49.40	49.15	55.52
USD	0.1493	0.1563	0.1666	0.9480	0.9471	0.9454	52.80	54.70	44.46
ZAR	1.1745	1.7958	1.7709	0.6423	0.6767	0.5937	54.00	49.30	46.89

Table C.3: For the volatility surface implied by the thirty different currency options in our sample, the table reports out–of–sample average prediction errors across models. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). VAR corresponds to equation (9), with d selected via successive applications of the log–likelihood ratio test, starting with a maximum value of 10. RW stands for the random walk model that sets four–day's ahead implied volatility forecast equal to today's level. MSE is the mean squared error, R^2 is the coefficient of determination from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and MCP the mean correct prediction of the direction of change.

	Five–day ahead forecasts								
Currency	MSE			R^2			MCP		
Code	RMPQ	VAR	RW	RMPQ	VAR	RW	RMPQ	VAR	RW
AUD	0.1465	0.1608	0.1399	0.9612	0.9579	0.9653	56.22	53.71	44.26
BRL	1.6750	1.9800	1.6390	0.8397	0.8370	0.8669	53.98	51.74	46.65
CAD	0.0915	0.0947	0.0845	0.8521	0.8493	0.8805	52.59	52.46	45.53
CHF	0.0837	0.0875	0.0619	0.8814	0.8922	0.9256	54.35	52.22	45.67
CLP	0.3719	0.3208	0.2566	0.8210	0.8206	0.8551	49.92	46.50	61.96
CZK	0.2356	0.2343	0.1978	0.8396	0.8588	0.8987	52.58	52.09	47.09
DKK	0.1104	0.1255	0.1306	0.9638	0.9485	0.9661	53.28	53.11	47.49
GBP	0.0832	0.0871	0.0707	0.9552	0.9517	0.9632	53.16	52.83	45.82
HKD	0.1287	0.1206	0.1130	0.9116	0.9232	0.9265	44.26	47.53	64.31
HUF	0.3210	0.3283	0.3140	0.7076	0.7424	0.7436	53.69	50.48	44.16
IDR	0.7094	1.0712	0.6760	0.8374	0.8135	0.8582	55.16	52.72	43.90
INR	0.1824	0.2070	0.1841	0.8152	0.8454	0.8573	53.54	49.87	66.99
ISK	1.3070	2.0432	1.0557	0.6082	0.6516	0.8408	52.67	49.51	45.98
JPY	0.2233	0.2447	0.2069	0.9552	0.9530	0.9615	53.02	53.13	43.26
KRW	0.6459	0.6379	0.5232	0.9531	0.9553	0.9639	53.55	52.42	45.83
MXN	0.2831	0.2970	0.2863	0.8780	0.8909	0.8918	61.38	55.93	45.64
MYR	0.4215	0.3831	0.3282	0.0819	0.2861	0.2670	44.73	49.17	57.41
NOK	0.1331	0.1320	0.1273	0.8495	0.8773	0.8833	56.13	52.68	46.18
NZD	0.1657	0.1929	0.1743	0.7552	0.7874	0.8023	58.22	52.63	44.06
PHP	3.0464	11107.3518	2.3808	0.8456	0.0211	0.8738	50.24	51.09	55.13
PLN	0.1481	0.1448	0.1369	0.9087	0.9216	0.9252	57.84	51.26	41.28
RUB	0.0525	0.0548	0.0470	0.8598	0.8675	0.9082	56.92	54.39	48.12
SEK	0.1816	0.1577	0.1372	0.8807	0.9099	0.9247	53.13	52.65	48.28
SGD	0.0700	0.0721	0.0712	0.9086	0.9184	0.9193	54.80	55.73	52.42
SKK	0.1822	0.1803	0.1787	0.7984	0.8262	0.8291	56.20	54.28	45.27
THB	1.2399	1.3166	0.8548	0.9146	0.8998	0.9402	43.88	52.60	56.45
TRY	5.1569	4.3875	3.9255	0.9355	0.9430	0.9501	54.03	53.65	48.76
TWD	0.0898	0.0929	0.0882	0.8382	0.8357	0.8603	43.73	43.31	63.69
USD	0.1755	0.1832	0.1747	0.9384	0.9378	0.9431	52.70	54.76	44.12
ZAR	1.5102	2.4375	1.3349	0.5474	0.6135	0.7144	55.16	50.14	46.17

Table C.4: For the volatility surface implied by the thirty different currency options in our sample, the table reports out–of–sample average prediction errors across models. RMPQ corresponds to equation (10) with (r, m, p, q) set equal to the order selected by the BIC criterion on estimation of the model on the first 100 observations of the IVS (starting with maximum values of $(10, 10, 4, 4)$). VAR corresponds to equation (9), with d selected via successive applications of the log–likelihood ratio test, starting with a maximum value of 10. RW stands for the random walk model that sets five–day's ahead implied volatility forecast equal to today's level. MSE is the mean squared error, R^2 is the coefficient of determination from a univariate regression of the observed implied volatilities on the model–predicted implied volatilities, and MCP the mean correct prediction of the direction of change.

References

ALEXANDER, C. (2001): "Principal component analysis of volatility smiles and skews," Working Paper, University of Reading.

BANK OF INTERNATIONAL SETTLEMENTS (2007): *OTC derivatives market activity in the second half of 2006*. Monetary & Economic Department, BIS.

BARTLETT, M. (1937): "The statistical conception of mental factors," *British Journal of Psychology*, 28(xx), 97–104.

BLACK, F., AND M. SCHOLES (1973): "The pricing of options and corporate liabilities," *Journal of Political Economy*, 81(3), 637–654.

CAMPA, J. M., AND K. P. CHANG (1995): "Testing the expectations hypothesis on the term structure of volatilities," *Journal of Finance*, 50(2), 529–547.

——— (1998): "The forecasting ability of correlations implied in foreign exchange options," *Journal of International Money and Finance*, 17, 855–880.

CANINA, L., AND S. FIGLEWSKI (1993): "The informational content of implied volatility," *Review of Financial Studies*, 6(3), 659–681.

CHRISTOFFERSEN, P., AND K. JACOBS (2004): "The importance of the loss function in option valuation," *Journal of Financial Economics*, 72(xx), 291–318.

CHRISTOFFERSEN, P., AND S. MAZZOTTA (2005): "The accuracy of density forecasts from foreign exchange options," *Journal of Financial Econometrics*, 3(4), 578–605.

CONT, R., AND J. DA FONSECA (2002): "Dynamics of implied volatility surfaces," *Quantitative Finance*, 2(1), 45–60.

DERMAN, E., AND M. KAMAL (1997): "The patterns of change in implied index volatilities," Working Paper, Goldman Sachs, Financial Strategies Group.

DIEBOLD, F., AND C. LI (2006): "Forecasting the term structure of government bond yields," *Journal of Econometrics*, 130(xx), 337–364.

DUMAS, B., J. FLEMING, AND R. E. WHALEY (1998): "Implied volatility functions: Empirical tests," *Journal of Finance*, 53(6), 153–180.

ENGLE, R. F. (1982): "Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation," *Econometrica*, 50, 987–1008.

GARCIA, R., R. LUGER, AND E. RENAULT (2003): "Empirical assessment of an intertemporal option pricing model with latent variables," *Journal of Econometrics*, 116(xx), 49–83.

GARMAN, M. B., AND S. W. KOHLHAGEN (1983): "Foreign currency option values," *Journal of International Money and Finance*, 2(3), 231–237.

GONÇALVES, S., AND M. GUIDOLIN (2006): "Predictable dynamics in the S&P 500 index options implied volatility surface," *Journal of Business*, 79(3), 1591–1635.

HARVEY, C., AND R. WHALEY (1992): "Market volatility prediction and the efficiency of the S&P 100 index options market," *Journal of Financial Economics*, 31(1), 43–73.

HENTSCHEL, L. (2003): "Errors in implied volatility estimation," *Journal of Financial and Quantitative Analysis*, 38(xx), 779–810.

HEYNEN, R. C., A. KEMNA, AND T. VORST (1994): "Analysis of the term structure of implied volatilities," *Journal of Financial and Quantitative Analysis*, 29(1), 31–56.

JACKSON, J. E. (1991): *A User's Guide to Principal Components*. John Wiley & Sons Inc., New York.

KNEZ, P., R. LITTERMAN, AND J. SCHEINKMAN (1994): "Exploration into factors explaining money market returns," *Journal of Finance*, 49(xx), 1861–1882.

KRYLOVA, E., J. NIKKINEN, AND S. VÄHÄMAA (2005): "Cross–dynamics of volatility term structures implied by foreign exchange options," Working Paper, European Central Bank.

LITTERMAN, R., AND J. SCHEINKMAN (1988): "Common factors affecting bond returns," Working Paper, Goldman Sachs, Financial Strategies Group.

MERTON, R. C. (1973): "The theory of rational option pricing," *Bell Journal of Economics*, 4(1), 141–183.

Mixon, S. (2002): "Factors explaining movements in the implied volatility surface," *Journal of Futures Markets*, 22(10), 915–937.

Nelson, C. R., and A. F. Siegel (1987): "Parsimonious modeling of yield curves," *Journal of Business*, 60(4), 473–489.

Rebonato, R. (2000): *Interest–Rate Option Models*. John Wiley & Sons Inc., New York.

Rubinstein, M. (1994): "Implied binomial trees," *Journal of Finance*, 49(3), 781–818.

Schwarz, G. E. (1978): "Estimating the dimension of a model," *Annals of Statistics*, 6(2), 461–464.

Sims, C. A. (1980): "Macroeconomics and reality," *Econometrica*, 48(xx), 1–48.

Skiadopoulos, G., S. D. Hodges, and L. Clewlow (1999): "The dynamics of the S&P 500 implied volatility surface," *Review of Derivatives Research*, 3(1), 263–282.

Velicer, W. F. (1976): "Determining the number of components from the matrix of partial correlations," *Psychometrica*, 41(3), 122–148.

Wilson, T. (1994): "Debunking the myths," *Risk*, 7(April), 67–73.

Xu, X., and S. J. Taylor (1994): "The term structure of volatility implied by foreign exchange options," *Journal of Financial and Quantitative Analysis*, 29(1), 57–74.