BANCA COMMERCIALE ITALIANA
Research Division

R99-9

EXCHANGE RATES STATISTICAL PROPERTIES IMPLIED IN FX OPTIONS

M. Esposito – E. Laruccia

BANCA COMMERCIALE ITALIANA
Research Division

EXCHANGE RATES STATISTICAL PROPERTIES IMPLIED IN FX OPTIONS

M. Esposito – E. Laruccia

R99-9

Editing: S. Papandrea

Milan, December 1999

INTRODUCTION

From FX options' prices one can infer some aspects of the statistical model the market is implicitly assuming for the exchange rate dynamics. The Finance literature has already shown how to compute implied volatilities and also how to use the apparently redundant cross-rates to extract implied correlations. In the following, we take a step further and use this information to derive the principal components of the exchange rate system. Once the covariance matrix is known, this is a simple mathematical operation which retains practical relevance because it allows to understand the relationships among currency areas as perceived by the market.

From the FX options' prices, we can derive also a term structure of implied volatilities and correlations. This topics, too, has been already investigated, proving the importance of modeling time-varying volatility. What we want to show, however, is that the volatility term structure slope might shed some light also on the implied auto-correlation structure of the exchange rates. This, too, is an important information about the way in which the FX dynamics is perceived by the market.

1. EXTRACTING IMPLIED VOLATILITIES

The market's uncertainty about the future value of a currency can be derived inverting the Garman-Kohlhagen (a variant of the Black-Scholes formula with the modification for the foreign interest rate) option pricing formula. Formally, the implied volatility ($\sigma*$) is the value of σ such that the theoretical price of an option is equal to its market price:

$$\sigma* : C_{th} = C_{mkt} \tag{1}$$

$$C_{th} = f(Y, T, X, r, \sigma)$$

Where:

$\sigma*$ = implied volatility;

C_{mkt} = call option market price;

C_{th} = call option theoretical price;

Y = currency price

T = time to maturity; X = strike price

r = risk free rates spread; σ = volatility of Y

It's is important to stress that this doesn't mean that traders believe in Black and Scholes assumptions, but their form solution simply serves as convenient mapping between implied volatilities and option premia. In addition, there's not a closed formula to derive $\sigma*$, so it has to be computed numerically. Fortunately, option prices are quoted directly in terms of implied volatility in the over- the-counter (OTC) market. In particular, prices refers to at the money forward straddles (i.e. one call plus one put, both European) at different maturities (f.e. 1, 3, 6 months and 1 year).

It is also possible to infer the implied volatility from official markets (like the Philadelphia Stock Exchange), but with several drawbacks with respect to the OTC market:

- in official markets, options are rarely traded at the money. This can induce some smile effects on the implied volatility, i.e. dependence on the in/out the moneyiness of the option;
- OTC markets are more liquid;
- for OTC options the time to expiration is constant, while official markets quote options at fixed dates, resulting in a variable time to expiration.

On the theoretical side, one possible objection to the use of implied volatility is that the Black- Scholes pricing model assumes that the underlying asset price has constant volatility. Theoretically and empirically, it has been proved that:

- the true value of an at-the-money option in presence of time varying volatility is approximately equal to the Black-Scholes valuation (Feisten, 1989), if the volatility used in the pricing formula is the average expected volatility for the time to maturity.

There is an high correlation between the implied volatility and the true volatility under different models with time-varying volatility (Heynen, Kemma and Vorst 1994).

2. THE TERM STRUCTURE OF IMPLIED VOLATILITIES

The different maturities of the options traded on the OTC market allow to derive a term structure of volatilities. This has been used to validate the different models of stochastic volatility. As we will see, however, the shape of the term structure of volatilities can tell something interesting about the auto-correlation structure of the (risk-neutral) stochastic process that the market believe the exchange rate will follow in the future. More precisely, we can understand from the term structure of volatilities if the market believes that the exchange rate will exhibit some sort of mean reversion or instead if it will behave as a random walk. We will use basic spectral analysis to see exactly what we mean by these statements.

Our time unit is constituted by the month. The implied volatility over a T months horizon is referred to the log changes of the exchange rate over such a time span. Let's then define:

$$y_{t,t+T} = \ln(Y_{t+T} / Y_t) \tag{2}$$

where Y denotes the exchange rate. We can re-write the log-change over the *T* months horizon as the sum of the *T* log-changes over the single months:

$$y_{t,t+T} = \sum_{j=0}^{T-1} y_{t+j,t+j+1} \tag{3}$$

As one can see, $y_{t,t+T}$ can be interpreted as a constant-weight moving average of the unit changes $y_{t+j,t+j+1}$. As such, it operates as a sort of "filter" that amplifies the low frequencies' contribution to the variance (i.e. the role of the components that displays the longer periodicity) and dampens the high frequencies. Imagine that the original unit changes $y_{t+j,t+j+1}$ are white noise (i.e. the exchange rate behaves like a random walk), then the spectrum is a flat line, constant at $\dfrac{\sigma^2}{2\pi}$.

Fig. 1 – GAIN FUNCTION OF $y_{t,t+T} = \sum_{j=0}^{T-1} y_{t+j,t+j+1}$ **, WITH T=2**

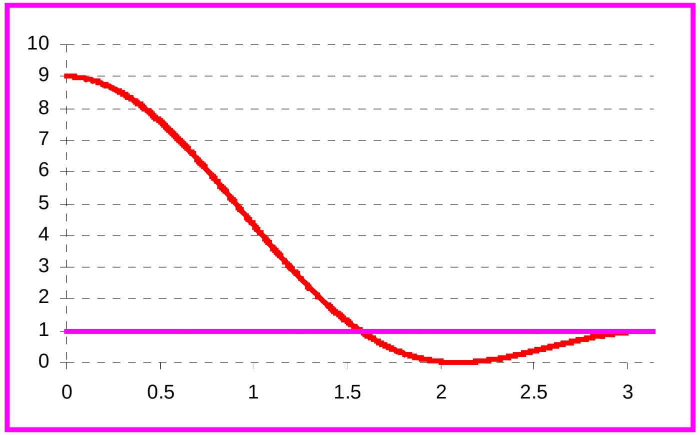

The spectrum of the multi-period changes will be equal to the spectrum of the unit changes (i.e. the flat line) multiplied by the square of the so-called gain function of the filter.

The square of the gain function, $|B(\omega)|$, of the filter depicted in eq. (3) is equal to (see Granger-Newbold, 1986):

$$|B(\omega)|^2 = \frac{1-\cos(T\omega)}{1-\cos(\omega)} \tag{4}$$

As one can see from the figures above and below, the filter of eq. (3) has the characteristics of a low-band pass filter and it is not a case that it can be used to estimate the trend components of a time series. The higher the number of terms inside the summation, the lower the frequencies which pass through the filter.

Fig. 2 – GAIN FUNCTION OF $y_{t,t+T} = \sum_{j=0}^{T-1} y_{t+j,t+j+1}$ **, WITH T=5 AND T=11**

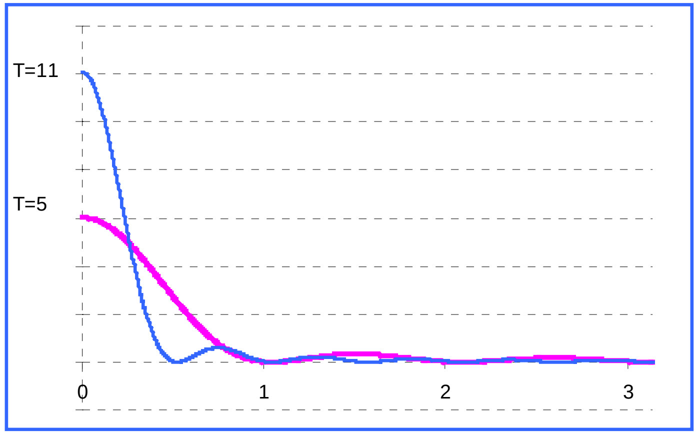

If our original time series is white noise and we divide the area below the spectrum of the filtered series by the area below the gain function of the filter, we should find as a result the variance of the original data. If, instead, we find a number which is bigger, this implies that the original time series was not white noise and was characterized by an higher contribution of low frequencies with respect to high frequencies. Viceversa, if we find a number which is smaller.

Clearly, in the options market we observe only the implied volatility, whose value is equal to (2 times) the area below the spectrum, and not the spectrum itself. So we are unable to validate, for example, the hypothesis of a white noise series over the unit period. However, since we know the shape of the filter and we know the implied volatilities for different values of *T (i.e. we observe the term structure of volatilities*), we can extract further information about the dynamic behavior of the statistical model that option traders are using for the exchange rate.

Let's consider the simplest alternative to the white noise model for the log-changes over the unit period, i.e. the AR(1) model.

$$y_{t,t+1} = \rho \cdot y_{t-1,t} + \varepsilon_t \tag{5}$$

Its spectrum is either positively or negatively sloped, depending on the negative or positive sign of the AR(1) parameter (see fig. 3).

Fig. 3 – SPECTRUM OF AN AR(1) PROCESS

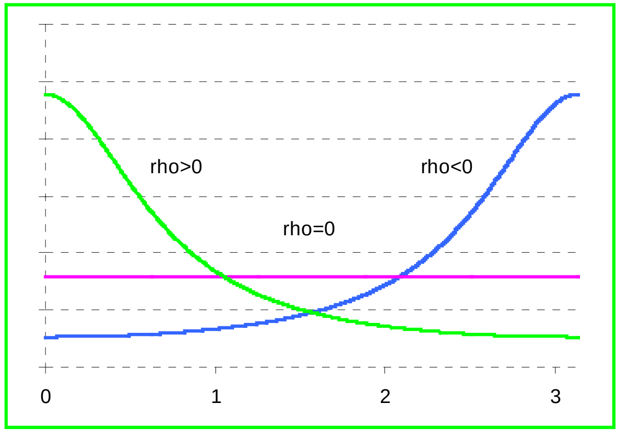

The application of a low-band pass filter to an AR(1) input process will have a very different impact on the output spectrum. Depending on the input time series being positively or negatively correlated, the variance of the output process can be magnified or dampened.

Denoting with F(y) the filtered series and with s(.) the spectrum as a function of the frequency, ω, we can have different cases:

$$\begin{cases} \int_0^\pi s(\omega, F(y))\, dw \Big/ \int_0^\pi |B(\omega, F(y))|^2\, dw \geq \int_0^\pi s(\omega, y)\, dw & \rho \geq 0 \\ \int_0^\pi s(\omega, F(y))\, dw \Big/ \int_0^\pi |B(\omega, F(y))|^2\, dw < \int_0^\pi s(\omega, y)\, dw & \rho < 0 \end{cases} \tag{6}$$

In words, the set of eq. (6) tells that the three basic shapes (flat, positively sloped, negatively sloped) of the volatility term structure can be obtained from a very simple mathematical model for the log-changes of the exchange rate, the AR(1).

We do not necessarily need a GARCH or a stochastic volatility model to obtain such a result. As an example, in the figure below, we show the shape of the term structure of volatilities that can be obtained under the hypothesis that the (log-changes of the) exchange rate are distributed according to an AR(1), with a 0.3 correlation coefficient.

Fig.4 – VOLATILITY TERM STRUCTURE IN AR(1) PROCESSES

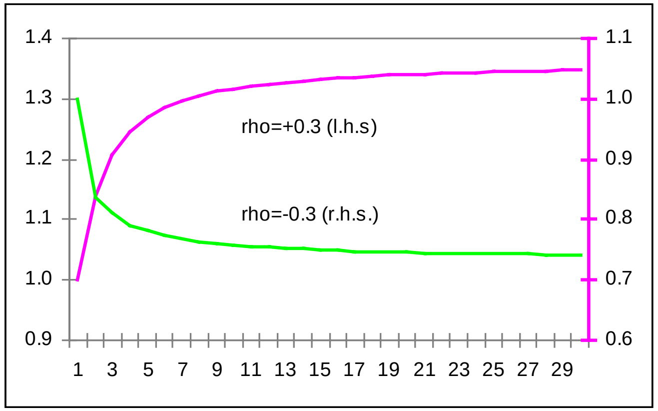

Note that in our simple case the one-period volatility is constant. The application of the Campa-Chang (1995) methodology to extract forward volatilities from the term structure would then be misleading in this case. In other words, an upward sloping volatility term structure does not necessarily predict a future increase in one-period volatilities.

These results may have deep consequences on the practical use of information extracted from the FX options, so that two fundamental questions need to be answered now. The first one is if it is possible that the (log-changes of the) exchange rate follow an AR(1) process. The second one is if there is a shape of the volatility term structure that we can consider as “normal”.

As for the first question, we have to recall that the volatility we use is derived from option prices. As such, the probability measure we have to refer to is the risk-neutral one. In this case, the drift of the stochastic process governing the exchange rate dynamics is given by the interest rate spread. An AR(1) process for the interest spread can then cause an AR(1) behavior in the exchange rate process. Note that an AR(1) process in the interest rate spread is not difficult to obtain, because in the no-arbitrage interest rate term structure models the spot interest rate is assumed to follow a mean reverting process (i.e. an AR(1) process with negative correlation). In those cases, the spread, too, will follow a mean-reverting process.

The answer to the first question partially answers the second question. If the interest rate spread is mean reverting, this implies that the (log-changes) exchange rate can be modelled as an AR(1) with negative correlation parameter. This implies that, “normally” the volatility term structure should be downward sloping. However, this seems more likely to happens for the exchange rates inside a common currency area, where there is a dominating country that set the pace of the economic and monetary cycles (and that has also a strong political influence on the smaller neighborhood). If we consider two countries belonging to different currency areas, there is no guarantee that the spread shows mean reverting properties. Consider the extreme case of the Russian rouble with respect to the US dollar.

In figures 5, we show the actual implied volatilities derived from OTC options on the exchange rate of the euro against four currencies (dollar, yen, pound, and Swiss franc). The source of the data are the Reuters multicontributor pages: EURxxx1MO, EURxxx3MO, EURxxx6MO, EURxxx1YO; where xxx stands for one of the four currencies’ codes. Data have been collected on the 12 of August, 1999 and are ask quotations.

Fig. 5- ACTUAL AND FITTED IMPLIED VOLATILITIES ON THE EURO EXCHANGE RATE (13/8/99)

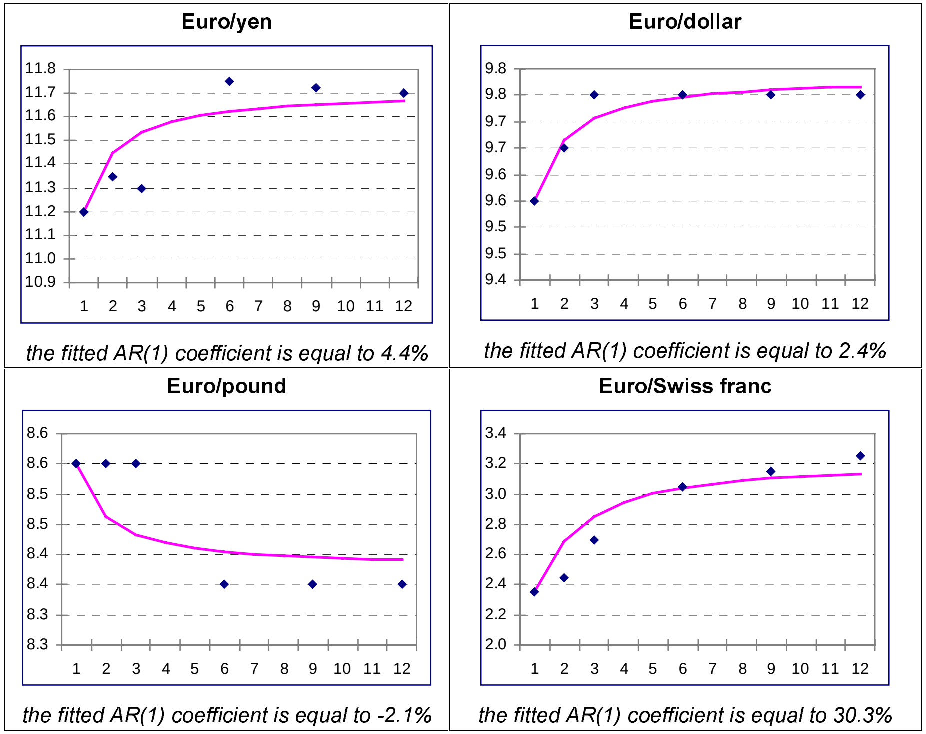

Actual market data are shown in the graph as a star point. The continuous line is instead the theoretical volatilities, calculated assuming that the log changes of the euro/xxx exchange rate follows an AR(1) process, with correlation parameter estimated (so as to minimise the square of the difference between actual and fitted implied volatilities).

As one can see, with the exception of the euro/Swiss franc, the implied autocorrelation coefficient is quite low and, differently from what we would have expected, there does not seem to be a prevailing negative slope shape. The only negative slope that we can find is the one relative to the euro/pound volatilities. Is this coherent with the analysis of the historical time series of the interest rate spreads and the exchange rates?

Table 1- INTEREST RATE SPREADS AND EXCHANGE RATE STATISTICAL PROPERTIES

8/86--8/99	**USD**				**D-Mark**		
	D-Mark	**Yen**	**Pound**	**Sw. Fr.**	**Yen**	**Pound**	**Sw. Fr.**
Interest rate spread (a)							
st. dev.	1.23%	1.18%	1.68%	1.77%	1.17%	1.29%	1.69%
autocorr. coeff.	12.8%	1.4%	11.4%	-6.7%	-13.1%	12.0%	-16.0%
Exchange rate (b)							
st. dev.	4.80%	5.47%	4.71%	5.24%	4.90%	3.53%	1.88%
Ratio of variances: (a)/(b)	*6.61%*	*4.64%*	*12.69%*	*11.45%*	*5.71%*	*13.40%*	*80.27%*
Exchange rate (risk adj.)							
Theor. autocorr. coeff.	0.84%	0.07%	1.44%	-0.76%	-0.75%	1.61%	-12.85%

In table 1, we considered the longest sample available of monthly data on the exchange rate and the spread between 1 month interbank interest rates (source British Bankers Association). We considered the D-mark as a proxy for the euro. We are interested in two variables. The first variable that is of interest is the sign of the auto-correlation in the interest rate spread. The second one is the ratio between the variances of the interest rate spread and the exchange rate. In order to have autocorrelation in the risk-adj. exchange rate process, the volatility of the spread has to be not negligible. In fact, if we have that Y=X+u and X is AR(1) with coefficient ρ, then the auto-correlation of Y will be equal to $\rho \sigma_X^2 / \sigma_Y^2$.

As one can see, the ratio between the two variances is in general small but not negligible. In the case of the Swiss franc/euro, it is equal to 80%. At least in this case, one should care about the impact of the spread autocorrelation. As far as the sign of the latter is concerned, we do not see a clear pattern to emerge from the data. The Swiss franc spread is negatively autocorrelated both when measured with respect to the euro and the dollar. The yen-euro spread is negatively autocorrelated, but not the yen-dollar. The pound-euro and the pound-dollar are both positively autocorrelated. The euro-dollar spread is positively autocorrelated.

If we calculate the autocorrelation coefficient induced on the risk-adjusted exchange rate by the autocorrelation in the interest rate spread, we find a value which is extremely low with the notable exception of the Swiss franc/euro exchange rate.

If we compare these data with the result we obtained from the options' market, we can see that the Swiss franc/euro volatility term structure is positively sloped (contrary to what historical data would suggest), but it is the only one which express a value of the implied correlation significantly higher than zero. The only inverted term structure is the one of the pound/euro exchange rate. Once again, the sign of the slope is opposite to the historical one, but in this case the contrast might be attributed to a sort of convergence game for the UK joining the EMU.

An analogous exercise can be made considering the term structure of implied volatilities on dollar-based exchange rates (see fig. 6).

Fig. 6- ACTUAL AND FITTED IMPLIED VOLATILITIES ON THE DOLLAR EXCHANGE RATE (13/8/99)

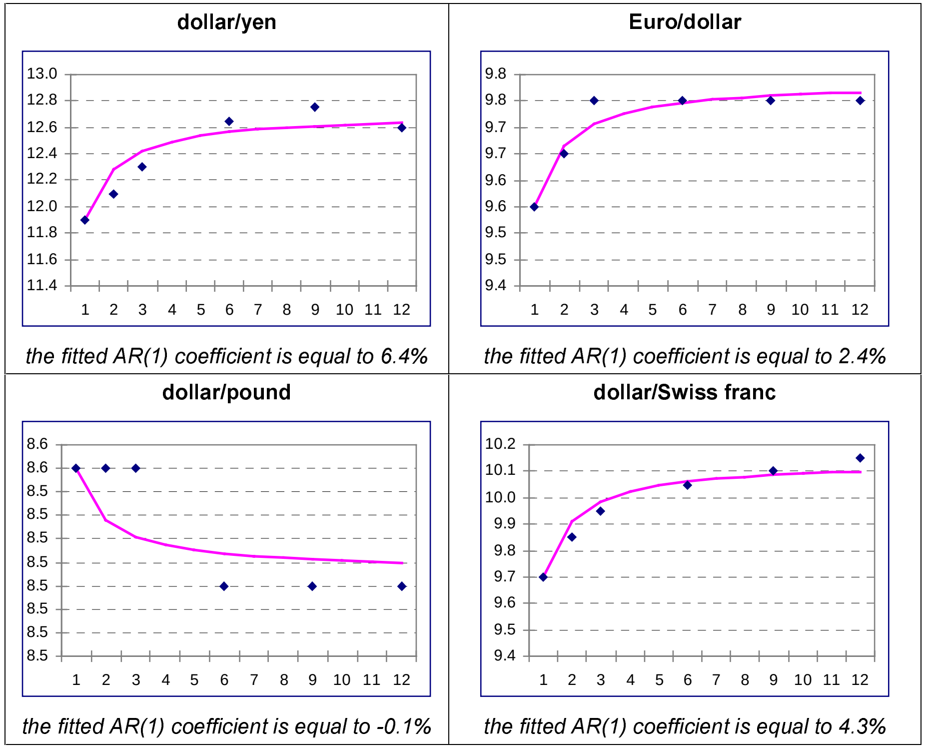

the fitted AR(1) coefficient is equal to 6.4% (dollar/yen)

the fitted AR(1) coefficient is equal to 2.4% (Euro/dollar)

the fitted AR(1) coefficient is equal to -0.1% (dollar/pound)

the fitted AR(1) coefficient is equal to 4.3% (dollar/Swiss franc)

3. IMPLIED FX CORRELATIONS

Another important statistical property of exchange rates that can be derived from FX options markets is the implied correlation between them. This is a unique feature of the exchange rates and its derivation is possible thanks to the apparently “redundant” information originated by the cross rates.

The correlations between currencies can be simply obtained by implied volatility (see, among others, C. Butler and N. Cooper, 1997). For example, let:Y_1 = euro/dollar; Y_2

= pound/dollar; Y_3 = euro/pound. Then, if we denote the log-change of each Y_i over the time unit by $y_i = \ln(Y_{i,t+1}/Y_{i,t})$, from the definition of the cross rates we have:

$$y_3 = \ln\frac{S_{1,t+1}/S_{2,t+1}}{S_{1,t}/S_{2,t}} = \ln\frac{S_{1,t+1}}{S_{1,t}} - \ln\frac{S_{2,t+1}}{S_{2,t}} = y_1 - y_2 \qquad (7)$$

It follows that:

$$\sigma^2_3 = \sigma^2_1 + \sigma^2_2 - 2\mathrm{cov}(y_1, y_2) = \sigma^2_1 + \sigma^2_2 - 2\rho_{1,2}\sigma_1\sigma_2 \qquad (8)$$

The implied correlation between y_1 and y_2 is given by:

$$\rho_{1,2} = \frac{\sigma_1^2 + \sigma_2^2 - \sigma_3^2}{2\sigma_1\sigma_2} \qquad (9)$$

With this formula we can easily compute the implied correlation over time. For example, Fig 7 shows the time series of the 1 month implied correlation between pound, euro and yen in the first half of 1999.

Fig. 7- IMPLIED CORRELATION OVER TIME (6/1/99 TO 29/07/99)

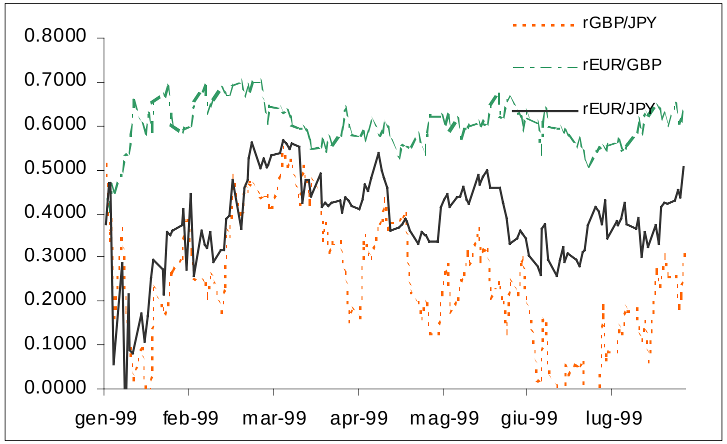

4. IMPLIED PRINCIPAL COMPONENTS

Principal Component analysis is often used to reduce the dimensionality of complex systems and to individuate those (groups of) variables that tend to "move" together. In our case, principal components can help understand the nature of the factors driving the international exchange rate system and, more precisely, to individuate currency areas.

Formally, once the implied correlation matrix **A** for *n* currencies has been derived, it is possible to extract its eigenvectors, x, and its eigenvalues λ working out the following equation:

$$\mathbf{A} \cdot x = \lambda \cdot x \tag{10}$$

Since the correlation matrix is always positive definite, there are always real solutions to (10). Principal components are, then, simply derived from the eigenvectors of the correlation matrix.

The application of this method to a large number of currencies needs a wide availability of quoted implied volatilities. We know that in the OTC market only a few cross-rate options are quoted. We then focus only on the four major currencies, considering their exchange rate with the dollar.

Let suppose that the implied correlation matrix of the exchange rates against the US dollar is given by:

Tab.2 - IMPLIED CORRELATION MATRIX (13/8/99)

	EUR/USD	GBP/USD	JPY/USD	CHF/USD
EUR/USD	1.0000	0.5651	0.4974	0.9760
GBP/USD	0.5651	1.0000	0.3038	0.5061
JPY/USD	0.4974	0.3038	1.0000	0.4344
CHF/USD	0.9760	0.5061	0.4344	1.0000

If we compute the eigenvector and the eigenvalues, we obtain:

Tab.3 - EIGENVECTOR AND EIGENVALUES

Eigenvalues	**1**	**2**	**3**	**4**
Value	2.6983	0.7084	0.5740	0.0192
% of variability	0.6746	0.1771	0.1435	0.0048
Cumulative %	0.6746	0.8517	0.9952	1.0000
Vectors :	**1**	**2**	**3**	**4**
EUR	0.5833	-0.0973	-0.3365	0.7329
GBP	0.4325	-0.4755	0.7640	-0.0565
JPY	0.3926	0.8637	0.3112	-0.0548
CHF	0.5645	-0.1358	-0.4541	-0.6758

From table 3, we can see that the first component explains the 67.46% of the total variance. The coefficient of each currency on that component are respectively 0.5833 (euro), 0.4325 (pound), 0.3926 (yen) and 0.5645 (Swiss franc). The first principal component can then be interpreted as the common factor behind the movements of the exchange rate systems. In simple words, in the 70% of the cases the four currencies are expected to contemporaneously depreciate or appreciate with respect to the dollar.

However, from the magnitude of the coefficients, we can also infer that the first factor seems to be determined mostly by the behavior of the euro and the Swiss franc. The fact that the euro and the Swiss franc constitute a currency area is partially confirmed by the analysis of the magnitude and the sign of their coefficients with respect to the second and the third principal components. The second principal component is dominated by the idiosyncratic movements of the yen, whose coefficient has, in fact, a sign opposite to the one of the other three currencies. The third principal component is instead dominated by the pound.

BIBLIOGRAPHY

Campa Josè M. and P.H. Kevin Chang, 1995, “Testing the Expectations Hypothesis on the Term Structure of Volatilities in Foreign Exchange Options”, *Journal of Finance*, 50, p. 529-47.

C. Butler and N. Cooper, 1997, “Implied exchange rate correlation and market perceptions of EMU”, *Bank of England Quarterly Bulletin*, November 1997.

S. Feinstein, 1989, “The Black-Scholes Formula is Nearly linear for all ATM Options”, Presentation at the AFA Conference, Dec 1990, Washington, D.C.

R. Heynen, A. Kemna and T.Vorst, 1994, “*Analysis of the term structure of implied volatilities*”, Journal of Finance and Quantitative Analysis, vol. 29, n. 1, march 1994, p. 31-56.