

Detecting intraday periodicities with application to high frequency exchange rates

Article

Accepted Version

Brooks, C. ORCID: <https://orcid.org/0000-0002-2668-1153> and Hinich, M. J. (2006) Detecting intraday periodicities with application to high frequency exchange rates. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 55 (2). pp. 241-259. ISSN 1467-9876 doi: 10.1111/j.1467-9876.2006.00534.x Available at <https://centaur.reading.ac.uk/20509/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1111/j.1467-9876.2006.00534.x>

To link to this article DOI: <http://dx.doi.org/10.1111/j.1467-9876.2006.00534.x>

Publisher: Wiley

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

This is the authors' accepted manuscript of an article published in the *Journal of the Royal Statistical Society, Series C*. The definitive version is available at www3.interscience.wiley.com

Detecting Intraday Periodicities with Application to High Frequency Exchange Rates

Chris Brooks¹, Cass Business School
and Melvin J. Hinich, University of Texas at Austin

Abstract

Many recent papers have documented the existence of periodicities in returns, return volatility, bid-ask spreads and trading volume, in both equity and foreign exchange markets. In this paper, we propose and employ a new test for detecting subtle periodicities in time-series data based on a signal coherence function. The technique is applied to a set of seven half-hourly exchange rate series. Overall, we find the signal coherence to be maximal at the 8 hour and 12 hour frequencies. Retaining only the most coherent frequencies for each series, we implement a trading rule based on these observed periodicities. Our results demonstrate in all cases except one that, in gross terms, the rules are able to generate returns considerably greater than those of a buy-and-hold strategy, although they are unable to retain their profitability net of transactions costs. We conjecture that this methodology could constitute an important tool for financial market researchers, which will enable them to better detect, quantify and rank the various periodic components in financial data.

February 2005

J.E.L. Classifications: C32, C53, F31

Keywords: spectral analysis, periodicities, seasonality, forecasting, exchange rates, trading rules.

¹ Chris Brooks (Corresponding author), Faculty of Finance, Cass Business School, 106 Bunhill Row London EC1Y 8TZ, UK; Tel. (+44) 20 7040 5168; Fax: (+44) 20 7040 8881; e-mail: C.Brooks@cass.ac.uk The authors are grateful for comments from two anonymous referees and from the Editor, Alan Watkins. The usual disclaimer applies.

1. Introduction

Many recent papers have documented the existence of periodicities in returns, return volatility, bid-ask spreads and trading volume, in both equity and foreign exchange markets. Such systematically recurring features or regularities have sometimes been termed calendar anomalies or seasonal effects. Examples include open and close effects, lunchtime effects, day-of-the-week effects, and holiday effects. Studies of day-of-the-week effects include French (1980), Gibbons and Hess (1981), and Keim and Stambaugh (1984), for example, all of whom have found that the average market close-to-close return in the US is significantly negative on Monday and significantly positive on Friday. Moreover, Rogalski (1984), and Smirlock and Starks (1986) observed that this negative return between the Friday close and Monday close for the Dow Jones Industrial Average (DJIA) occurs on Monday itself during the 1960's but moves backward to the period between the Friday close and Monday open in the late 1970's. By contrast, Jaffe and Westerfield (1985) found that the lowest mean returns for the Japanese and Australian stock markets occur on Tuesdays. Harris (1986) also examines weekly and intraday patterns in stock returns and finds that most of the observed day-of-the-week effects occur immediately after the open of the market, with a price drop on Mondays on average at this time and rises on all other weekdays.

Recent research has also exploited the increasing availability of very high frequency and tick data together with more powerful computational abilities in order to analyse more closely the intraday patterns in financial markets. Wood *et al.* (1985) examine minute-by-minute returns data for a large sample of NYSE stocks. They find that significantly positive returns are on average earned during the first 30 minutes of trading and at the market close, a result echoed by Ding and Lau (2001) using a sample of 200 stocks from the Stock Exchange of Singapore. Andersen and Bollerslev (1997b) show that the return volatility of the German mark – dollar exchange rate exhibits the same general intraday pattern as trading volume and bid-ask spreads. Using the same set of data, Andersen and Bollerslev (1998) also study the effects of macroeconomic announcements on the behaviour of the series².

² An extensive survey of the literature on intraday and intraweek seasonalities in stock market indices and futures market contracts is given in Yadav and Pope (1992).

The existence of such periodicities has typically been reconciled with the efficient markets hypothesis by appealing to market microstructure arguments (e.g. cyclical variations in market depth or liquidity; price discovery; inventory management), information arrival, macroeconomic announcements, or tax effects. Many theoretical models of investor and market behaviour have also been proposed to explain these stylised features of many financial time series, including those that account for the strategic behaviour of liquidity traders and informed traders (see, for example, Admati and Pfleiderer, 1988). An alternative method for reconciling a finding of recurring seasonal patterns in financial markets is the possible existence of time-varying risk-premia, implying that expected returns need not be constant over time, and could vary in part systematically without implying market inefficiency.

Traditionally, studies concerned with the detection of periodicities in financial time series would either use a regression model with seasonal dummy variables (e.g., Chan *et al.*, 1995) or would apply spectral analysis to the sample of data (e.g. Bertoneche, 1979; Upson, 1972). Spectral analysis may be defined as a process whereby a series is decomposed into a set of mutually orthogonal cyclical components of different frequencies. Calculating the spectrum involves fitting by least squares a set of sinusoidal curves, equivalent to a regression with trigonometric transformations of the independent variable. The spectrum, a plot of the signal amplitude against the frequency, will be flat for a white noise process, and evidence of periodic behaviour is indicated by statistically significant amplitudes at any given frequency. By examining the spectrum, Upson (1972) observed periodicities representing cycles of duration 32, 3.8 and 2.5 weeks in US dollar – British pound data for the 1960's, while Bertoneche (1979) was not able to detect any significant departures from randomness in an application of spectral analysis to a set of weekly European stock returns. Spectral analysis was a popular tool for data analysis in economics and finance in the 1960's and 1970's (see also Granger, 1966; Granger and Hatanaka, 1964; Granger and Morgenstern, 1969), although it has been largely discarded in the empirical economics and finance literature more recently. This seems to stem partly from the perceived inability of the spectral methods that were previously available to detect the relevant temporal dependencies in financial time series data.

In this paper, we propose and employ a new test for detecting periodicities in financial markets based on a signal coherence function. Our approach can be applied to any fairly large, evenly spaced sample of time series data that is thought to contain periodicities. A periodic signal can be predicted infinitely far into the future since it repeats exactly in every period. In fact, in economics and finance as in nature, there are no truly deterministic signals and hence there is always some variation in the waveform over time. The notion of partial signal coherence, developed in this paper into a statistical model, is a measure of how much the waveform varies over time. The coherence measures calculated are then employed to hone in on the frequency components of the Fourier transforms of the signal that are the most stable over time. By retaining only those frequency components displaying the least variation over time, we are able to detect the most important seasonalities in the data, and these are then used to derive a trading rule for buying or selling a currency in a hold-out sample. The performance of the trading rule is then compared with that of buy-and-hold and randomised trading strategies. Our approach is able to capture a much broader range of regularities than a linear regression with time-dependent dummy variables. The model that we propose based on the signal coherence function has been shown to provide additional detectability relative to a Fisher test (see Hinich, 2003).

The remainder of this paper is organised as follows. Section 2 introduces some notation and defines the test statistics employed to detect the periodicities, while Section 3 describes the data. Section 4 presents and analyses the results while Section 5 concludes and offers suggestions for extensions and further research.

2. Methodology

2.1 Development of a Test for Signal Autocoherence

This paper develops below a model for a signal with randomly modulated periodicity, and a measure known as a signal coherence function, which embodies the amount of random variation in each Fourier component of the signal. Let $\{x(t_n), n = 0, 1, 2, \dots\}$ be a time series of interest sampled at equally spaced times $t = n \delta$. The series would be said to exhibit randomly modulated periodicity with period T if it is of the form

$$x(t) = \frac{a_0}{K} + \frac{1}{K} \sum_{k=1}^K (a_{1k} + u_{1k}(t)) \cos(2\pi f_k t) + \frac{1}{K} \sum_{k=1}^K (a_{2k} + u_{2k}(t)) \sin(2\pi f_k t) \quad (1)$$

where $f_k = k/T$ are called Fourier frequencies and $K = T/\delta$. The u_{ik} ($i=1,2$) are jointly dependent zero mean random processes that are periodic block stationary and satisfy finite dependence. The modulations are in the fundamental and harmonic frequency components and so a_0 is a constant. Note that we do not require $u_{1k}(t_n)$ and $u_{2k}(t_n)$ to be Gaussian. It is apparent from (1) that the random variation occurs in the modulation rather than being additive noise; in statistical parlance, the specification in (1) would be termed a random effects model. The modulations are produced by the data generating mechanism and may in fact be deterministic but in most cases the statistician analysing the data does not know the modulations process and thus they are treated as random processes. The signal $x(t_n)$ can be expressed as the sum of a deterministic (periodic) component and a stochastic process term as follows:

$$x(t) = \frac{a_0}{K} + \frac{1}{K} \sum_{k=1}^K a_k \exp(i2\pi f_k t) + \frac{1}{K} \sum_{k=1}^K u_k(t) \exp(i2\pi f_k t) \quad (2)$$

where $a_k = a_{1k} - ia_{2k}$ and $u_k(t_n) = u_{1k}(t_n) - iu_{2k}(t_n)$. The task at hand then becomes one of quantifying the relative magnitude of the k th modulations to the magnitude of the fixed effect a_k for each k .

A common approach to processing signals with a periodic structure is to portion the observations into M frames, each of length $T = N \delta$, so that there is exactly one waveform in each sampling frame. There could alternatively be an integer multiple of T observations in each frame. The periodic component of $a(t)$ is the mean component of $x(t)$. In order to determine how stable the signal is at each frequency across the frames, the notion of signal coherence is employed. Signal coherence is loosely analogous to the standard R^2 measure used in regression analysis, and quantifies the degree of association between two components for each given frequency. It is worth noting that the methodology that we propose here is based on the coherence of the signal across the frames for a single time series (which may also be termed

autocoherence). This is quite different from the tests for signal coherence across markets used, for example, by Hilliard (1979) and Smith (1999)³.

The discrete Fourier transform of the m^{th} frame, beginning at observation $\beta_m = ((m-1)T) + \delta$ and ending at observation mT for frequency $f_k = k/T$ is given by

$$x_m(k) = \sum_{t=0}^{T-\delta} x(\beta_m + t) \exp(-i2\pi f_k t) = a_k + U_m(k) \quad (3)$$

where $U_m(k) = \sum_{t=0}^{T-\delta} u(\beta_m + t) \exp(-i2\pi f_k t)$. The variance of $U_m(k)$ is given by

$$\sigma_u^2(k) = \sum_{\tau=0}^{T-\delta} \exp(-i2\pi f_k \tau) \sum_{t=0}^{T-\tau-1} c_u(t, t + \tau) \quad (4)$$

where $c_u(t_1, t_2) = E[u_m^*(t_1)u_m(t_2)]$, and the variance is of order $O(T)$. Provided that $u_m(t_n)$ is weakly stationary, (4) can be written

$$\sigma_u^2(k) = T[S_u(f_k) + O(1/T)] \quad (5)$$

where $S_u(f)$ is the spectrum of $u(t_n)$.

Although the model for randomly modulated periodicity is not in general a cyclostationary one, when the assumption of weak stationarity is added, it becomes so (see Gladyshev (1961), Gardner and Franks (1975), Gardner (1985), and Gardner (1994) for extensive writings on cyclostationarity). The signal coherence function, $\gamma_x(k)$, measures the variability of the signal across the frames, and is defined as follows for each frequency f_k

$$\gamma_x(k) = \sqrt{\frac{|a_k|^2}{|a_k|^2 + \sigma_u^2(k)}} \quad (6)$$

³ Both of these papers employ the frequency domain approach in order to examine the extent to which stock markets co-move across countries. Our technique is also distinct from that proposed by Durlauf (1991) and used by Fong and Ouliaris (1995) to detect departures from a random walk in five weekly US dollar exchange rate series. Both Fong and Ouliaris (1995) and Andersen and Bollerslev (1997a) detect long memory effects in the currency rates.

It is obvious from the construction of $\gamma_x(k)$ in (6) that it is bounded to lie on the $[0,1]$ interval. The endpoint case $\gamma_x(k) = 1$ will occur if $a_k \neq 0$ and $\sigma_u^2(k)=0$, which is the case where the signal component at frequency f_k has a constant amplitude and phase over time, so that there is no random variation across the frames at that frequency (perfect coherence). The other endpoint, $\gamma_x(k) = 0$, will occur if $a_k=0$ and $\sigma_u^2(k) \neq 0$, when the mean value of the component at frequency f_k is zero, so that all of the variation across the frames at that frequency is pure noise (no coherence).

The signal coherence function is estimated from the actual data by taking the Fourier transform of the mean frame and for each of the M frames. The mean frame will be given by

$$\bar{x}(t_n) = \frac{1}{M} \sum_{m=1}^M x(\beta_m + t_n) \quad , \quad n = 0, 1, \dots, N-1 \quad (7)$$

Letting $\hat{a}(k)$ denote the mean frame estimate, with its Fourier transform being $\hat{A}(k)$, and letting $X_m(k)$ denote the Fourier transform for the m^{th} frame, then $D_m(k) = X_m(k) - \hat{A}(k)$ is a measure of the difference between the Fourier transforms of the m^{th} frame and the mean frame for each frequency. The signal coherence function can then be estimated by

$$\hat{\gamma}_x(k) = \sqrt{\frac{|\hat{A}_k|^2}{|\hat{A}_k|^2 + \frac{1}{M} \sum_{m=1}^M |D_m(k)|^2}} \quad (8)$$

and $0 \leq \hat{\gamma}_x(k) \leq 1$. It can be shown (see Hinich, 2000) that the null hypothesis of zero coherence at

frequency f_k can be tested using the statistic $M \frac{\hat{\gamma}_x(k)^2}{1 - \hat{\gamma}_x(k)^2}$, which is asymptotically distributed under the

null as a non-central chi-squared with two degrees of freedom and non-centrality parameter given by

$\lambda_k = \frac{M a_k^2}{TS_u(f_k)}$, where $S_u(f_k)$ is the spectrum of $\{u(t)\}$ at the frequency f_k . We also employ a joint test of

the null hypothesis that there is zero coherence across the M frames for all $K/2$ frequencies examined. This test statistic will asymptotically follow a non-central Chi-squared distribution with K degrees of freedom.

2.2 Development and Testing of a Trading Rule

The sample is split into ($M=$) 52 non-overlapping frames each of length one week, with each week of observations containing 240 half-hourly observations⁴. This implies that a total of 120 periodicities will be examined: 240, 120, 80, 60, 48, ..., 240/119, 2, using the whole sample of data, and the autocohereance measures are calculated across the 52 frames. Following initial exploratory analysis and estimation of the signal coherence function for the whole year's set of frames data, the sample is then split into two portions. The first 26 weeks (6,240 half-hourly observations) are used for in-sample estimation of the coherence function (across the resulting 26 frames), and the remaining data is held back for out-of-sample trading rule evaluation. The out-of-sample period begins with the return of 00:30am-01:00am on 1 July 1996.

For the out-of-sample trading rule analysis, the signal coherence function is re-estimated using the first half of the sample only, and the mean frame is "cleaned" by removing all frequency components with signals whose random variation implies that they are not statistically coherent at the 1% significance level. Statistically, if $\Pr(\hat{\gamma}_x(k)=0)<0.01$, \hat{a}_{jk} ($j = 1,2$) are kept, otherwise \hat{a}_{jk} are set to zero. The coherent part of the mean frame is analogous to the fitted value of the dependent value in a standard regression model. Retaining only the most coherent frequencies for each series, we implement a trading rule based on these observed periodicities, "buying" the foreign currency if the return is predicted to be positive and assuming no position if it is predicted to be negative. The trading rule is then compared with a buy-and-hold the foreign currency strategy and also in a novel way involving a simulation. In order to evaluate the statistical significance of the strategies, we generate random binomial 0/1 draws equal in number to the out-of-sample observations. This column of zeros and ones is then multiplied by the actual return series and added over the hold-out sample to generate a profit from a randomised market entry and exit rule. This is repeated 10,000 times to generate a distribution of artificial rules and subsequent artificial returns, which is then compared with the profit generated by the rules based on the coherent periodic signal. If the actual rule

⁴ The choice of frame length is bound to be somewhat arbitrary, although in our case it represents a trade-off between having a sufficient number of frames over which to compute the tests, while having a long enough frame to detect interesting periodic components.

generates a profit larger than 95% of those generated by artificial timing, it is considered to produce statistically significant abnormal returns.

We then exploit the information contained in the lower half of the signal coherence function by short-selling the foreign currency when it is periodically expected to exhibit negative returns, as well as taking a long position when it is expected to yield positive returns. The relevant simulation comparator is now one where the column of zeros and ones is modified to a column of -1 and $+1$ to be multiplied by the actual returns to give a set of returns from randomised long and short positions.

3. Data

The high frequency financial data provided by Olsen and Associates as part of the HFDF-96 package includes 25 exchange rate series sampled half-hourly for the whole of 1996, making a total of 17,568 observations for each series. However, this series contains observations corresponding to weekend periods when all the world's exchanges simultaneously have virtually no trading. This period is the time from 23:00 GMT on Friday when North American financial centres close until 23:00 GMT on Sunday when Australasian markets open. The incorporation of such prices would lead to spurious zero returns and would potentially render trading strategies which recommended a buy or sell at this time to be nonsensical. Removal of these weekend observations leaves 12,575 observations for subsequent analysis and forecasting.

We do not account for differences in the dates that different countries switch to daylight saving time, since the effect of this one-hour difference is likely to be negligible as trading occurs virtually around the clock at some destination around the world. Moreover, it is not clear how such an adjustment could be made. This problem would be much more serious if we were examining, say, cross-correlations between equity returns for stocks on markets that were in different time zones. Such an approach appears to be consistent with the existing literature - for example, Andersen and Bollerslev (1997a) do not correct for bank holiday effects. Finally, it is worth noting that if asynchronous switches to daylight saving time did make a

difference to our results, since it is not possible to create coherence, only to destroy it, the impact would be to render our results less strong than they otherwise would have been rather than causing spurious findings.

The price series used are the average of the most recent best bid and best ask in that half-hour interval, and are transformed into a set of continuously compounded half-hourly percentage returns in the standard fashion. The first observation in the sample corresponds to the period between 00:30am and 01:00am GMT on 1 January 1996, while the last corresponds to 11:30pm-midnight GMT on 31 December of the same year.

Of the 25 exchange rate series provided by Olsen, only 7 are used in this study to avoid repetition, and due to space constraints. These are (using the usual neumonic) DEM_JPY, GBP_DEM, GBP_USD, USD_CHF, USD_DEM, USD_ITL, and USD_JPY. These are all quoted as units of the second (foreign) currency per unit of the first (domestic) currency. For example, xxx_yyy would be the units of yyy per unit of xxx . Hence a positive return implies that the number of units of the foreign currency yyy per unit of the domestic currency (xxx) has increased. This would imply a profit for an investor who's domestic currency is xxx , but who had switched into yyy at the start of the period and converted the terminal sum back at the end.

Some summary statistics for these seven returns series are presented in Table 1. It is clearly evident that all series are non-normal (predominantly due to fat tails rather than asymmetry), and all exhibit evidence of negative first order autocorrelation, and conditional heteroscedasticity (as the Ljung Box and Engle tests respectively show). The BDS statistic therefore rejects the null hypothesis of independent and identical distribution at the 0.1% level of significance.

4. Results

4.1 Evidence for Periodicities

Table 2 gives the p -values for tests of the joint null hypothesis that there is zero coherence at all frequencies examined for the whole 52-week sample. Clearly, there is significant evidence of coherence at one or more frequencies for the pound-mark and the dollar-yen, while the result for the dollar-lira is marginal. The joint test results would also suggest that there is no coherence at any frequency for the other four exchange rate series. A non-rejection from the joint test does not in practice imply that there is actually no coherence at any frequency, however, since the effect of significance at one or two frequencies could be diluted by many insignificant frequencies. Hence Table 3 presents the periodicities for which the individual autocohereance estimates are statistically significant, together with the associated p -values for the chi-squared test, and the log of the spectrum (in dB). Whilst there is no single periodicity where all seven series show significant coherence simultaneously, there are several common features across the exchange rate returns. The mark-yen and dollar-yen both have statistically significant (at the 0.3% level or higher) autocohereance at a periodicity of 120 half-hourly units (i.e. 60 hours, or half a week), while all of the series denominated against the dollar, except the lira, show autocohereance at the 8-hour periodicity. Interestingly, the dollar-Swiss franc and the dollar-yen also have significant coherence at the 15-hour periodicity, corresponding to 8 cycles per week. In the latter case, this is the most empirically stable periodicity, with a coherence statistic of 0.362 (on a 0-1 scale), and an associated p -value of less than 0.1%. What is also evident from Table 3 is that none of the coherence statistics are larger than 0.362, implying that there is still a considerable amount of variation in the waveform over the frames even for the most coherent parts. The log-spectrum in dB is 20 times the natural log of the spectral amplitude or equivalently, it is 10 times the spectrum (in variance units). The measure in dB, as shown in Table 3, gives on a log-scale the average sizes of the periodic movements in terms of the heights of the peaks and troughs of the coherent periodicities. Whilst autocohereance quantifies how stable these periodicities are, the amplitude measures how big the cyclical fluctuations are. It is evident from the second column of Table 3 that, in general, the higher frequency components have the largest amplitudes, with some even being positive (on a log scale) at frequencies of lower than three hours for the dollar-Swiss, dollar-lira, and dollar-yen.

How can one explain the observed 8-, 12-, and 15-hour periodicities? A number of factors could justify their existence, including the opening and closing times of the three major markets in different time zones (London, New York and Japan), and related changes in market volatility and liquidity through each 24-hour period. For example, Andersen and Bollerslev (1998, Figure 3) report cycles in intra-day volatility, where it is highest from 1pm-5pm, and lowest from 3am to 5am (GMT). It may therefore simply be that the cycles in returns are rewards for bearing time-varying intra-daily risks, which are themselves cyclical. A cycle that repeats every 8 hours is consistent with an effect driven by the opening and closing of the three markets, while a 15-hour cycle⁵ is consistent with an effect that is attributable to the currency being heavily traded in two of the three major centres (e.g., the yen-dollar being heavily traded in New York and Tokyo but not London). This is exactly the kind of result that Baillie and Bollerslev (1991) observed when hourly dummies were applied to intra-daily exchange rate volatility. The Asian market generated much less volatility than the other two, with noticeable increases in volatility occurring around the start of the trading days for the London and New York markets.

A 12-hour cycle is arguably harder to explain but may arise from behaviour caused by the full trading day of one market combined with the morning of another (e.g., if the German mark – yen) were actively traded during the full European trading day and for the morning of the Japanese trading day, but with very thin trading for the Japanese afternoon and the whole of the North American trading day.

Figures 1 and 2 present plots of the coherence statistics at each frequency for the dollar-denominated and non-dollar denominated currencies respectively. The first of these seems to suggest some correlation between the coherences for the dollar-mark and dollar-Swiss franc, while in all cases the low frequency components (those with the highest periods) show less variation across frames; most coherence statistics appear to lie within the range (0.1, 0.2). Figure 2, on the other hand, seems to show no visible relationship

⁵ Periodicities of 15 rather than 16 hours would result if the time between the time when one regional market “opens” and the regional market in the next time zone “closes” were 15 hours rather than 16 due to slightly shorter busy trading hours.

across the currencies in the coherences of the waveforms at any frequency. Again, the statistics appear largest at the lowest frequencies, and smallest at the hourly frequency.

Figures 3 to 9 plot the coherent part of the mean frame for frames of length one week for each of the returns series. As one would expect, the graphs look very different from one another, at least at first glance, since different frequencies have been retained for different currencies, and even when the same frequencies are included, differences in their relative amplitudes will alter the shape of the plot. In all cases, however, the cyclical patterns are obvious in the figures now that the mean frame has been purged of frequencies with higher amounts of random variation. In Figure 3, it is evident that the mark-yen has a low-frequency component (with 2 full cycles per week) and a high frequency cycle around that with a period of approximately 2 hours. In Figure 4, the low-frequency component of the pound-mark is harder to visualise since it is smaller in amplitude than that of the mark-yen and there are two high-frequency cycles around it. Figures 5 to 7, showing plots of the coherent parts of the mean frame for the pound-dollar, dollar-Swiss franc, and dollar-mark respectively, look similar in shape as all have two coherent frequencies close together that have exceeded the 99% threshold. The dollar-Swiss also has two additional coherent high frequency components. In Figure 8, the coherence plot for the dollar-lira has a very jagged appearance, due to the high frequency (approximately one-hour-and-twenty-minutely) component with far larger amplitude than the other two. Finally, Figure 9 showing the dollar-yen is arguably the most complex, since the highest number of frequencies (5) have been retained.

It is also of interest to examine Figures 3 to 9 for the purpose of determining on average at which times of the week the exchange rate moves in particular directions. The x -axis shows the half-hourly observations for a week, starting with the return for the interval midnight-00:30am, and plotted against it the average continuously compounded return, in percent, during that period. Examining the low-frequency cyclical component of Figure 3, the returns are negative (i.e. the mark weakens) for the first 60 half-hours (1 and $\frac{1}{4}$ trading days), then the mark rises until mid-week, before falling again for $1\frac{1}{4}$ days and rising on Friday. The most easily interpretable behaviour can be examined where there are time-zone differences between

the foreign and domestic currencies, and coherent frequency components with periods of 8, 12, or 15 hours, which there appears to be for many of the other series. For example, the pound-dollar exchange rate (Figure 5) on average rises (the pound becomes stronger) from the early hours of the morning GMT, until mid-day GMT, followed by falls when the US East coast markets start trading. The dollar-German mark exchange rate on average rises (the dollar strengthens) until around 10am GMT, and then falls (the mark strengthens) while the European markets are open until 4pm GMT, followed by further dollar appreciation. Similarly, the mark-Swiss franc exchange rate (Figure 6) on average rises (the dollar strengthens) until 10am, with falls (the franc appreciates) until around 3pm, followed by franc depreciation into the European evening. The dollar-yen (Figure 9) appears to fall (dollar weakens) from midnight until 4am GMT, when the US markets are closed, and then rises until around 2pm GMT before falling back again. Finally, the dollar-lira has such a large number of high-frequency movements that any associations with market opening and closing times are indiscernible.

These results support those of Ito and Roley (1987), who find a systematic dollar appreciation against the yen during the US trading hours, and a systematic depreciation during Japanese trading hours. Our results are also consistent with those of Baillie and Bollerslev (1991), who observed hourly patterns in foreign exchange market volatility related to major market opening and closing times. This is, perhaps, to be expected, since the weight of trading volume will move between the different world centres through the day.

4.2 Evaluation of Trading Rule Profitability

Given the results presented above showing evidence that there are periodically recurring patterns in all seven exchange rate return series, we now continue to investigate whether trading opportunities arise from exploiting such structure in the first half of the sample, as described above. The total gross profits, continuously compounded, and expressed as a percentage of the initial investment, for the 6-month out-of-sample period, are presented in Table 4. The row immediately following the column headers presents the gross profits for a rule based on a long position only in the foreign currency or no position (i.e. it is

assumed that no short sales of the foreign currency are permitted). For comparison, the last panel of the table presents the profitability of buying the foreign currency at the start of the 6-month-period, holding it until the end of the year and then converting it back to the domestic currency. This is termed a “buy-and-hold-the-foreign-currency” rule. As can be seen, the trading rules based on the notion of signal coherence appear to generate useful entry and exit rules since the returns in all cases exceed those obtained by simply buying and holding the foreign currency, except for the pound-dollar. For example, a British investor buying marks and holding them for the 6 months to December 1996 would have experienced a foreign currency appreciation of 10.8%, but switching in and out of marks using the coherent parts of the signal as a guide would have increased returns to 17%. The lira depreciated against the dollar over the period by 0.11%, so that a US investor would have lost money by buying and holding lira. However, the switching rules would have generated gross profits of 7.8% over the period.

The sixth row of Panel A in Table 4 presents the percentage of times that randomly generated entry times into the foreign currency would have led to higher returns than those generated by the signal coherence-generated rules. It is then argued that figures below, say, 5%, imply that the rule generates profits that are statistically significant at the 5% level, since this would indicate that it is less than 5% likely that the rules could generate such high profits by chance alone. The figures in the third row of the table are lower than 1% for four of the cases, and are only larger than 5% for the pound-dollar. In the case of the dollar-yen, none of the 10,000 sets of randomised rules were able to generate a higher return than that of the rules based on spectral methods. Clearly, the gross profitability of the entry and exit rules from foreign currency investment appear to show evidence of market timing ability. This is examined statistically using the Pesaran-Timmerman (1992, 1994) non-parametric test for market timing ability, which is generalised from the Henriksson-Merton (Henriksson and Merton, 1981; Merton, 1981) test for independence between the signs of the forecast and realised values. The statistic is not derived here⁶, but is asymptotically distributed standard normal under the null hypothesis that the signs of returns and the buy-sell signals are independent.

⁶ Interested readers are referred to the original papers for a derivation of the test statistic.

The number of correct sign predictions varies between 53% and 55% across the currencies, while the number of correct predictions that would be expected to result if the buy-sell signals and actual returns were independent is very close to 50% in all cases. Note that in the use of very high frequency data implies that a non-negligible number of actual returns per half-hour interval are zero. Although it is not clear from an examination of the formulation of the test, we define zero returns as implying correct sign predictions if the rule signal for that observation had been a zero (so that the transactions costs associated with buying the foreign currency are correctly avoided). However, incorrect sign predictions are argued to result if the rule generated no buy signal, but the return turned out to be positive. The relatively high proportion of correct sign predictions combined with the large sample size leads the Pesaran-Timmerman statistic to convincingly reject the null hypothesis at the 1% level in all cases.

The trading profitability of rules exploiting the sell-signals generated for those observations when the coherence plots are negative are also examined in Panel B of Table 4. In this case, a positive expected return for a given observation is taken to imply a buy order, while a negative expected return is taken to imply short-selling, with all results being expressed as a percentage of the initial position. In all cases except one, enabling short selling leads to considerable increases in gross profitability – in fact, gross profitability is typically doubled, as one may have expected. For example, the 18.4% and 7% gross profits from trading the yen against the dollar and the mark against the dollar respectively, become 31.8% and 12.5%. However, the profit made by taking long positions only in the dollar against the pound actually fell when short positions are permitted. This result does not seem to arise from anomalous behaviour in the series, such as a larger trend component or larger outliers than were present in the other series. Rather, it is possible simply that the behaviour of the pound-dollar exchange rate was more different than the others between the in-sample and out-of-sample periods; an examination of the coherence statistics revealed that this was indeed the case.

Comparing the standard deviations of the long only, long and short, and buy-and-hold rules, the latter two have very similar risk profiles, while the long only rule implies, by definition, less risk since there will be a

flat position in the foreign currency for part of the period. Also presented are the largest half-hourly gains and largest half-hourly losses for each of the three strategies. In all cases except that of the dollar-lira, the largest gains and largest losses are identical for the long only and long-short rules, suggesting that the extreme returns in both directions arise from being long rather than short the foreign currency. The Sharpe ratios for each rule and currency are also calculated by taking the average half-hourly risk-free rate of return (proxied by the domestic country 3-month Treasury bill rate) from the average half-hourly rule rate of return, and dividing by its standard deviation. Again, across the three panels of Table 4, in all cases except the pound-dollar, the Sharpe ratios are larger for the coherence-based rules than for the buy-and-hold approach. This suggests that after accounting for any differences in the riskiness inherent in the rules, they are still able to outperform a buy-and-hold the foreign currency rule. Also, the long and short rule Sharpe ratio is typically only very slightly higher than that of the long only approach. Thus, even though gross profitability generally increases considerably when short positions are permitted, the increased risk counters this to leave the risk-adjusted performance only slightly altered.

Figure 10 shows the cumulative returns for the long and short trading rules based on the signal coherence concept. Buying and selling German marks is the only activity that is seen to lose an American investor money for a protracted period. Trading in DEM using the methods described above would have been a loss-making activity from July until mid-September when there was a substantial reversal of fortune. UK investors trading marks or US investors trading yen would have earned steady profit streams, while the rules corresponding to German investors would have given a fairly flat return profile indicating a relatively lacklustre performance.

Finally, the last row of Panel A in Table 4 displays the number of “round trips” required to apply the coherence-based rules, that is – the number of purchases and sales of the foreign currency per week. For a large inter-bank trade, the bid-ask spreads on foreign exchange trades are very small – of the order of 0.05%, and there are no other fees involved. The average bid-ask spreads for each currency over the sample period, expressed as a percentage of the mid-point quotes, are given in the last row of the table. It is clear

that, in spite of such small transactions costs, the rules are not able to generate any positive net returns, since such a large number of trades are made. This is true for all exchange rates tested, including both the dollar-yen, where gross profits were quite phenomenal, and for the dollar-mark, where the smallest number of trades was suggested.

The lack of net trading profits as a result of excessive transactions costs suggests the use of a filter rule, so that the number of round trips is reduced. We experimented with various filter rules, including those that imply buying or selling the foreign currency when the predicted return fell outside of threshold limited equivalent to the bid-ask spread. Whilst the filter rules were successful in reducing the number of trades, they also considerably reduced the gross profitability, so that these were unable to generate net positive returns either.

5. Conclusions

This paper has proposed and employed a new methodology for evaluating and quantifying the autocohereance of financial time series, which was then tested on a set of seven half-hourly exchange rate returns. Significant coherence for at least one frequency across frames was revealed for all series. Overall we find the signal coherence to be maximal at the 15 hour, 12 hour and especially 8 hour frequencies. In the latter case, this may be attributable to the opening and closing of the world's three financial market time zones.

The mean frame estimate was cleaned by removing all incoherent frequency components, and the remaining estimates were used to generate trading rules based on these most stable cyclical features. The strategies were employed to construct a number of performance evaluation measures commonly used by financial market practitioners. These trading rules were able to generate in most cases phenomenal gross profits that were statistically significant and showed evidence of significant market timing ability, but the very large number of required trades meant that these were more than wiped out by transactions costs.

Our analysis could be extended and enhanced in a number of ways. First, in the out-of-sample evaluation, improved trading performance is likely to result from the use of a rolling window of data, where a one-frame ahead forecast is produced and the trading rule implemented, and then the sample rolled forward by one frame. Second, the use of a 1% significance level cut-off for determining the coherent parts of the signal is somewhat arbitrary, and sensitivity analysis could be conducted to determine the impact on profitability of increasing or reducing this threshold. If this threshold is set too high, then important cycles in the data will not be employed, while the situation if it is set too low is comparable to the effect of including irrelevant regressors in a standard time-series regression model used for forecasting.

A third possibility would be an examination of a much longer run of lower frequency data, so that the frequency of transacting would be reduced and potential returns per trade increased in order to mitigate the effect of the bid-ask spread on profitability. These factors suggest that the profitability demonstrated in this paper is probably an understatement of that which is possible if the approach is further refined and optimised. In any case, there are many financial agents who do not require immediacy when making foreign exchange trades. US Multinational companies, for example, may desire to reduce their holdings in marks and increase their holding in pounds, and they may be willing to make the trades at any time in the near future, although the exact time is not of concern to them. Such agents would be forced to incur the transactions costs whenever they trade, and our research suggests that there are better methods for selecting the time to trade than doing so at random. Some large financial institutions (for example, hedge funds,) may be able to trade at much lower transactions costs than the figures that we quote, thus significantly improving net returns. We conjecture that the methodology employed in this paper could be a widely applicable tool for market microstructure researchers and statistical arbitrageurs, which will enable them to better detect, quantify and rank the various periodic components in financial data.

References

Admati, A.R. and Pfleiderer, P. (1988) A Theory of Intraday Patterns: Volume and Price Variability
Review of Financial Studies 1(1), 3-40.

- Andersen, T.G. and Bollerslev, T. (1997a) Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns *Journal of Finance* 52, 975-1005.
- Andersen, T.G. and Bollerslev, T. (1997b) Intraday Periodicity and Volatility Persistence in Financial Markets *Journal of Empirical Finance* 4, 115-158.
- Andersen, T.G. and Bollerslev, T. (1998) Deutsche-Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements and Longer Run Dependencies *Journal of Finance* 53, 219-265.
- Betoneche, M.L. (1979) Spectral Analysis of Stock Market Prices *Journal of Banking and Finance* 3, 201-208.
- Baillie, R.T. and Bollerslev, T. (1991) Intra-Day and Inter-Market Volatility in Foreign Exchange Rates *Review of Economic Studies* 58, 565-585.
- Chan, K.C., Christie, W.G. and Schultz, P.H. (1995) Market Structure and the Intraday Pattern of Bid-Ask Spreads for NASDAQ Securities *Journal of Business* 68(1), 35-60.
- Ding, D.K. and Lau, S.T. (2001) An Analysis of Transactions Data for the Stock Exchange of Singapore: Patterns, Absolute Price Change, Trade Size and Number of Transactions *Journal of Business Finance and Accounting* 28, 151-174.
- Durlauf, S.N. (1991) Spectral Based Testing of the Martingale Hypothesis *Journal of Econometrics* 50, 355-376.
- Fong, W.M. and Ouliaris, S. (1995) Spectral Tests of the Martingale Hypothesis for Exchange Rates *Journal of Applied Econometrics* 10, 255-271.
- French, K.R. (1980) Stock Returns and the Weekend Effect *Journal of Financial Economics* 8, 55-69
- Gardner, W.A. (1985) *Introduction to Random Processes and Application to Signal and Systems* Macmillan: New York.
- Gardner, W.A. (1994) Ed., *Cyclostationarity in Communications and Signal Processing* IEEE Press: New York.
- Gardner, W.A. and Franks, L.E. (1975) Characterisation of Cyclostationary Random Signal Processes *IEEE Transactions on Information Theory* IT-21, 4-14.
- Gibbons, M.R. and P.J. Hess (1981) Day of the Week Effects and Asset Returns *Journal of Business* 54, 579-596
- Gladyshev, E.G. (1961) Periodically Correlated Random Sequences *Soviet Math. Dokl.* 2, 385-388.
- Granger, C.W.J. (1966) The Typical Spectral Shape of an Economic Variable *Econometrica* 34(1), 150-161.
- Granger, C.W.J. and Hatanaka, M. (1964) *Spectral Analysis of Economic Time Series* Princeton University Press, Princeton, NJ.
- Granger, C.W. J. and Morgenstern, O. (1963) Spectral Analysis of Stock Market Prices *Kyklos* 16, 1-27.
- Harris, L. (1986) A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns *Journal of Financial Economics* Vol. 16, pp. 99-118.
- Henriksson, R.D. and Merton, R.C., (1981) On Market Timing and Investment Performance: II. Statistical Procedures for Evaluating Forecasting Skills *Journal of Business* 54, 513-533.
- Hilliard, J.E. (1979) The Relationship Between Equity Indices on World Exchanges *Journal of Finance* 34, 103-114.
- Hinich, M.J. (2003) Detecting Randomly Modulated Pulses in Noise *Signal Processing* 83, 1349-1352.
- Hinich, M.J. (2000) A Statistical Theory of Signal Coherence *Journal of Oceanic Engineering* 25, 256-261.
- Ito, T. and Roley, V.V. (1987) News for the US and Japan: Which moves the yen-dollar Exchange Rate? *Journal of Monetary Economics* 19, 255-277.
- Jaffe, J.R. and R. Westerfield (1985) The Weekend Effect in Common Stock Returns: The International Evidence *Journal of Finance* 40, 432-454
- Keim, D.B. and R.F. Stambaugh (1984) A Further Investigation of the Weekend Effect in Stock Returns, *Journal of Finance* 39, 819-840.
- Merton, R.C. (1981) On Market Timing and Investment Performance I. An Equilibrium Theory of Value for Market Forecasts *Journal of Business* 54, 363-406
- Pesaran, M.H. and Timmerman, A. (1994) A Generalisation of the Non-Parametric Henriksson-Merton Test of Market Timing *Economics Letters* 44, 1-7.

- Pesaran, M.H. and Timmerman, A. (1992) A Simple Non-Parametric Test of Predictive Performance *Journal of Business and Economic Statistics* 10(4), 461-465
- Rogalski, R.J. (1984) New Findings Regarding Day of the Week Returns over Trading and Non-trading Periods: A Note *Journal of Finance* Vol. 39, pp. 1603-1614.
- Smirlock, M, and Starks, L. (1986) Day-of-the-Week Effects and Intraday Effects in Stock Returns *Journal of Financial Economics* Vol. 17, pp. 197-210.
- Smith, K.L. (1999) Major World Equity Market Interdependence a Decade After the October 1987 Crash: Evidence from Cross-Spectral Analysis 26, 365-392.
- Upson, R.B. (1972) Random Walk and Forward Exchange Rates: A Spectral Analysis *Journal of Financial and Quantitative Analysis* 7(4), 1897-1905.
- Wood, R.A., McUnish, T.H. and Ord, J.K. (1985) An Investigation of Transactions Data for NYSE Stocks *Journal of Finance* 40(3), 723-739.
- Yadav, P.K. and Pope, P.F. (1992) Intra-week and Intraday Seasonalities on Stock Market Risk Premia: Cash and Futures *Journal of Banking and Finance* 16, 233-270.

Table 1: Summary Statistics for Half-Hourly Exchange Rate Returns

	DEM_ JPY	GBP_ DEM	GBP_ USD	USD_ CHF	USD_ DEM	USD_ ITL	USD_ JPY
Mean	3.4E-4	9.7E-4	5.6E-4	8.6E-4	4.5E-4	-2.1E-4	6.5E-4
Variance	6.5E-3	4.6E-3	4.8E-4	8.5E-3	5.1E-3	9.0E-3	6.2E-3
Skewness	-0.049	-0.004	-0.167	-0.156	-0.190	-0.011	-0.019
Kurtosis	5.642	83.516	13.014	79.408	11.200	15.719	9.723
Minimum	-0.707	-1.966	-1.137	-2.431	-0.698	-0.924	-0.770
Maximum	0.659	1.992	1.203	2.403	0.777	0.966	0.758
acf lag 1	-0.198	-0.306	-0.205	-0.189	-0.097	-0.315	-0.150
acf lag 2	-0.013	-0.0053	-0.001	-0.004	0.005	-0.019	-0.002
acf lag 3	0.008	0.007	-0.000	0.002	0.015	-0.000	0.005
acf lag 4	-0.006	0.000	0.004	-0.009	-0.005	-0.005	-0.008
acf lag 5	0.006	0.004	-0.000	0.032	0.002	-0.017	-0.005
LB-Q(10)	500**	3144**	536**	476**	129**	1261**	288**
ARCH(4)	601.1**	23.6**	1355**	2559**	693**	1616**	462**
BJ Norm	4E+5**	2E+10**	9E+4**	3E+6**	7E+4**	9E+4**	2E+4**
BDS	32.47**	30.68**	41.00**	37.52**	38.95**	44.12**	33.27**
% zeros	7.5	6.1	5.0	5.7	5.1	5.4	4.9

Notes: Kurtosis represents excess kurtosis, LB-Q(10) is a Ljung Box test for autocorrelation of all orders up to 10, and is asymptotically distributed as a χ^2 (10) under the null hypothesis; ARCH(4) is Engle's (1982) Lagrange multiplier test for ARCH which is asymptotically distributed as a χ^2 (4); BJ norm is the Bera Jarque normality test, which is asymptotically distributed as a χ^2 (2) under the null of normality; BDS is the Brock, Dechert, and Scheinkman (1987) test for iid, which is distributed asymptotically as a standard normal under the null (statistic shown is for $m = 5$ and $\varepsilon / \sigma = 1$); % zeros gives the percentage of returns that are zero (i.e. no price change).

Table 2: P-values for Joint Test of Null Hypothesis that there is no signal coherence for all 120 Frequencies

	DEM_ JPY	GBP_ DEM	GBP_ USD	USD_ CHF	USD_ DEM	USD_ ITL	USD_ JPY
<i>p</i> -value	0.1532	0.0221	0.5342	0.2574	0.5109	0.0772	0.0000

Table 3: Periodicities with Coherence Statistics that are Significant at the 1% Level

Period (in units of half an hour)	Log Spectrum (dB)	Coherence statistic	Coherence statistic <i>p</i> -value
<u>Panel A: DEM_JPY</u>			
120.0 (half a week = 60 hours)	-2.164	0.314	0.003
4.068 (2 hours & 2 minutes)	-0.165	0.307	0.004
<u>Panel B: GBP_DEM</u>			
12.00 (6 hours)	-1.357	0.312	0.003
4.286 (2 hours & 8 minutes)	-0.775	0.329	0.002
3.288 (1 hour & 38 minutes)	-1.534	0.317	0.003
<u>Panel C: GBP_USD</u>			
24.00 (12 hours)	-3.688	0.302	0.005
16.00 (8 hours)	-2.882	0.295	0.007
<u>Panel D: USD_CHF</u>			
30.00 (15 hours)	-1.320	0.336	0.001
16.00 (8 hours)	-2.110	0.345	0.001
7.059 (3 hours & 32 minutes)	-0.803	0.299	0.006
4.000 (2 hours)	0.096	0.283	0.010
<u>Panel E: USD_DEM</u>			
30 (15 hours)	0.390	0.352	0.001
16.00 (8 hours)	-0.365	0.331	0.002
<u>Panel F: USD_ITL</u>			
24.00 (12 hours)	-4.166	0.308	0.004
11.43 (5 hours & 43 minutes)	-3.135	0.329	0.002
2.824 (1 hour & 49 minutes)	0.139	0.353	0.001
<u>Panel G: USD_JPY</u>			
120.0 (half a week = 60 hours)	-1.669	0.341	0.001
30.00 (15 hours)	-0.956	0.362	0.000
16.00 (8 hours)	-1.399	0.349	0.001
5.217 (2 hours & 37 minutes)	0.046	0.298	0.006
2.087 (1 hour & 5 minutes)	0.907	0.297	0.006

**Table 4: Total Profit of Trading Rules derived from Coherent Part of Signal
For 6-months: 1 July - 31 December 1996**

Rule	DEM_ JPY	GBP_ DEM	GBP_ USD	USD_ CHF	USD_ DEM	USD_ ITL	USD_ JPY
<u>Panel A: Long Trades Only</u>							
Coherence rule total 6 month return (long only)	7.16%	17.02%	6.70%	15.75%	6.99%	7.75%	18.43%
Standard deviation of returns (%)	5.59	5.80	4.91	6.76	5.00	7.30	5.59
Largest ½-hour gain	0.51	1.99	0.72	0.66	0.62	0.99	0.62
Largest ½-hour loss	0.56	0.64	0.85	0.72	0.59	0.92	0.45
Sharpe Ratio	0.012	0.038	0.012	0.030	0.013	0.010	0.044
Percentage of randomised rules with higher return	0.00	0.11	38.05	0.13	0.40	1.42	0.00
Percentage of correct direction predictions	53.09	53.76	54.32	53.05	54.03	55.03	55.50
Percentage of correct predictions that would arise if the buy-sell signals and returns were independent	50.14	50.04	50.06	50.06	49.99	50.03	50.05
Pesaran-Timmerman statistic	4.70**	5.94**	6.80**	4.78**	6.51**	7.99**	8.69**
Number of round trip trades per week	52	46	15	21	10	59	52
<u>Panel B: Long and Short Trades</u>							
Coherence rule total 6 month return (long and short)	9.28	23.20	4.55	24.56	12.52	15.62	31.84
Percentage of randomised rules with higher return	0.00	0.02	22.25	0.06	0.65	1.60	0.00
Standard deviation of returns(%)	7.79	8.41	7.34	10.31	6.79	9.67	7.94
Largest ½-hour gain	0.51	1.99	0.72	0.66	0.62	0.97	0.62
Largest ½-hour loss	0.56	0.64	0.85	0.72	0.59	0.92	0.45
Sharpe Ratio	0.013	0.038	0.004	0.033	0.022	0.021	0.058
<u>Panel C: Buy-and-hold foreign currency rule results</u>							
Buy-and-hold foreign currency	5.02	10.84	8.85	6.94	1.45	-0.11	5.02
Bid-ask spread	0.054	0.054	0.048	0.056	0.041	0.069	0.052
Standard deviation of returns(%)	7.79	8.42	7.34	10.31	6.80	9.67	7.96
Sharpe Ratio	0.004	0.015	0.013	0.006	-0.003	-0.005	0.042

Note: Bid-ask spread is taken as the difference between the bid and ask, divided by the average of the bid and ask, and multiplied by 100 to express is as a percentage of the exchange rate, and the average is taken over the sample. The Pesaran-Timmerman statistic is asymptotically distributed standard normal under the null of return and entry/exit timing independence. ** denotes significance at the 1% level.

Figure 1: Coherence against Period for USD_CHF, USD_DEM, USD_ITL, and USD_JPY




Figure 2: Coherence against Period for DEM_JPY, GBP_DEM and GBP_USD




Figure 3: DEM_JPY Coherent Part of the Mean Frame for a Week




Figure 4: GBP_DEM Coherent Part of the Mean Frame for a Week




Figure 5: GBP_USD Coherent Part of the Mean Frame for a Week




Figure 6: USD_CHF Coherent Part of the Mean Frame for a Week




Figure 7: USD-DEM Coherent Part of the Mean Frame for a Week




Figure 8: USD_ITL Coherent Part of the Mean Frame for a Week




Figure 9: USD_JPY Coherent Part of the Mean Frame for a Week




Figure 10: Cumulative Profit for July-December 2001 using Long and Short Positions

