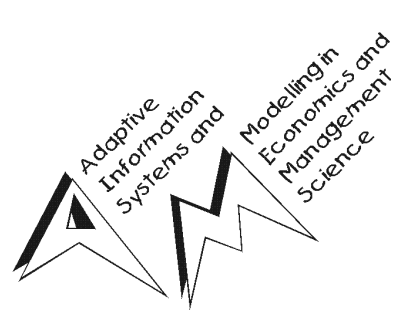Working Paper Series
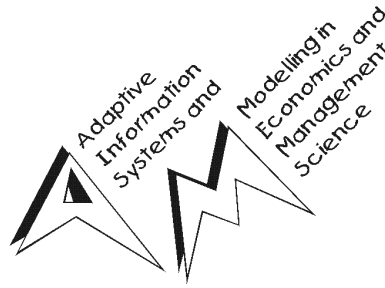
Cointegration and Exchange Market Efficiency: An Analysis of High Frequency Data

Adrian Trapletti
Alois Geyer
Friedrich Leisch

Working Paper Series

August 1999

SFB
'Adaptive Information Systems and Modelling in Economics and
Management Science'

Vienna University of Economics
and Business Administration
Augasse 2–6, 1090 Wien, Austria

in cooperation with
University of Vienna
Vienna University of Technology

http://www.wu-wien.ac.at/am

Cointegration and Exchange Market Efficiency: An Analysis of High Frequency Data

Adrian Trapletti [a,1], Alois Geyer [a], and Friedrich Leisch [b]

[a] *Department of Operations Research, Vienna University of Economics and Business Administration, Augasse 2-6, A-1090 Vienna, Austria*

[b] *Department of Statistics and Probability Theory, Vienna University of Technology, Wiedner Hauptstr. 8-10, A-1040 Vienna, Austria*

Abstract

A cointegration analysis on a triangle of high frequency exchange rates is presented. Market efficiency requires the triangle to be cointegrated and the cointegration term to be a martingale difference sequence. We find empirical evidence against market efficiency for very short time horizons: The cointegration term does not behave like a martingale difference sequence. In an out-of-sample forecasting study the cointegrated vector autoregressive (VAR) model is found to be superior to the naive martingale. Finally, a simple trading strategy shows that the VAR also has a significant forecast value in economic terms even after accounting for transaction costs.

Key words: Cointegration, Exchange rates, High frequency data, Market efficiency
JEL classification: C32,F31,G14

1 Introduction

The development of cointegration theory has motivated a large number of empirical studies investigating the efficiency of foreign currency exchange (FX) markets. Two major lines of research can be discerned. One investigates the relationship between spot and forward exchange rates which are required to

---

[1] *Correspondence to*: Adrian Trapletti, Department of Operations Research, Vienna University of Economics and Business Administration, Augasse 2-6, A-1090 Vienna, Austria, phone ++43-1-313364561, fax ++43-1-31336708.

be cointegrated as a necessary condition for the unbiasedness of market expectations. The second line examines the relations across different currency markets. According to this perspective efficiency implies that no FX rate should be cointegrated with any linear combination of other rates against which it floats freely. If cointegration holds there are long-run relationships among FX rates which tend to eliminate any deviations from equilibrium. Whereas cointegration implies that such deviations are useful in predicting future FX rates, in an efficient market nothing should be useful for prediction; in other words, FX rates should evolve as a vector martingale.

The conclusions obtained from empirical cointegration tests using (mainly) daily data are highly controversial, however. Whereas Baillie and Bollerslev (1989, 1994) have argued that FX main rates are cointegrated, Copeland (1991), Diebold et al. (1994), and Sephton and Larsen (1991) found that the evidence for cointegration is much less strong. The purpose of this paper is to provide new empirical evidence on the issue of FX market efficiency using high-frequency data. However, it is not merely the nature of the data, but an implication of cointegration, so far largely ignored, that motivates our analysis.

Any definition of market efficiency is based, in some way or another, on the notion that asset prices fully reflect relevant information. However, in practical terms, it is necessary to allow for some time to elapse such that new information can be processed by the market and be incorporated into prices. The availability of high-frequency data offers new perspectives for market efficiency tests. It allows to focus on a property implied by market efficiency that cannot or can hardly be accounted for using, e.g., daily data.

We make use of the so-called triangular identity, stating that the ratio of two FX rates, called the main rates, must be equal to the cross rate. In other words, direct trading of the cross rate should be equivalent to carrying out the trade through the main rates. For daily data this condition is almost the definition by which the cross rate is calculated from the main rates. This relationship forces the three FX rates to be cointegrated, provided that each rate is integrated. However, looking at higher frequencies, the triangular identity does not continuously hold. A necessary condition for cointegration is the stationarity of the residuals of the cointegration regression (the so-called cointegration term). Market efficiency, however, requires the cointegration term to be a martingale difference sequence! Viewed from this perspective, a new and stronger (yet not sufficient) condition for market efficiency can be derived: The cointegration term must not have any discernible structure in the sense of a martingale difference, otherwise information is not fully and immediately reflected in the rates.

This setup has considerable advantages over efficiency tests based on low-

frequency data. As Hakkio and Rush (1991) note, long spans of data rather than a large number of observations are required to give tests for cointegration much power. In addition, Diebold et al. (1994) remark that cointegration tests are sensitive to the assumption regarding the presence of a drift in the data. As we will show below, tests of cointegration on a FX rate triangle are not affected by these issues. Efficiency tests can be based on the properties of the cointegration term, which are largely unaffected by the length of the available time span and the drift assumption.

Focusing on a FX triangle the situation in an efficient market can be summarized as follows:

- There exists exactly one cointegration relation within a FX rate triangle.
- The cointegration term behaves like a martingale difference sequence.
- It is not possible to develop a model with predictive performance significantly better than a naive martingale.

Thus, an obvious way to reject market efficiency is to develop a FX rate model with predictive performance significantly better than a naive martingale. Jong et al. (1995) found that the cross rate has predictive power for the main rates and vice versa. Furthermore, Bolland and Connor (1995) report arbitrage opportunities investigating a FX rate triangle.

Our analysis is based on the high frequency FX rate triangle USD-DEM, USD-JPY, and DEM-JPY. In order to take the effects of seasonal volatility into account, we apply a deseasonalization procedure similar to the method proposed by Dacorogna et al. (1993). In particular, we apply a volatility based time scale to the price generating process. Based on the work of Johansen (1991) we, thereafter, estimate a vector autoregressive (VAR) model and analyze the resulting cointegration term. We compare the out-of-sample forecasting performance of the VAR model to that of a martingale. Finally, the economic value of the VAR forecasts is evaluated in terms of two trading strategies.

In the next section we introduce high frequency FX rates and briefly summarize the deseasonalization procedure and its effects. The results of the VAR analysis appear in Section 3. In Section 4 we outline the experimental design and present the results of the forecasting and trading experiments. Section 5 offers a summary and conclusions.

2 High Frequency Exchange Rates

2.1 Data

The basic data for this study are the exchange rate quotes for USD-DEM, USD-JPY, and the cross rate DEM-JPY. The data set covers the period from October 1, 1992 to September 30, 1993 on a tick-by-tick basis and contains 1,472,241 (USD-DEM), 570,813 (USD-JPY), and 158,978 (DEM-JPY) data records. Each record consists of the time the record was collected, the bid and ask price, the identification of the reporting institution, and a validation flag.

To construct an equally spaced time series, we take the most recent valid price record as a proxy for the current price record. The price values are computed as the average of the log of bid and ask prices, the returns as the price changes over some fixed time interval, and the volatility as the absolute values of the returns. Note that the sampling frequency is arbitrary and strongly influences the results. In order to keep the intra-daily character of the time series, but not getting a bid-ask spread of the same order as the returns, we use one hour time steps.

2.2 Time Scale

One of the major characteristics of high frequency data is the strong intra-week and intra-day seasonal behavior of the volatility. This is shown, e.g., by the autocorrelation function of the volatility for the USD-DEM rate in figure 1.

A data generation process with strong seasonal distribution patterns cannot be stationary. Therefore, controlling these seasonalities before fitting any time series model should improve the overall model quality.

A very promising approach to filter these seasonal patterns is to apply a new time scale to the price generating process similar to the one suggested by Dacorogna et al. (1993). In probability theory, this is formally equivalent to subordinated process modelling of the observed process. In this model the returns follow a subordinated process in physical time and are non-stationary. However, they follow the stationary parent process in another time scale which we call operational time scale.

The first step towards obtaining such a time scale is to cut the weekends. We call this time scale business time. Unfortunately, this step does not remove all seasonalities. Mainly the hour-of-the-day effect is still highly significant. To
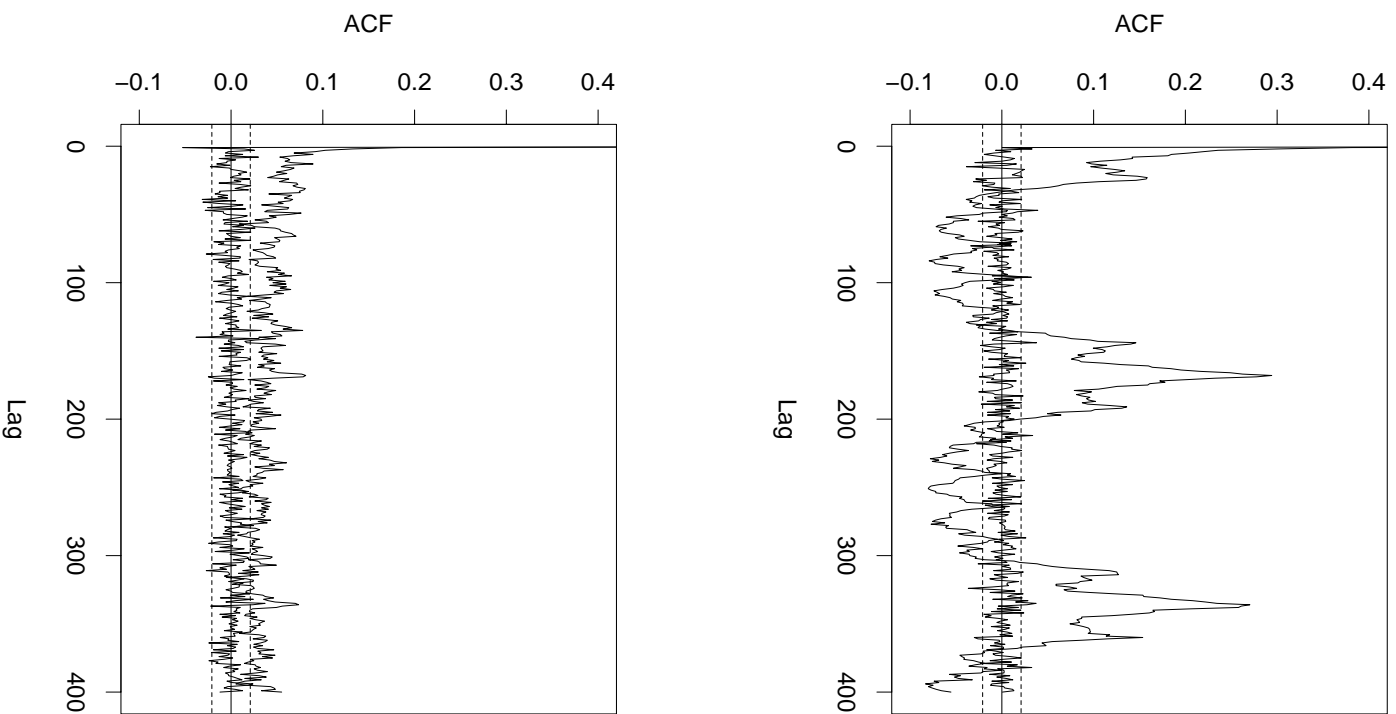
Fig. 1. Correlogram for USD-DEM returns (inside asymptotic 95% confidence bounds for white noise) and volatilities (outside confidence bounds) in physical time (above) and in operational time (below) for the period from October 1, 1992 to September 30, 1993.

5

Table 1

Augmented Dickey-Fuller test statistics with the 5% significance levels for the observations 1 to 4387.

Single unit root [a]			Second unit root [b]			
USD-DEM	USD-JPY	DEM-JPY	USD-DEM	USD-JPY	DEM-JPY	$\tau_{0.95}$
0.04	0.00	0.01	744.46	1037.23	819.64	1.95 [c]
1.23	0.02	0.44	744.91	1037.31	820.94	2.86 [d]

[a] no lagged differences

[b] one lagged difference

[c] without drift

[d] with drift

control for the remaining seasonal patterns, we apply another time scale to the business time scale which is computed by stretching highly volatile market periods, whereas less volatile periods are shortened. A more detailed description of the computation of the operational time scale is given in Appendix A. We sampled from the three series equally spaced in operational time with a sampling frequency corresponding to one hour time intervals. When the volatility in the market is high, these one hour time intervals in operational time correspond to about 15 minutes in physical time. On average, however, both time intervals have equal length.

The effect of changing the time scale can be seen in figure 1. Although the conditional heteroscedasticity is still present, most of the seasonal effects have been removed. Since the results for the two other rates are very similar to those of the USD-DEM rate, they are not reported here. It is important to note that the change of time scale does not affect our conclusions: a fair game cannot be made unfair by sampling it at certain stopping times, i.e., if the FX rates are a vector martingale, then optional sampling cannot improve the prediction quality over the martingale model. From this section onwards we use the operational time scale.

3   Cointegration Analysis

To establish the order of integration of the time series, unit root tests are computed. The results are reported in table 1. Based on the Augmented Dickey-Fuller tests (e.g., Cromwell et al., 1994) all series are found to be $I(1)$ which is hardly surprising.

We then estimate a k'th order VAR for our system of three series by the maximum likelihood procedure following Johansen (1991). The ordering of

Table 2
Order selection using the AIC and FPE criteria.

	VAR order (k)						
	1	2	3	4	5	6	7
AIC	-42.353	-42.379	-42.387	-42.390*	-42.388	-42.384	-42.383
FPE	-41.352	-41.377	-41.385	-41.388*	-41.386	-41.383	-41.381

*Minimum

the variables is (USD-DEM, USD-JPY, DEM-JPY). Allowing for linear trends and for a varying number of independent unit roots the VAR system can be written as

$$\Delta \boldsymbol{X}_t = \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \boldsymbol{X}_{t-i} + \boldsymbol{\Pi} \boldsymbol{X}_{t-k} + \boldsymbol{\mu} + \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\boldsymbol{\varepsilon}_t$ are $\boldsymbol{NID(0, \Lambda)}$ with the 3×3 covariance matrix $\boldsymbol{\Lambda}$. Further parameters are the 3×1 vector $\boldsymbol{\mu}$, the 3×3 matrices $\boldsymbol{\Gamma}_i$, and the 3×3 matrix $\boldsymbol{\Pi}$. The rank of the latter is equal to the number of cointegrating vectors.

We select the lag length using the Akaike information and final prediction error (AIC and FPE) criteria. Therefore, the unrestricted model (1) with Rank($\boldsymbol{\Pi}$) = 3 is estimated using the first half of the sample having 4387 observations. The results of order selection are presented in table 2. AIC and FPE both yield an optimal lag length of $\hat{k} = 4$.

3.1 Cointegrating Vectors

The hypothesis that there are r independent cointegrating vectors, i.e., $0 <$ Rank($\boldsymbol{\Pi}$) = $r < 3$ can be expressed as

$$\boldsymbol{\Pi} = \boldsymbol{\alpha}\boldsymbol{\beta}', \tag{2}$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are $3 \times r$ matrices. Furthermore, the hypothesis about the absence of a linear trend in the process is expressed as

$$\boldsymbol{\mu} = \boldsymbol{\alpha}\boldsymbol{\beta}_0', \tag{3}$$

or alternatively $\boldsymbol{\alpha}_{\perp}'\boldsymbol{\mu} = \boldsymbol{0}$, where $\boldsymbol{\alpha}_{\perp}$ is a $3 \times (3-r)$ matrix of vectors chosen orthogonal to $\boldsymbol{\alpha}$. Then we can write

$$\boldsymbol{\alpha}\boldsymbol{\beta}'\boldsymbol{X}_{t-k} + \boldsymbol{\mu} = \boldsymbol{\alpha}\boldsymbol{\beta}^{*'}\boldsymbol{X}_{t-k}^*,$$

where $\boldsymbol{\beta}^* = (\boldsymbol{\beta}', \boldsymbol{\beta}_0')'$ and $\boldsymbol{X}_{t-k}^* = (\boldsymbol{X}_{t-k}', 1)'$.
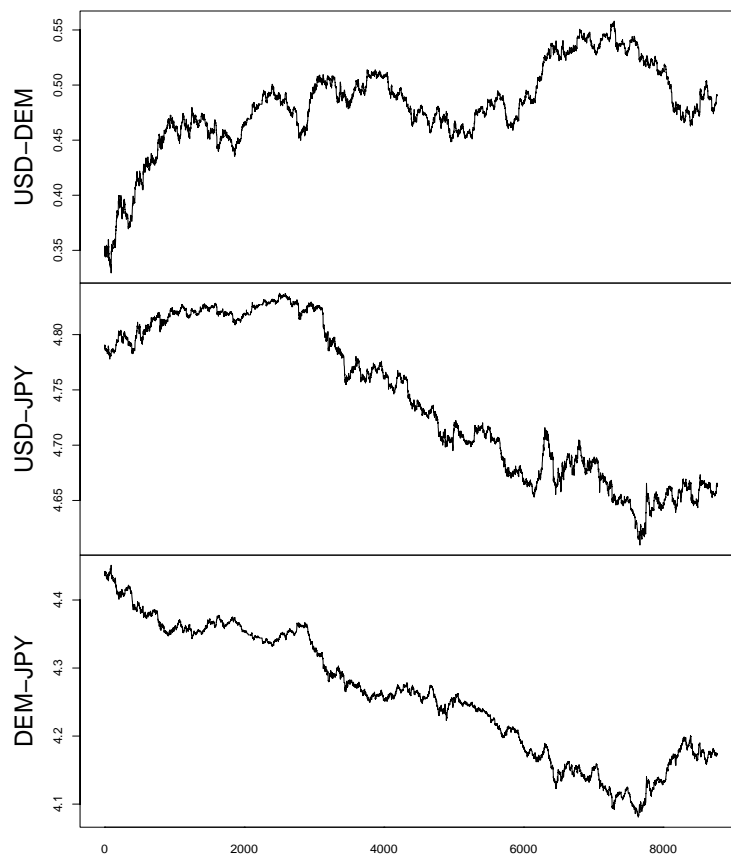
7

Fig. 2. The price series sampled equally spaced in operational time with a sampling frequency corresponding to one hour time intervals for the period from October 1, 1992 to September 30, 1993.

Table 3
Trace and maximum eigenvalue test statistics for various values of r with the 5% significance levels.

	trace	trace (0.95)	λ_{max}	λ_{max} (0.95)
$r \leq 2$	3.35	9.09	3.35	9.09
$r \leq 1$	14.01	20.17	10.66	15.75
$r = 0$	608.98	35.07	594.97	21.89

On the basis of the plots of the series (see figure 2) a model without a linear trend, i.e., under hypothesis (3) was assumed. The results of the trace and maximum eigenvalue test statistics are given in table 3 (cf. Johansen and Juselius, 1990). The support for the existence of exactly one cointegrating vector is strong. Moreover, assuming the absence of a linear trend is data con-

8

Table 4
The estimated cointegrating vector $\widehat{\boldsymbol{\beta}^*}'$. normalized on USD-DEM.

USD-DEM	USD-JPY	DEM-JPY	Constant
1.000	-0.999	1.000	-0.005

sistent. The likelihood ratio test statistic is LR $= 5.05$ which is asymptotically $\chi^2(2)$ and, thus, not significant.

Table 4 presents the normalized estimate of $\boldsymbol{\beta}^*$. The cointegrating vector $\boldsymbol{\beta}^*$ can be interpreted as an error correction mechanism. According to our hypothesis, all rates of the FX triangle should enter the error correction mechanism with the same weight. Given the chosen normalization we consider the hypothesis

$$\boldsymbol{\beta}^* = (1, -1, 1, 0)' \, \boldsymbol{\varphi}, \tag{4}$$

where $\boldsymbol{\varphi}$ is a real-valued parameter. The likelihood ratio test is given by LR $=$ 6.89 which should be compared with the quantiles of the $\chi^2(3)$ distribution. It is not significant and, therefore, hypothesis (4) is accepted with $\hat{\boldsymbol{\varphi}} = 2.612 \cdot 10^3$.

In the remainder of the paper hypotheses (1)-(4) with $\hat{k} = 4$ and $\hat{r} = 1$ are maintained. The estimates of the other parameters are presented in Appendix B.

3.2 Cointegration Term

Under the efficient market hypothesis the cointegration term should behave like a martingale difference. In figure 3 the autocorrelation function for the cointegration term $(1, -1, 1)\boldsymbol{X}_{t-k}$ is plotted.

While the process is apparently $I(0)$, it is immediately clear that this process is not uncorrelated. This implies that non-trivial forecasts of the FX rates are possible which violates the efficient market hypothesis. The relatively rapid decay of the autocorrelation coefficients indicates the purely intra-day character of this phenomenon.

The average speed of adjustment towards equilibrium is determined by the corresponding estimate of the adjustment parameters

$$\hat{\boldsymbol{\alpha}} = (-0.164, 1.764 \cdot 10^{-4}, -0.343)'.$$

The estimated coefficients indicate a much faster reaction for the DEM-JPY cross rate and for the USD-DEM rate than for the USD-JPY rate. In other words, arbitrage trades are typically not carried out using the USD-JPY rate.
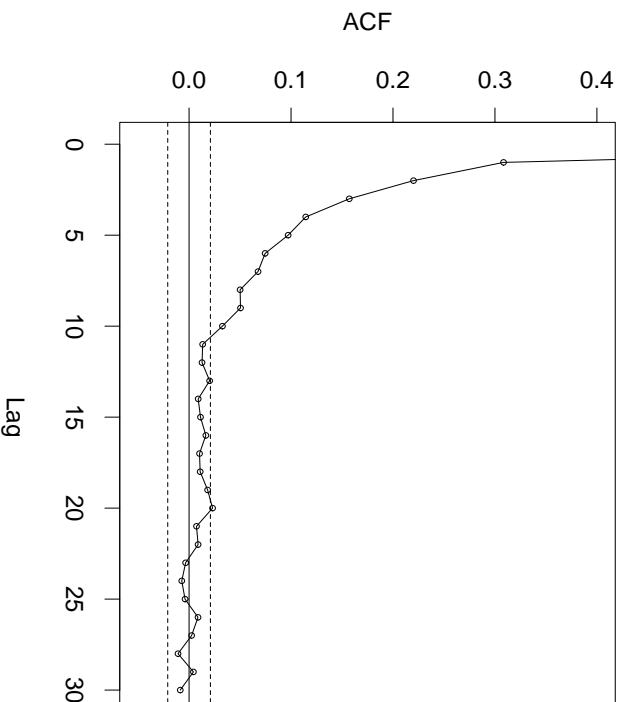
Fig. 3. Correlogram for the cointegration term $(1, -1, 1) \mathbf{X}_{t-k}$ with asymptotic 95% confidence bounds for white noise estimated from observations 1 to 4387.

Since the ultimate test for market efficiency is the predictive power of the model, we do not consider any further hypotheses regarding the parameters.

4 Prediction

4.1 Experimental Design

The out-of-sample forecasting exercise is performed on the second half of the sample, i.e., starting with observation 4387 while fixing the parameters at their estimated values. At each time period 1-, 2-, 4-, and 8-step ahead forecasts are computed by iterating forward through time. We compare these forecasts with the naive martingale model. The horizons correspond to about 15 minutes, 30 minutes, 1 hour and 3 hours in physical time if the seasonal volatility in the market is high, i.e., in the afternoon (Greenwich Mean Time). The average forecasting horizon in operational time, however, is the same as in physical time.

We evaluate the forecasts in terms of the mean squared prediction error (MSPE) and the mean absolute prediction error (MAPE). Since the main

goal is to compare the VAR with the naive martingale, we only report the relative performance of the VAR compared to the martingale. These accuracy measures are, however, parametric in the sense that they rely on the desirable properties of mean and variance. Therefore we also apply a distribution free procedure. A measure having this desirable property is the percentage of forecasts in the right direction which will be called the direction quality (DQ). In conjunction with the DQ the signal correlation (SC) between the forecasting signal and the actual price signal is computed.

In order to avoid decisions based on point estimates of the forecast accuracy, we test for the equality of two forecasts as proposed by Diebold and Mariano (1995). We give a brief explanation of how this test works:

Consider two scalar forecasts producing the errors $\{e_{1t}, e_{2t}\}, t = 1, \ldots, T$. The null hypothesis of the equality of the expected forecast performance is $E[g(e_{1t})] = E[g(e_{2t})]$, or $E[g(d_t)] = 0$, where $d_t = g(e_{1t}) - g(e_{2t})$ and $g(\cdot)$ is some loss function. It is natural to base a test on the sample mean $\overline{d} = T^{-1} \sum_{t=1}^{T} d_t$ of the loss-differential series. If the loss-differential series is weakly stationary and short memory, we have the asymptotic result $\overline{d} \overset{d}{\to} N(\mu, \sigma^2)$, where $\mu = E[g(d_t)]$, $\sigma^2 = T^{-1} \sum_{h=-\infty}^{\infty} \gamma(h)$, and $\gamma(h)$ is the autocovariance function of $\{d_t\}$. The obvious test statistic for testing the null is then $S = \overline{d}/\hat{\sigma}$, where $\hat{\sigma}$ is a consistent estimator of σ. We also computed a modified version of the above test as suggested by Harvey et al. (1997). As the results are very similar they are not reported here.

Choosing $g(\cdot)$ as the square function or as the absolute value function results in the MSPE or the MAPE, respectively. To evaluate the DQ we set e_{1t} to one if the forecast is in the right direction, and to zero otherwise. Zero changes either of the forecast or of the actual price are omitted. The DQ of the VAR is compared to a coin flip with probabilities 1/2 for each direction. Hence, $e_{2t} = 0.5$ and $g(\cdot)$ is chosen as the identity function. For the SC the forecast errors are chosen as $e_{1t} = (\Delta X_t \widehat{\Delta X_t})/\sqrt{\sum_{t=1}^{T}(\Delta X_t)^2 \sum_{t=1}^{T}(\widehat{\Delta X_t})^2}$ and $e_{2t} = 0$ for the benchmark of no signal correlation. ΔX_t and $\widehat{\Delta X_t}$ are observed and predicted returns, respectively.

To conclude, we note that these forecast measures are primarily of academic interest. However, they neither provide necessary nor sufficient conditions for forecast value in economic terms, i.e., a profitable trading strategy yielding positive returns. To evaluate the forecast value in that sense we implement two simple trading strategies based on the forecasts of the VAR model. The first "trader" takes a long position if the forecast direction is up and a short position otherwise. However, this trader is very busy and changes the trading position frequently. Therefore, we implement a more realistic second strategy which accounts for transaction costs. Since actual transaction costs are not available we used the quoted bid-ask spreads for simplicity. The motivation is

11

Table 5

Out-of-sample forecasting results of the VAR(4) for the observations 4388 to 8774. The results are given in terms of the MSPE, the MAPE, the DQ, and the SC relative to the benchmark model.

FX rate	Forecast horizon (hours)	MSPE ($\cdot 10^{-8}$)	MAPE ($\cdot 10^{-6}$)	DQ ($\cdot 10^{-2}$)	SC ($\cdot 10^{-2}$)
USD-DEM	1	0.31	-2.15	3.03*	6.01*
	2	-0.58	-4.96*	1.58†	2.91†
	4	-3.58†	-11.02*	-0.24	0.10
	8	-9.53*	-17.89*	-1.32	-1.22
USD-JPY	1	1.55†	6.22*	5.22*	8.51*
	2	-0.60	2.92	4.13*	2.74
	4	-2.25	-0.54	1.73*	0.50
	8	-1.26	-1.90	1.08	1.25*
DEM-JPY	1	4.78*	-0.79	5.88*	15.52*
	2	6.40*	1.68*	3.80*	13.32*
	4	5.55	1.26*	2.52†	9.06*
	8	1.42*	8.59	1.38†	9.06*

* Significant at the 5% level

† Significant at the 10% level

to combine the VAR direction forecasts with a "technical" trading rule which significantly reduces the trading frequency. If the current position is neutral, then the direction forecast is used to take a position. Now, the position is held until a so called "stop-loss" rule is violated. In an add-hoc way we used twice the transaction costs as a limit. If cumulative negative returns exceed this limit the position is immediately cleared. Next, the direction forecast is again used and so on. Transaction costs are subtracted from returns whenever a position is changed. For this strategy, winning runs are unlimited. On the other hand, long negative runs are not possible, since they are "stopped out".

4.2 Results

The results of the out-of-sample forecasting exercise are reported in table 5. For the 1- and 2-hour horizons the message is clear: the forecasting performance of the VAR is significantly better than that of the martingale. With a few exceptions, the VAR tends to outperform the martingale, and almost every

estimate is significant. The most significant results are those for the DQ and SC. The percentage of forecasts in the right direction is between 53.1% for the USD-DEM and 55.9% for the DEM-JPY for the 1-hour horizon. For the 2-hour forecasts the DQ is between 51.6% (USD-DEM) and 54.1% (USD-JPY). The corresponding SC varies from 6% to 15% (1-hour) and from 2.7% and 13.3% (2 hours). For the longer horizons the quality of the forecasts decreases continuously. Nevertheless, the DQ for an 8-hour forecast of the cross-rate DEM-JPY is 51.4%. Furthermore, the ranking of the rates from lower to higher predictability is USD-DEM, USD-JPY, and DEM-JPY.

Concerning the MSPE and the MAPE, the results are not that clear. For the shorter horizons there seems to be some predictability, whereas for the longer horizons the martingale significantly outperforms the VAR. We have two explanations for this result: First, the underlying distributions all have fat-tails. This means that the MSPE and the MAPE performances depend very much on a few large observations, turning the results very unstable. Second, the volatility, i.e., in our case either the MSPE or the MAPE loss series, has long memory (e.g., Trapletti and Fischer, 1997). When making l-step predictions the situation gets even worse. Thus, the asymptotics for the sample mean could change dramatically as shown in Taqqu (1975). E.g., for one of the MSPE loss series, we estimated the order of fractional integration as 0.22. Consequently, one should not over-interpret the results from the MSPE and the MAPE significance tests. Concerning the DQ and the SC the situation is less critical since the magnitude of the forecasts is irrelevant.

The cumulative return of the first trading strategy without taking transaction costs into account is shown in figure 4. However, to achieve this return frequent trading is necessary. Taking transaction costs into account yields a negative overall return (not shown). The second trading strategy is more successful as shown in figure 5. Although we have not optimized the stop-loss rule in any sense, the overall net-return is positive. Moreover, it is a well known fact among of FX traders, that the actual transaction costs are much smaller than the quoted bid-ask spreads. Therefore, our estimates of transaction costs can be considered conservative estimates of the actual costs.

5 Conclusions

We have presented a cointegration analysis on the triangle (USD-DEM, USD-JPY, DEM-JPY) of high frequency FX rates and studied the implications of our empirical findings on market efficiency. Theoretically, market efficiency requires the triangle to be cointegrated and the cointegration term to be a martingale difference sequence. In this article we have found that:
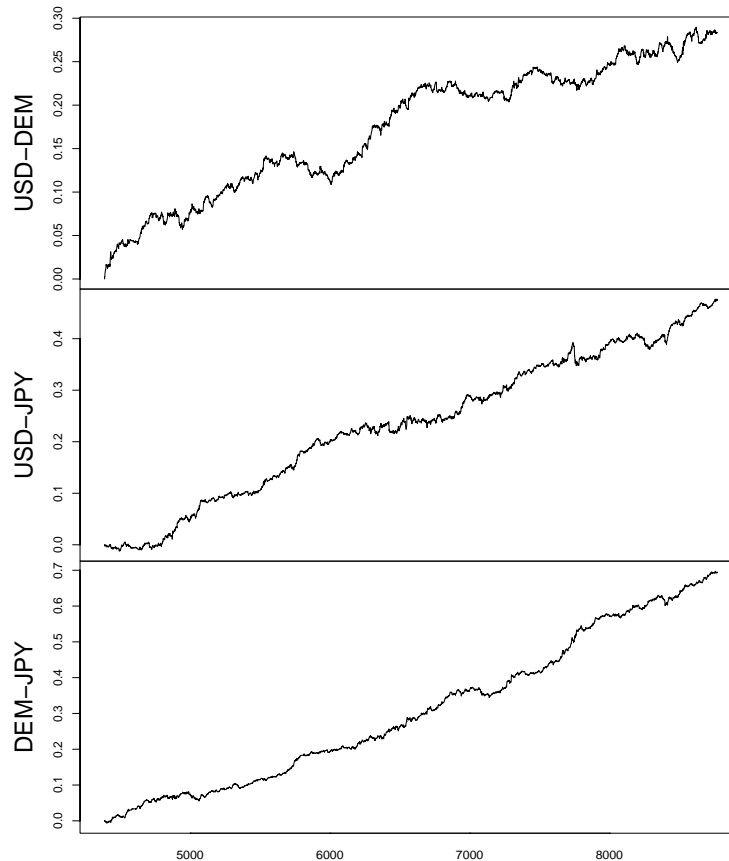
Fig. 4. Cumulative returns (ignoring transaction costs) of the first trading strategy for the out-of-sample observations 4388 to 8774.

- The relationship within the exchange rate triangle is characterized by exactly one cointegration relation. The cointegrating vector is unbiased in the sense that each rate has the same weight and no systematic deviation from equilibrium occurs.
- The residuals from the cointegration regression do not behave like a martingale difference sequence. However, the relatively rapid decay of the autocorrelation coefficients indicates the purely intra-day character of this phenomenon.
- The out-of-sample forecasting exercises and the improved performance testing procedures reject the null hypothesis of a martingale process, at least for short time horizons. On a 1-hour forecast horizon the percentage of forecasts in the right direction was found to be between 53.1% for the USD-DEM and 55.9% for the DEM-JPY.
- A simple trading scheme shows that the VAR also has a significant forecast value in terms of a profitable trading strategy yielding significantly positive returns even after transaction costs.
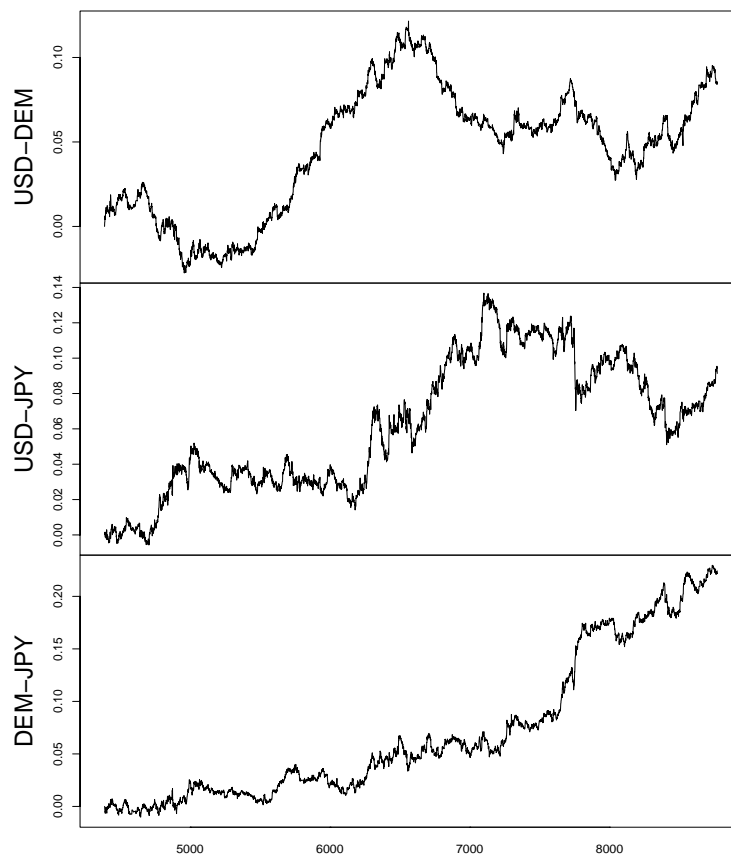
14

Fig. 5. Cumulative net-returns of the second trading strategy for the out-of-sample observations 4388 to 8774.

These results reject market efficiency in a weak form, i.e., the martingale hypothesis is not accepted for these high frequency exchange rates on such short forecasting horizons. This is not surprising when taking into account that some time is needed to integrate new price information into the different markets, and considering how diverse the interpretation of information can be. Concerning the profitability of such an approach it is doubtful if FX markets are fully efficient on very short time horizons. Therefore, a stronger version of market efficiency that accounts for transaction costs cannot be fully accepted.

Of course, more empirical work is necessary to support our findings. In particular, the extension of the analysis to a longer time span, and using more rates may be useful.

Table 6
Parameter estimates for the scaling power law. The numbers in parentheses are estimated standard errors. Δt is measured in seconds.

$\widehat{\log(c)}$	$\hat{e}$	R^2 fit
-11.703	0.521	0.998
(0.055)	(0.006)	

Appendix

A   Time Scale Transformation

The basis for the computation of the operational time scale is the empirical scaling power law found by Müller et al. (1990). It expresses the average volatility $\bar{V}$ over some time interval Δt as a function of the length of the time interval Δt

$$\bar{V} = c\,(\Delta t)^e. \tag{5}$$

Since we focus on multivariate modelling, $\bar{V}$ is computed as the volatility of a weighted average of the underlying time series. c and e are estimated from a regression of $\log(\bar{V})$ on $\log(\Delta t)$ and depend on the underlying time series and the averaging weights. The results of the regression are reported in table 6.

To express time as a function of the volatility we invert the scaling power law as suggested in Schnidrig and Würtz (1995). Application to each hour of an average week gives the length of this hour in operational time

$$\Delta t_i^{\text{op}} = \left(\frac{\bar{V}(\Delta t_i)}{c}\right)^{1/e}, \quad i = 1, \dots, 168, \tag{6}$$

where $\bar{V}(\Delta t_i)$ is the average volatility over the i'th hour of the week. Δt_i^{op} is the resulting length of the i'th hour in operational time.

The complete time mapping is found by interpolating between these 168 tabulated points $(\Delta t_i^{\text{op}}, \Delta t_i)$ and synchronizing the two time scales each week. We have used a cubic spline interpolation procedure as described in Press et al. (1995). Note that business time instead of physical time was used as the underlying time scale for these computations. The time mapping from business time to physical time is shown in figure 6.

The averaging weights are chosen in such a way that the sum of the seasonal fluctuations in operational time over all time series becomes minimal. We use the root mean square error of the average volatility around the mean (RMSEV) as a measure of the seasonal fluctuations (see Dacorogna et al., 1993). Table 7 shows the averaging weights which minimize the RMSEV criterion. Furthermore, the corresponding RMSEV in operational time are reported as percentages of the RMSEV in physical time.

Table 7
The averaging weights which minimize the sum of the seasonal fluctuations in operational time over all time series. The RMSEV-ratios are computed as the ratio of the RMSEV in operational time to the RMSEV in physical time.

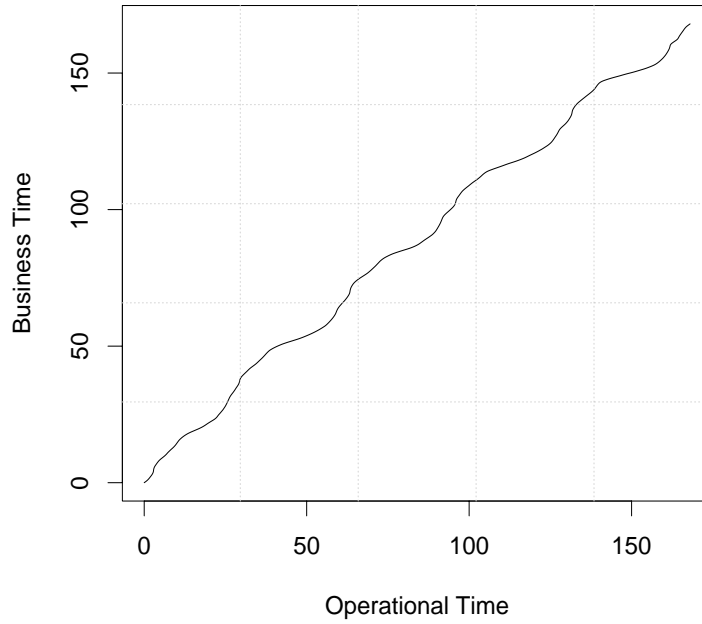|	USD-DEM	USD-JPY	DEM-JPY
weights	0.68	0.13	0.19
RMSEV-ratio	0.29	0.33	0.34

Fig. 6. Time mapping from operational time to business time.

B VAR Parameter Estimates

The maximum likelihood procedure gives the following estimates:

$$\hat{\boldsymbol{\Gamma}}_1 = \begin{bmatrix} -0.129 & 0.089 & -0.096 \\ (0.040) & (0.043) & (0.043) \\ 0.156 & -0.209 & 0.169 \\ (0.029) & (0.031) & (0.031) \\ -0.440 & 0.420 & -0.452 \\ (0.035) & (0.037) & (0.037) \end{bmatrix}, \hat{\boldsymbol{\Gamma}}_2 = \begin{bmatrix} -0.050 & 0.070 & -0.034 \\ (0.050) & (0.052) & (0.051) \\ 0.073 & -0.095 & 0.072 \\ (0.036) & (0.038) & (0.038) \\ -0.460 & 0.410 & -0.472 \\ (0.043) & (0.045) & (0.045) \end{bmatrix},$$

17

and

$$\hat{\boldsymbol{\Gamma}}_3 = \begin{bmatrix} -0.096 & 0.102 & -0.123 \\ (0.055) & (0.056) & (0.056) \\ 0.086 & -0.087 & 0.073 \\ (0.040) & (0.041) & (0.041) \\ -0.356 & 0.337 & -0.335 \\ (0.048) & (0.049) & (0.049) \end{bmatrix},$$

where the numbers in parentheses are estimated standard errors. Furthermore,

$$\hat{\boldsymbol{\Pi}} = \begin{bmatrix} -0.164 & 0.164 & -0.164 \\ (0.037) & (0.037) & (0.037) \\ 1.764 \cdot 10^{-4} & 1.764 \cdot 10^{-4} & 1.764 \cdot 10^{-4} \\ (0.027) & (0.027) & (0.027) \\ -0.343 & 0.343 & -0.343 \\ (0.032) & (0.032) & (0.032) \end{bmatrix}$$

and

$$\hat{\boldsymbol{\Lambda}} = \begin{bmatrix} 1.940 & 0.653 & -1.121 \\ 0.653 & 1.035 & 0.306 \\ -1.121 & 0.306 & 1.465 \end{bmatrix} \cdot 10^{-6}.$$

Acknowledgements

References

Baillie, R. T. and Bollerslev, T. (1989). Common stochastic trends in a system of exchange rates. *The Journal of Finance*, **44**, 167–181.

Baillie, R. T. and Bollerslev, T. (1994). Cointegration, fractional cointegration, and exchange rate dynamics. *The Journal of Finance*, **49**, 737–745.

Bolland, P. J. and Connor, J. T. (1995). Identification of FX arbitrage opportunities with a non-linear multivariate kalman filter. In *Proceedings of the Third International Conference on Neural Networks in the Capital Markets*, London.

Copeland, L. S. (1991). Cointegration tests with daily exchange rate data. *Oxford Bulletin of Economics and Statistics*, **53**, 185–198.

Cromwell, J. B., Labys, W. C., and Terraza, M. (1994). *Univariate Tests for Time Series Models*. No. 07-099 in Sage University Paper series on Quantitative Applications in the Social Sciences. Thousand Oaks, CA: Sage.

Dacorogna, M. M., Müller, U. A., Nagler, R. J., Olsen, R. B., and Pictet, O. V. (1993). A geographical model for the daily and weekly seasonal volatility in the FX market. *Journal of International Money and Finance*, **12**, 413–438.

Diebold, F. X., Gardeazabal, J., and Yilmaz, K. (1994). On cointegration and exchange rate dynamics. *The Journal of Finance*, **49**, 727–735.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, **13**, 253–263.

Hakkio, C. S. and Rush, M. (1991). Cointegration: How short is the long run? *Journal of International Money and Finance*, **10**, 571–581.

Harvey, D., Leybourne, S., and Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, **13**, 281–291.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica*, **59**, 1551–1580.

Johansen, S. and Juselius, K. (1990). Maximum likelihood estimation and inference on cointegration - with applications to the demand for money. *Oxford Bulletin of Economics and Statistics*, **52**, 169–210.

Jong, F. D., Mahieu, R., and Schotman, P. (1995). The dynamics of the actual and Dollar implied Yen/Dmark cross exchange rate. In *Proceedings of the First International Conference on High Frequency Data in Finance*, Zürich.

Müller, U. A., Dacorogna, M. M., Olsen, R. B., Pictet, O. V., Schwarz, M., and Morgenegg, C. (1990). Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking and Finance*, **14**, 1189–1208.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1995). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge, MA: Cambridge University Press, 2nd edn.

Schnidrig, R. and Würtz, D. (1995). Investigation of the volatility and autocorrelation function of the USD-DEM exchange rate on operational time scales. In *Proceedings of the First International Conference on High Fre-*

quency Data in Finance, Zürich.

Sephton, P. S. and Larsen, H. K. (1991). Tests of exchange market efficiency: Fragile evidence from cointegration tests. *Journal of International Money and Finance*, **10**, 561–570.

Taqqu, M. S. (1975). Weak convergence to fractional Brownian motion and to the Rosenblatt process. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, **31**, 287–302.

Trapletti, A. and Fischer, M. M. (1997). Bayesian modelling of high frequency data in finance. *Neural Network World*, **7**, 417–426.