Original Article

An eigenvalue approach to risk regimes in currency markets

Received (in revised form): 30th June 2009

Pierre Lequeux

is head of currency management at AVIVA Investors which he joined in April 2010. Before this, Pierre was CIO currencies at fortis Investments (ABN AMRO Asset Management Legacy business) from March 1999 to March 2010. Pierre began his career in 1987 with Banque Nationale de Paris, London branch, where his last role; was head of the quantitative research and trading desk. Pierre, born 1961, is a graduate in international trade from ISTE (1983) and holds a diploma of the Forex Association (1990). He is also an associate researcher at the Centre for International Banking and Finance of Liverpool Business school. He is an active contributor to academic and practitioner and also acted as an Expert reviewer for the European Commission.

Murali Menon

is a senior currency manager at Aviva Investors. He worked in various telecommunications firms for over 10 years before transitioning to the financial industry in 2005. He holds a Master of Engineering degree in Computer Science from the Indian Institute of Science (1994) and a Master of Science degree in Mathematical Trading and Finance from the Cass Business School (2006).

Correspondence: Pierre Lequeux, Head of Currency Management, AVIVA Investors, London, UK
E-mail: Pierre@lequeux.org

ABSTRACT Over the past decades, both academics and market practitioners have produced much research on forecasting financial risk. In view of the levels of risk that the markets have operated over in recent times, this research is very topical. One growing area of investigation is the construction and use of risk indicators that aim to detect changes in market risk regimes. For investors, these have tangible applications as risk affects the accuracy of asset valuations, portfolio construction and risk management. In this article, in contrast to other studies that deal with cross-asset risk indicators, we propose a risk indicator that derives solely from implied volatilities of currency exchange rates. Our logic is that derivative markets have been shown to exhibit a degree of forward-looking element into their pricing. In addition, foreign exchange is the largest financial market, suggesting that a large amount of information is therefore reflected in the time series of exchange rate volatilities. In our article, we first propose an eigenvalue-based approach to construct a risk indicator that focuses on the reduction of the number of significant factors during periods of market stress. We then go on to study its informational content, and conclude on possible avenues of research and applications.

Journal of Derivatives & Hedge Funds (2010) **16**, 123–135. doi:10.1057/jdhf.2010.10

Keywords: exchange rates; leading indicators of risk; forecasting risk; risk aversion; principal component analysis; Kaiser criterion

INTRODUCTION

Over the past decades, there has been much research produced by both academics and market practitioners on the subject of forecasting financial market risk. The introduction of Risk-Adjusted Return on Capital (RAROC) to global financial markets by Bankers Trust in the 1970s¹ and the subsequent developments by JP Morgan² of the RiskMetrics (Value at Risk) technology prompted significant research in forecasting volatility, correlation and the higher moments of the distributions of financial time series. Clearly, research towards a global and standardised metric of *ex ante* portfolio risk makes sense and should be attractive to most investors. However, it is also true that the real value of such a metric resides in both the accuracy and the forward-looking element embedded in the estimates of covariance and other moments used in its computation. The caveat is that most financial times price series used as input are not time-invariant. Therefore, they cannot really be described by a single distribution. They are rather more likely to be represented by a complex mixture of joint distributions. Hence, one can argue that the value of most risk management technologies comes mainly from the ability to forecast and pinpoint possible change of regime in the financial markets, and thereby transition from one state distribution to another.

At first, most of the research effort was concentrated on developing a robust statistical methodology to estimate risk and correlation. There is a vast amount of research produced in the field of Autoregressive Control Heteroskedasticity/Generalised Autoregressive Conditional Heteroskedasticity modelling of volatility and correlations.³ In addition, techniques such as Markov Regime Switching

to detect financial market regime changes have been investigated, and applied for many years in various asset classes.^{4–7} More recently, one growing and avidly discussed area of research by market practitioners has been the construction and use of risk indicators aiming to detect changes in market risk regime. This is not a trivial issue for investors as structural changes in risk have direct implications on valuation of assets, portfolio construction and risk management. In view of the levels of risk under which financial markets have operated over recent times, the issue has become very topical. In our research, we propose a risk indicator that focuses solely on the implied volatilities of currency exchange rates. Our rationale is that derivative markets have been shown to exhibit some degree of forward-looking element into their pricing.⁸

Foreign exchange is the largest and most liquid financial market by far.⁹ It is also very effective in absorbing new information rapidly, suggesting that much of that information flow is reflected in the underlying time series of exchange rates volatilities. We also believe that the design approach undertaken by many banks that provide Global Risk Aversion Indicators (GRAI)¹⁰ may be flawed as they tend to incorporate too many inputs by focussing on instruments – from a variety of asset classes – that are highly correlated to risk. Even ignoring the possibility of curve fitting, this may lead to some perverse effects with regard to their optimality. Using too many risk inputs may have the effect of reducing the variance of the indicators and hence their informational content. In addition, it is likely that their efficiency varies from one period to another as the correlation between markets change. A GRAI would probably be of little use to an active currency manager when the typical institutional diversified portfolio has a

Figure 1: XY plot of equity market volatility and exchange rate delivered volatility.

higher domestic content or is greatly hedged. The orthogonal relationship between exchange rates volatility and the Chicago Board Options Exchange Volatility Index (VIX) in 1987 tends to support our view. Figure 1 shows the XY plot of the average 21-day delivered volatility of main currencies versus the US Dollar and the level of the VIX (Implied volatility of S&P 500). Clearly, the relationship has been quite correlated across recent times, whereas in 1987 the relationship was pretty much independent as highlighted in the chart. We contend that the recent positive and significant relationship is possibly explained by the increased globalisation of financial markets resulting in a higher content of foreign assets within the institutional portfolio. In 1987, the asset allocation of most pension funds and large investors tended to be highly geared

towards domestic assets, except for a few countries such as the United Kingdom. Therefore, when the US stock market crashed, resulting in large spikes in its volatility, there was little effect on the currency markets, which remained calm.

This, we believe, supports the need for market-specific indicators. In addition, we contend that designing a risk indicator specific to the Foreign Exchange markets may still be of a use to other asset class managers because of the global nature of Foreign Exchange markets. Not only do Foreign Exchange markets have great sensitivity to economic fundamentals, but their dynamics are also highly intertwined with many other asset classes. In this article, we use a Principal Components Analysis (PCA)-based approach to construct a risk indicator while

using foreign exchange implied volatilities as a source of information. Our methodology focuses on the reduction of dimensions, which occurs during periods of market stress, and the informational content derivable from the shift of one to another dimension.

DATA

In our analysis, we use all possible combinations of G10 currencies (USD, EUR, JPY, GBP, NOK, SEK, CHF, AUD, NZD and CAD). This results in a universe of 45 exchange rates. Our choice is driven by the fact that these currencies represent at least 75 per cent of the global turnover in the currency markets, according to the most recent survey (2007) from the Bank for International Settlements.⁹ It is therefore likely to provide us with a good representation of the information content driving the level of risk in the global currency markets. For each exchange rate, we calculate daily logarithmic returns for both the spot exchange rates and their 1-month associated implied volatilities. The period covered is from 14 March 1996 to 31 December 2008, or 3342 daily observations per time series. Our time series of spot exchange rates S_t^i and volatilities σ_t^i for the i th currency pair at time t are transformed as in equations (1) and (2), respectively.

$$R_t^i = \ln\left(\frac{S_t^i}{S_{t-1}^i}\right) \quad (1)$$

$$R_{\sigma_t^i}^t = \ln\left(\frac{1 + \sigma_t^i}{1 + \sigma_{t-1}^i}\right) \quad (2)$$

Although 1-month implied volatility is readily available for major currency exchange rates (that is, against the USD), this is not necessarily the case for less liquid exchange rates. We therefore calculated the 1-month implied volatility for the

less liquid exchange rates by using a triangulation method as shown in equation 3):

$$\sigma_{ij} = \sqrt{\sigma_i^2 + \sigma_j^2 - 2\rho_{ij}\sigma_i\sigma_j} \quad (3)$$

where σ_i represent the 1-month implied volatility of the USD exchange rate i and ρ_{ij} the 21-day rolling correlation of the daily logarithmic returns of the USD exchange rates involved in the non-USD exchange rate (for example, the implied volatility of GBP-JPY can be obtained from the volatilities of its split as GBP-USD and USD-JPY, for which option-implied volatilities are readily available).

As our estimation uses a rolling historical correlation of logarithmic returns of exchange rates, it has the effect of potentially introducing a lagging factor in our methodology beyond the window used in the analysis described in the sequel. We believe therefore that some of the predictive aspect of our risk regime indicator may be penalised because of this.

The summary statistics for the daily price returns series are presented in Tables 1 and 2. Not surprisingly the data are highly non-normal, as highlighted by the significant level of kurtosis and the P -values of the Jarque-Bera test, which reject the hypothesis of normality in all cases. In addition, the existence of significant autocorrelations at various lags demonstrates the likelihood of non-random walk behaviour. Some of the exchange rates behave in a mean-reverting way and others in a trending way. This is visually supported by Figure 2.

As seen in Figure 3 and Table 3, the implied volatilities have had several spikes, most notably in the last several months, as the credit crunch exacerbated global risk aversion. Interestingly, as shown in Table 4, the implied volatility series also exhibit negative autocorrelation levels,

Table 1: Summary statistics (USD pairs)

	<i>EUR- USD</i>	<i>USD- JPY</i>	<i>GBP- USD</i>	<i>USD- CHF</i>	<i>USD- CAD</i>	<i>AUD- USD</i>	<i>NZD- USD</i>	<i>USD- SEK</i>	<i>USD- NOK</i>
Mean	0.0000	0.0001	0.0000	0.0001	0.0000	0.0000	0.0001	0.0001	0.0000
Standard deviation	0.0059	0.0069	0.0053	0.0065	0.0051	0.0080	0.0080	0.0066	0.0068
Min	-0.039	-0.057	-0.037	-0.054	-0.051	-0.088	-0.067	-0.056	-0.054
Max	0.040	0.032	0.045	0.044	0.044	0.066	0.058	0.036	0.048
Skewness	0.07	-0.76	-0.12	-0.27	-0.17	-1.07	-0.50	-0.26	0.12
Kurtosis	2.45	4.80	4.00	2.88	10.58	16.99	6.71	4.65	5.85
JB(<i>P</i> -value)	835.53(0)	3529(0)	2239(0)	1194(0)	15 608(0)	40 791(0)	6411(0)	3054(0)	4767(0)
Observations	3335	3335	3335	3335	3335	3335	3335	3335	3335

Table 2: Autocorrelations of daily logarithmic price returns (USD pairs)

<i>Lag</i>	<i>EUR- USD</i>	<i>USD- JPY</i>	<i>GBP- USD</i>	<i>USD- CHF</i>	<i>USD- CAD</i>	<i>AUD- USD</i>	<i>NZD- USD</i>	<i>USD- SEK</i>	<i>USD- NOK</i>
1	0.027	0.018	0.046	-0.010	0.017	-0.023	0.028	0.021	0.013
2	-0.003	0.015	0.002	-0.004	-0.038	-0.007	-0.031	-0.053	-0.015
3	-0.003	-0.032	-0.041	-0.010	0.035	-0.038	-0.024	-0.012	-0.024
4	0.041	-0.012	0.039	0.025	0.032	0.061	0.014	0.040	0.040
5	-0.012	-0.004	-0.007	-0.001	-0.043	-0.048	-0.028	-0.048	-0.026
6	-0.025	-0.011	-0.004	0.012	-0.041	-0.011	-0.037	-0.013	-0.023
7	0.020	0.001	0.002	0.014	0.009	0.021	0.019	0.049	-0.005
8	0.029	0.020	0.012	0.019	0.029	0.024	0.019	0.058	0.031
9	-0.029	-0.015	-0.026	-0.017	-0.066	-0.054	-0.031	-0.049	-0.029
10	-0.005	0.043	-0.039	0.019	-0.020	0.032	0.029	-0.008	0.003

confirming the volatility clustering phenomenon described by Taylor.¹¹

CONSTRUCTION OF THE RISK REGIME INDICATOR

PCA has long been used both by market practitioners and academics to explain the multi-

dimensionality of variance structure within complex system. The application to risk analysis is not new. As an example, Loretan¹² investigated the use of PCA to generate market risk scenarios. In addition, the Bank of England¹³ has made use of PCA in their methodology to determine risk aversion. The logic behind PCA is to isolate specific factors (principal

Figure 2: Cumulative logarithmic returns of dollar exchange rates March 1996 – December 2008.

components (PC)) that are orthogonal to each other and which explain most of the variance of the system analysed. PCA is a technique invented by Karl Pearson,¹⁴ dating back to 1901. It transforms a system of possibly correlated variables into one with uncorrelated factors. The methodology is elementary: create the correlation matrix of the variables in question, compute the eigenvalue decomposition of the matrix and sort the eigenvectors in decreasing order of eigenvalue. Each eigenvector is orthogonal to the others, and each explains a proportion of the variation of the overall system given by the size of its eigenvalue. A variety of criteria exist to select an appropriate number of explanatory factors, and the one we use is the *Kaiser criterion*.¹⁵ The selection focuses only on those factors that explain variability in the system

by an amount at least equal to one original variable; in other words, on the factors whose corresponding eigenvalues exceed 1.

These orthogonal factors are often informative in a financial sense, that is, they lend themselves to meaningful interpretation.¹⁶ For example, Sløk and Kennedy¹⁷ interpret the first two components (derived from stock and bond markets in developed and emerging economies) as a measure of the contribution of industrial production to changes in investor risk aversion. Similarly, McGuire and Schrijvers¹⁸ show that the first PC drives most of the risk premium in 15 emerging markets.

In our research, we utilise a matrix of daily logarithmic returns for the 1-month implied volatilities of the 45 exchange rates derived from the combination of the currencies mentioned

Figure 3: 1-month implied volatility of dollar exchange rates March 1996 – December 2008.

Table 3: Summary statistics of the daily 1-month implied volatility changes

	<i>EUR- USD</i>	<i>USD- JPY</i>	<i>GBP- USD</i>	<i>USD- CHF</i>	<i>USD- CAD</i>	<i>AUD- USD</i>	<i>NZD- USD</i>	<i>USD- SEK</i>	<i>USD- NOK</i>
Mean	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Standard deviation	0.004	0.007	0.003	0.004	0.003	0.005	0.005	0.004	0.004
Min	-0.026	-0.087	-0.033	-0.036	-0.024	-0.048	-0.041	-0.031	-0.023
Max	0.040	0.128	0.038	0.051	0.032	0.067	0.055	0.049	0.029
Skewness	1.24	3.63	0.97	1.46	0.95	1.76	0.82	1.56	1.14
Kurtosis	14.39	84.84	23.99	19.83	19	39.88	25.43	18.3	9.89
JB(P-value)	29 656(0)	10 08 457(0)	80 604(0)	55 905(0)	50 730(0)	2 22 936(0)	90 361(0)	47 915(0)	14 338(0)
Observations	3334	3334	3334	3334	3334	3334	3334	3334	3334

above. The matrix size is also defined by the length of the rolling windows (61 days) we use while performing our PCA. We run the analysis

over the whole sample. Rather than focussing on the financial meaning of the PC, we take the number of eigenvalues that satisfy the Kaiser

Table 4: Autocorrelations of the daily 1-month implied volatility changes

Lag	EUR- USD	USD- JPY	GBP- USD	USD- CHF	USD- CAD	AUD- USD	NZD- USD	USD- SEK	USD- NOK
1	-0.023	-0.043	-0.037	-0.084	-0.138	0.039	-0.089	-0.097	-0.068
2	-0.094	-0.135	-0.009	-0.030	-0.056	-0.161	-0.046	-0.103	-0.051
3	0.035	-0.038	-0.050	-0.043	-0.006	-0.102	-0.015	0.032	-0.019
4	-0.016	-0.034	-0.025	-0.030	0.046	0.031	-0.054	0.062	-0.062
5	-0.024	-0.060	-0.046	0.002	0.014	0.037	0.094	-0.045	0.032
6	-0.079	-0.027	-0.013	-0.046	-0.098	0.009	-0.017	-0.085	-0.047
7	-0.035	-0.016	0.008	-0.033	0.019	0.084	-0.058	0.027	-0.074
8	-0.042	-0.006	-0.009	0.012	0.035	-0.107	-0.029	-0.024	-0.019
9	-0.042	-0.057	-0.056	-0.048	-0.043	-0.096	0.008	-0.078	0.010
10	0.090	0.076	0.043	0.030	0.058	0.028	0.025	0.090	0.078

criterion (that is, greater than one) and then calculate their mean eigenvalue as shown in Figure 4 and Table 5. Our logic relies on the contention that when implied volatilities rise, there is usually a loss of dimensionality in the risk structure expressed by those implied volatilities. Put simply, what market participants call ‘contamination of risk’, a combination of both rising volatility and correlation across assets, can be interpreted as a loss of diversification, leading to a reduction in the number of mutually independent and significant explanatory factors. Therefore, the loss of meaningful orthogonality in the risk structure may possibly apply as an indicator of market risk regime.

While performing our analysis, we observed a negative relationship between the number of eigenvalues and their mean: the higher the values of the eigenvalues satisfying the Kaiser criterion, the fewer of them, and vice-versa. This is not surprising: the correlation matrix is positive definite (its eigenvalues are positive); the sum of the eigenvalues is the trace of the matrix, which in this case is n , the order of the matrix. As the diagonal elements of the correlation

matrix are all ones, the mean of all the eigenvalues is one. Therefore, if some eigenvalues are much larger than one, there will be only a few of them.

MEAN EIGENVALUE AND MARKET RISK

In this section, we investigate if the previously calculated mean eigenvalue has any explanatory and predictive power over the VIX and the 21-day average delivered exchange rate volatility of the exchange rates under consideration. To do so, we use both visual inspection as well as the well-known Granger causality test.¹⁹ A simple formulation for Granger causality is as follows: a variable X is said to Granger-cause a variable Y if Y is forecast better with the histories of X and Y combined rather than using the history of Y alone. We can set this up as a regression with lags in X and Y :

$$Y_t = \alpha_0 + \sum_{i=1}^M \alpha_i Y_{t-i} + \sum_{i=1}^K \beta_i X_{t-i} + \varepsilon_t \quad (4)$$

Figure 4: 61-day mean eigenvalue rolling value.

Table 5: Summary statistics of mean of significant eigenvalues

	<i>Mean eigenvalue</i>
Mean	4.16
Standard deviation	0.60
Min	2.98
Max	6.71
Skewness	1.22
Kurtosis	2.29
Observations	3302

We then perform an F -test to test the null hypothesis that the β_i are zero. If they are significant, then we can say that X Granger-causes Y . It is useful to check if Y Granger-causes

X as well; if it does not, we may infer some amount of informational flow from X to Y .

One of the problems with regressions using closing values of VIX and currencies is that these are not synchronous. The consequences of the use of non-synchronous data have been well documented.^{20–22} The VIX closes at 16:15 New York time; whereas the implied volatility close-of-day values we use are at 16:00 London time. This fact is sometimes ignored by market practitioners, leading to spurious conclusions of the VIX as a leading indicator. To mitigate this problem of synchronicity, we have used weekly sampling of the previously calculated daily time series.

In our analysis, we find that there is a relationship between the mean eigenvalue and the VIX as evidenced by the bivariate density

Figure 5: VIX and mean eigenvalue.

plot (Figure 5) at a lag zero. We observe that this relationship is multimodal, with the greatest peaks concentrated around the smaller means of eigenvalues (3.5–4.25). This also corresponds to the lower values of the VIX. Conversely, we observe that as the mean of eigenvalues increases (in effect, because of the higher values of the first few eigenvalues), the VIX is also notably higher. To demonstrate that is not merely a visual effect, we superimposed a LOWESS fit²³ on the bivariate plot, which shows that the level of the VIX is increasing as a function of the mean of eigenvalue.

Likewise, the bivariate density plot of the mean eigenvalue and the mean 21-day delivered volatilities of the exchange rates show that the higher mean of eigenvalues correspond to heightened delivered risk as well (Figure 6).

While performing our analysis, we also found that this relationship persisted at a sub-sample

level (that is, if we were performing the same visual inspection from smaller sample drawn from the original sample).

Although the above results are interesting from the perspective of classification of the risk regime, they only highlight the quasi-linear relationship that exists from time to time in the global financial market risk. We note, further, that the modes are also spread vertically in the plot. For each level of the mean of the significant eigenvalues, there are several possible levels of the VIX or delivered Foreign Exchange Volatility. We surmise that this indicates differing risk periods over time. Possibly because of the economic cycle or changes in international portfolio allocation, what the markets considered a fair value for risk was expressed at different numbers of explanatory factors. This bears further investigation in a later article.

Figure 6: Average G10 delivered volatility and mean eigenvalues.

Table 6: Granger causality *F*-statistics (*P*-values)

	<i>Lag</i> =1	<i>Lag</i> <=2	<i>Lag</i> <=3	<i>Lag</i> <=4
VIX → EIGENMEAN	0.0738 (0.7860)	0.3218 (0.7250)	0.2135 (0.8871)	2.4139 (0.0350)
EIGENMEAN → VIX	5.4146 (0.0203)	2.9298 (0.0541)	3.0375 (0.0286)	2.2498 (0.0480)

Probably a more interesting question is: what is the predictive power of our risk measure in relation to both delivered volatility and VIX. Table 6 shows the result of the Granger-Causality test with respect to the VIX. Clearly, we can see that, in our sample, the causality works in favour of the mean eigenvalue – it provides some degree of insight into the future level of the VIX at lags of 1–4 weeks.

Table 7 shows the results of the test with respect to the delivered volatility of Foreign Exchange rates. Here, the result is less striking as

it seems that at a lag of 2 weeks delivered volatility Granger-causes the Mean Eigenvalue.

It may well be that a shift in risk dimensionality in Foreign Exchange markets, as expressed by our risk indicator, operates under the same dynamics as one observes in the stock market, but that the Foreign Exchange (FX) option market is faster at adjusting its pricing because of its relatively lower cost of access. On the other hand, the linkage between the historical volatility and the change in the mean eigenvalue derived from the implied volatility

Table 7: Granger causality F -statistics (P -values)

	$Lag=1$	$Lag \leq 2$	$Lag \leq 3$	$Lag \leq 4$
HISTVOL $\rightarrow$ EIGENMEAN	2.3011 (0.1298)	4.2663 (0.0144)	2.8739 (0.0356)	3.4522 (0.0044)
EIGENMEAN $\rightarrow$ HISTVOL	2.3385 (0.1267)	1.1390 (0.3208)	0.7814 (0.5046)	1.0613 (0.3807)

probably will add to the ongoing debate as to whether implied volatility is or is not, an unbiased predictor of future volatility. It remains that even if the indicator may fail on the latter, it still provides some gauge in classifying under which risk regime financial markets operate.

FINAL REMARKS

In this article, we investigated the relationship between the loss of dimensionality in global foreign exchange market risk and the level of market risk. We used implied volatilities and PCA analysis for that purpose, creating a risk indicator specific to currency markets. We found that our indicator had explanatory power in both FX market and US stock market. We also found that the bivariate distribution of the mean Eigenvalue with the risk index (VIX or delivered FX volatility) is multimodal, suggesting the presence of regimes over time. A fruitful line of research would be into the determination of these regimes and their possible use in enhancing the utility of the mean of the significant eigenvalues as a risk proxy. Furthermore, we found that the indicator had some predictive power as shown in the result of the Granger-causality test but that this was not true in the case of the delivered currency risk.

Another possible line of investigation is how such a measure of risk could be used in the

context of portfolio construction. It has been shown elsewhere that the assumption that all pair-wise correlations between assets in a portfolio are equal at any time period results in a more robust estimation and forecast of portfolio risk.²⁴ As the correlation matrix can be reconstructed from its eigenvalue decomposition, one could use our risk indicator to obtain a more robust correlation matrix.

And, finally, it is useful to study the effect of our indicator on various investment styles. In general, investment styles may be termed as concave (those that rise when the underlying asset value falls) or convex (those that rise when the underlying asset values rise.) For example, the trend trading style is convex, whereas the carry trading style is concave. It is well known that the carry style is essentially a short volatility play, and therefore performs worse in periods of risk aversion. The utility of our risk indicator for this, and similar strategies, will be dealt with in a subsequent article.

REFERENCES

- 1 James, C. (1996) RAROC Based Capital Budgeting and Performance Evaluation: A Case Study of Bank Capital Allocation. Philadelphia, PA: Wharton Business School. Working Paper Series no. 96-40.
- 2 J.P. Morgan, (1996) *RiskMetrics Technical Document*, 4th edn. New York: J.P. Morgan bank in association with Reuters.

- 3 Engle, R. (2002) New Frontiers for ARCH Models. New York: New York University. Working Paper no. FIN-02-037.
- 4 Turner, C.M., Startz, R. and Nelson, C.R. (1989) A Markov model of heteroskedasticity, risk, and learning in the stock market. *Journal of Financial Economics* 25: 3–22.
- 5 Perez-Quiros, G. and Timmermann, A. (2000) Firm size and cyclical variations in stock returns. *Journal of Finance* 55(3): 1229–1262.
- 6 Bansal, R., Tauchen, G. and Zhou, H. (2004) Regime shifts, risk premiums in the term structure, and the business cycle. *Journal of Business & Economic Statistics* 22(4): 396–409.
- 7 Clarida, R.H., Sarno, L., Taylor, M.P. and Valente, G. (2006) The role of asymmetries and regime shifts in the term structure of interest rates. *Journal of Business* 79: 1193–1224.
- 8 Malz, A.M. (1997) Option-implied Probability Distributions and Currency Excess Returns. New York: Federal Reserve Bank of New York. Staff Report no. 32.
- 9 Bank for International Settlements. (2007) *Triennial Central Bank Survey: Foreign Exchange and Derivatives Market Activity in 2007*. Basel, Switzerland: BIS.
- 10 Kumar, S.M. and Persaud, A. (2003) Pure contagion and investors' shifting risk appetite: Analytical issues and empirical evidence. *International Finance* 5: 401–426.
- 11 Taylor, S.J. (1986) *Modelling Financial Time Series*. Chichester, UK: John Wiley and Sons.
- 12 Loretan, M. (1997) Generating market risk scenarios using principal component analysis: Methodological and practical considerations. In: *The Measurement of Aggregate Market Risk*. CGFS Publications, Bank for International Settlements, Vol. 7, pp. 23–60, <http://www.bis.org/publ/ecsc07.htm>.
- 13 Bank of England. (2006) Markets and operations, Bank of England Quarterly Bulletin, Spring. Vol. 46, No. 1, Bank of England, ISSN 0005–5166. February.
- 14 Pearson, K. (1901) On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* 2(6): 559–572.
- 15 Kaiser, H.F. (1960) The application of electronic computers to factor analysis. *Educational and Psychological Measurement* 20: 141–151.
- 16 Coudert, V. and Gex, M. (2008) Does risk aversion drive financial crises? Testing the predictive power of empirical indicators. *Journal of Empirical Finance* 15(2): 167–184.
- 17 Sløk, T. and Kennedy, M. (2004) Factors Driving Risk Premia. Paris, France: OECD Economics Department. Working Paper no 385.
- 18 McGuire, P. and Schrijvers, M. (2003) Common factors in emerging market spread. *BIS Quarterly Review* (December): 65–78.
- 19 Granger, C.W.J. (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3): 424–438.
- 20 Acar, E. and Lequeux, P. (1996) Dynamic strategies: A correlation study. In: C. Dunis (ed.) *Forecasting Financial Markets*. London, UK: Wiley, pp. 93–123.
- 21 Aggarwal, R. and Park, Y.S. (1994) The relationships between daily US and Japanese equity: Evidence from spot versus futures markets. *Journal of Banking and Finance* 18: 757–773.
- 22 Lin, W.L., Engle, R.F. and Ito, T. (1994) Do bull and bears move across borders? International transmission of stock returns and volatility. *The Review of Financial Studies* 7(3): 507–538.
- 23 Cleveland, W.S. and Devlin, S.J. (1988) Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association* 83: 596–610.
- 24 Engle, R.F. and Kelly, B.T. (2009) Dynamic Equicorrelation. New York: New York University. Working Paper no. FIN-08-038.