

From Market Micro-Structure to Macro Fundamentals: Is There Predictability in the Dollar Deutsche Mark Exchange Rate?

Robert Hillman and Mark Salmon*
Financial Econometrics Research Centre,
City University Business School
London

15th December 1999

Abstract

In this paper we apply the variogram to the analysis of irregularly spaced transactions data from a one week sample of the Reuters DM2000-2 electronic trading system for the Dollar D-Mark exchange rate. The variogram, defined as the variance of increments in the process, is the standard tool used in Geostatistics for the analysis of spatial data. This literature emphasises the role of *intrinsic random functions* and *intrinsic stationarity* which generalise notions of integrated processes and second order stationarity familiar to Econometricians to irregularly spaced data. Moreover the variogram remains well defined over a wide range of non-stationary processes unlike the acf and has better sampling properties than the standard autocorrelation function. In our empirical study we find some evidence of deviation from unit-root behaviour in very short term, that is not apparent by autocorrelation analysis. We also examine the cross-dependencies between a number of micro-structural variables and dynamics of these variables around two interesting news events during the week under study. It is also possible to move directly from an estimated variogram which captures the second order properties of the process to MMSE predictors (kriging) without the need to develop intermediate structural models to represent the underlying DGP.

*We would like to thank Richard Payne for invaluable help with the DM2000-2 data set and also Reuters staff for explaining various ambiguities in the data and Noel Cressie for several discussions relating to the interpretation and significance of intrinsic random functions. We are also grateful to Ian Marsh and Steve Satchell for comments made at earlier presentations of some of this material in seminars at Cambridge and the Institute of Finance at City University Business School.

1 Introduction

This paper is concerned with the detection of autoregressive structure or temporal dependence and hence predictability in high frequency foreign exchange rates determined on the DM2000-2 electronic trading system. We are faced with two immediate problems; the choice of statistical technique for the analysis of a very large sample of irregularly spaced data and secondly the question of which economic theory to call upon to model the data.

Despite the growth of market micro structure theory in this area the theoretical issue is not so easily resolved as it might appear since it raises the question of how to rationalise the large body of existing, “macro” results generated from daily, weekly... and other temporally aggregated data with those generated from transactions data. It seems that a *different* economic theory and *different* conditional information sets need to be brought to bear when modelling an exchange rate on a transactions time scale and when modelling the long run or equilibrium exchange rate using say, quarterly data. This may seem reasonable since critical information on the instantaneous structure of the order book or bid-ask spread is lost when aggregated to even a daily basis and hence is likely to be irrelevant when attempting to model the long run equilibrium value of the exchange rate. Similarly new information on macro fundamentals is simply not available on the second by second time scale relevant to modelling transactions data. However both approaches purport to model the *same* underlying data and hence DGP. We are faced with a fundamental issue of consistency and at some stage have to ask if there is any intellectually coherent smooth transition between the two modelling approaches. While we may be willing to assume that there is, in principle, a single data generation process for the Dollar- Dmark exchange rate it seems that there will not be a single econometric model or single statistical representation of that *DGP* that we would want to use for all purposes.

The degree of predictability of a single given variable may of course vary across different conditioning information sets and time horizons as the relevant model also varies. It is, for instance, widely believed that exchange rates are unpredictable in the very short run with expectational forces and the heterogeneity of trader’s beliefs and objectives driving the market and yet there is a growing body of evidence that *fundamental* information, in various forms, can serve as an attractor for *equilibrium* exchange rates, see for instance MacDonald(1998).

Market micro structure theory does not imply random trading and market efficiency does not imply randomness in the *realised* exchange rate unless the expected equilibrium exchange rate is constant (see for

instance Mussa (1990)) so given systematic and well established incentives for market traders and systematic views as to where the equilibrium exchange rate lies, relative to the current value, it is not unreasonable to expect that we should find a degree of predictability in the very short run *given* a suitable conditioning information set ¹. While traders clearly have different objectives at different times they will trade in different markets such as the forward market if they are explicitly taking a long term position in which fundamental information is likely to be more relevant than the spot market such as DM2000-2. So we might expect to see little direct impact of fundamentals in DM2000-2 data, except at times of major announcements which would allow the market to adjust it's view of the equilibrium until the next news on fundamentals arrives². One way in which we might see views of the equilibrium rate being expressed in the transactions data set could be in the shapes and positions of the bid and offer curves as traders take limit order positions away from the current market price in anticipation of the market moving towards their own "equilibrium" view. Whether such behaviour is driven by fundamentals or motivated by direct trading concerns would seem to be difficult if not impossible to determine but clearly looking at the way the bid and offer curves move around major announcements will be important.

Given these arguments regarding the mixture of economic forces that are likely to be at work in the market we have decided to adopt the sim-

¹Although this degree of predictability might be lost almost as a statistical artifact through aggregation in the mid-horizon as the signals from the transaction based conditioning information sets are aggregated and become imperfect only to be recovered again in longer term models where the fundamental information dominates.

²Cheung, Chin and Marsh (1999) in a recent survey of Foreign Exchange traders provide some interesting observations on the nature of the foreign exchange market which are useful to keep in mind in what follows:

- 50.6% of trades are via electronic brokers, 67.7% are interbank, 32.3% are customer business
- 37 technical trading rule driven,41 fundamental analysis driven,36 driven by customer orders,40 jobbing (in and out)
- fundamental news is largely assimilated within 1 minute where the most important fundamentals are seen to be,interest rates, inflation and unemployment.
- 97% believe *intraday* movements in forex do not accurately reflect changes in fundamentals but within 6 months only 38 % and then 12 % over six months.
- on a scale of 1 to 5 traders believe the market is *predictable* to the following degree,2.20 (intraday), 2.94 (within 6 months) 2.89 (over six months)

plest atheoretical modelling strategy at this stage which is to simply to investigate if there is any autoregressive dependence in the DM2000-2 data either from within the transacted prices themselves or more generally given other micro market information. This leads us to our second immediate concern which is what statistical techniques should we employ to capture temporal dependence in high frequency irregularly spaced transactions data.

The standard approach has been to aggregate such data into blocks of fixed intervals of time and use standard time series tools such as the autocorrelation function. Using this approach Jon Danielsson and Richard Payne (1999, Table 3) have found using 20 second aggregated intervals on the “same” data set as ourselves that the D2000-2 returns are essentially *uncorrelated* in that the only significant autocorrelation coefficient is from the overnight 6.00pm-6.00 am period. Table(1) reproduces their results for convenience. They also suggest this result is consistent with efficiency in the DM2000-2 market.

period.

Time	ρ_1	Q(5)
6am to 8am	-0.03	3.81
8am to 10am	-0.02	3.41
10am to 12pm	0.02	6.81
12pm to 2 pm	-0.05	5.18
2pm to 4pm	0.00	3.03
4pm to 6pm	-0.04	6.42
6pm to 6am	-0.13*	270.51*

Table 1: First order autocorrelation coefficient, DandP (1999)

In what follows we explore this conclusion further by considering the difficulties in using the autocorrelation function with irregular data and suggest the use of different techniques, notably the *variogram* in preference to the autocorrelation function when considering irregular data.

In the next section we discuss the concepts of *intrinsic stationarity* and *intrinsic random functions* and explain how they extend our standard notions of integrated processes and second order stationarity to irregular data. One implication of this discussion is that the simple first difference transformation may not be able to induce second order stationarity with irregular data and a more general, in fact a *generalised increment vector*, may be defined which effectively annihilates both stochastic and generalised polynomial time trends. The class of intrinsically stationary processes is wider and encompasses the class of

second order stationary processes. We then review several statistical difficulties with the use of the autocorrelation function on temporally aggregated data before turning to apply the variogram and cross variogram to the DM2000-2 data. Finally we turn to consider the impact of macro news on the micro market characteristics and the return process on the exchange rate itself to investigate “regimes of predictability” within the period as discussed for instance by Guarda and Salmon(1997), Hillman(1998) and Dacco and Satchell (1998), amongst others. In an appendix we explore the use of kriging (the term used in the geostatistical literature for ‘optimal prediction’) on simulated and real series.

2 Intrinsic Stationarity and Intrinsic Random Functions

Matheron(1973) introduced the concepts of *intrinsic stationarity* and *intrinsic random functions* in the context of the statistical analysis and subsequent prediction of mineral deposits where non-stationarity appears to be as commonplace as it is in economics. Spatial statistics treats observations as arising from some general continuous multi-dimensional coordinate system and while we are not immediately interested in the flexibility offered by *Random Field Theory* or indexing our observations in 3 dimensional space we are interested in exploiting the tools used by geostatisticians for the analysis of irregularly spaced data³. The geostatistical method differs in several important respects from econometric time series analysis. In the first place the main emphasis is on the use of the variogram rather than the autocorrelation function (acf). There seem to be three reasons for this choice. In the first place the variogram is well defined for a wider class of stochastic processes than the autocorrelation function and hence allows us to legitimately consider the temporal dependence of processes that are not second order stationary. The variogram is well defined for all intrinsically stationary processes which is a wide class which encompasses second order stationary processes and a range of processes that are *not* second order stationary. Secondly the approach allows us to consider generalised transformations to stationarity or extensions of simple differencing with regularly spaced data to irregularly spaced data such as that generated from point processes. The *variogram cloud* provides an indication of memory properties of the process at all potential lags whether they be integer or real valued corresponding to irregularly spaced data. Finally the statistical properties the variogram in terms of finite sample bias may be substantially better

³The following discussion is largely drawn from Matheron(1973), Cressie(1988) and Cressie(1991),

than those of the autocorrelation function. For our purposes then the variogram enables us to simultaneously consider the dynamic dependence in the series given a range of stationary and non-stationary DGPs for which the acf may not formally be defined.

We start by considering a general stochastic process

$$\{Z(s) : s \in D\}$$

defined in D a random set in Euclidean \mathcal{R}^d . The ability to draw observations from a d dimensional space is not one that we need to exploit but we shall use the fact that the data is indexed on the real line as opposed to a set of regularly spaced integers. We will also refer to irregularly spaced data as being observed at points $\{t_i : i = 1, \dots, n\}$ rather than the regular spacing of observations at $\{t_i = i : i = 1, \dots, n\}$.

The *variogram* is defined as the variance of the difference between two values of the stochastic process separated by some potentially non-integer valued distance h . Such increment processes have been studied for many years, see Kolmogorov(1941), Yaglom(55), Gel'fand and Vilenkin (1964). More familiar to econometricians will be Whittle (1983) and the classic Von Neuman Ratio (1941) apart from the recent literature on the use of Variance Ratio statistics for determining the presence of a unit root and independence in time series data, see Cochrane (1988), Poterba and Summers (1988), Diebold (1989), Richardson and Stock (1989), Lo and McKinley(1988). None of this more recent literature has it seems realised the deeper potential offered by the variogram in the analysis of irregularly spaced data and the power of the underlying concepts of intrinsic stationarity and intrinsic random functions.

A process is defined as being *intrinsically stationary* if

$$E[Z(s+h) - Z(s)] = 0 \tag{1}$$

and

$$Var [Z(s+h) - Z(s)] = 2\gamma(h) \tag{2}$$

In other words the variance of the h increment is simply a function of h and not the time origin. The function $\gamma(h)$ is known as the *semi-variogram*. The important properties of an intrinsic stationary process are defined in terms of the *increments*, $(Z(s+h) - Z(s))$. Comparing this definition with *second order stationarity* we need to consider both mean and variance stationarity.

The process is *mean stationary*, if

$$E[Z(s)] = \mu \tag{3}$$

and *variance stationary*, if

$$\text{var}[Z(s)] = \sigma^2 = \kappa(0) \quad (4)$$

is well defined and constant and

$$\text{Cov}[Z(s), Z(s+h)] = \kappa(h) \quad (5)$$

is only a function of the lag interval. In this latter case we have

$$\text{var}[Z(s+h) - Z(s)] = \text{var}[Z(s+h)] + \text{var}[Z(s)] - 2\text{Cov}[Z(s+h)Z(s)]$$

Assuming second order stationarity and given

$$\kappa(0) = \text{var}[Z(s+h)] = \text{var}[Z(s)] = \sigma^2$$

we may write

$$\gamma(h) = \kappa(0) - \kappa(h)$$

showing that all second order stationary processes are intrinsically stationary. Notice that *only* in the case of second order stationarity will the autocorrelation function be well defined and there will be a simple relation with the scaled semivariogram

$$\rho(h) = \frac{\kappa(h)}{\kappa(0)} = 1 - \frac{\gamma(h)}{\kappa(0)}$$

The converse is not true however since as Cressie(1998) shows, if $\{W_t : t = 1, 2, \dots\}$ is a Wiener process observed at $t = 1, 2, \dots$ then $2\gamma(h) = \sigma^2 h$ ($h = 1, 2, \dots$) but $\text{Cov}(W_t, W_u) = \sigma^2 \min(t, u)$ which is not a function of $|t - u|$. Similarly a fractional Brownian motion (see Mandelbrot and Van Ness(1968)) where

$$2\gamma(h) = b \|h\|^\lambda \quad b > 0, 0 < \lambda < 2$$

is an intrinsically stationary process but not second order stationary. The variogram is more general than the acf and enables the second moment dependence of a wider class of time series to be characterised.

We can also see that when the variance of the process is well defined and $\lim_{h \rightarrow \infty} \kappa(h) = 0$ so the dependence in the process goes to zero as the lag interval increases then

$$\kappa(0) = \lim_{h \rightarrow \infty} \gamma(h)$$

and the asymptote of the semivariogram is the variance of the process when second order stationarity holds. When the process is not variance

stationary the variogram retains its usual interpretation unlike the acf but does not tend to an asymptote. The autocovariance is obviously not defined in this case. Diebold(1989), Cressie(1991) and Beran(1994) discuss the use of the variogram in detecting long memory and the direct estimation of the Hurst exponent from an estimated variogram.

As we have emphasised above even if the process is not variance stationary it may nevertheless be a valid *intrinsically stationary process* which is an *intrinsically random function* of order 0.

Following Cressie(1988) we suppose

$$Z = (Z_{t_1}, Z_{t_2}, \dots, Z_{t_n})$$

represent observations at irregularly spaced time points $\{t_i : i = 1 \dots n\}$. The if we define an $n \times d$ matrix X with i 'th row

$$\{1, t_i, t_i^2, \dots, t_i^{d-1}\}$$

Suppose that λ is an $n \times 1$ vector of real numbers satisfying $X'\lambda = 0$ then λ is called a generalised increment vector of order $(d - 1)$ and $\lambda'Z$ is a generalised increment of order $(d - 1)$. An intrinsic random function of order $(d - 1)$ is defined as any process, $\{Z_t : t \geq 0\}$ for which

$$V_u = \sum_{i=1}^n \lambda_i Z_{t_i+u} \quad u \geq 0$$

is second order stationary for any $\{t_i : i = 1 \dots n\}$ and any generalised increment vector, λ of order $(d - 1)$. Notice that with regularly spaced data the first difference transformation provides the weights for the generalised increment process and this can now be generalised to irregularly spaced data. What is effectively achieved in the irregularly spaced data case above is that a particular weighted combination of the observations, given by the generalised increment vector, generates a second order stationary process. The generalised increment vector annihilates both the nonstationary mean (polynomials in time) and stochastic trends. The parallels with integrated processes are clear and the intrinsic random function idea appears to be more general and applicable to irregularly spaced data. Using results from Gel'fand and Vilenkin (1964), Mathéron(1973) showed that any intrinsic process of order d possesses a *generalised covariance matrix* $K(h) : h \in R$ such that given data Z and generalised increment vectors, λ and v (of order $d - 1$)

$$cov(\lambda'Z, v'Z) = \sum_{i=1}^n \sum_{j=1}^n \lambda_i v_j K(t_i - t_j)$$

This result can be used when considering prediction directly from the second moment structure of these intrinsically stationary processes (see appendix 3).

3 The Autocorrelation Function with Irregular Observations

We now turn to consider the use of the autocorrelation function in measuring the dependence in an irregularly spaced time series. The obvious problem is that the easy interpretation of autocorrelation functions on regularly spaced data is lost as we move to irregularly spaced data. Given regular data we can calculate directly the effect of a shock today on some variable h days ahead. However when we have irregular data we lose the natural time scale and although we can compute autocorrelations in the normal way we cannot easily interpret the results and will not know, for instance, whether a shock will reach its half life in the next minute or in the next day. Operational time autocorrelation will have different implications in terms of calendar time depending on the time of day or state of the market and it is difficult if not impossible to use operational time autocorrelations on the clock time scale.

This issue has been considered by Quenouille (1958), Brillinger (1972) and Clark (1975) amongst others. Autocorrelations computed in irregular operational time will be weighted averages of the autocorrelations on a regular scale. There are two ways forward; the first and by far the most common is simply to aggregate the irregular data into fixed intervals and the second is to use some form of time deformation transformation to extract a regular acf from the irregular acf. We need to emphasise here that we are not dealing with an underlying continuous time process but a discrete point process so the alternative of treating the realised data as irregular observations on an underlying diffusion process is not one we wish to exploit.

If we write the standard autocovariance estimator (assuming a zero mean for the moment) as

$$\kappa_n^e(j) = \frac{1}{n-j} \sum_{i=j+1}^n Z(t_i)Z(t_{i-j})$$

based on the observed but irregular data $Z(t_i)$, then Clark(1975) shows that this estimator is a mean square consistent estimator of

$$E[\kappa_n^e(j)] = \int \kappa(v)dF_j(v)$$

where the distribution function $F_j(\Delta t)$ of the operational time intervals $(t_{i+j} - t_i)$ provides a weighting on the true autocovariance function, $\kappa(v)$. So in principle if the distribution of the durations, $(t_{i+j} - t_i)$ was known we could work back from this expression to an estimate of the underlying autocovariance structure.

3.1 Time Deformation

The idea behind time deformation is to construct a suitable transformation of clock time into economic time so that each interval of clock time captures the same quantity of information about the underlying stochastic process. Brillinger (1972) proposed an estimator in which he used the observed intervals between the observations, v_i , to estimate the autocovariance function for the stationary point process which generates the intervals which could then be used to improve the estimate of the autocovariance function of the underlying transactions. This estimator is relatively inefficient as the observations are not related to the operational time gap between them. Clark proposed an estimator based on specifying a poisson process which effectively mapped operational time to clock time and enabled direct estimation of a continuous AR(1) model from the irregular data. Ghysels, Gouriéroux and Jasiak, (1996) have approached the problem in a similar way by constructing a transformation that effectively attempts to translate an irregular sample in clock time into a regular spaced sample in economic time. They are interested in measuring the autocovariance structure in economic time and work from acf calculated on aggregated clock time to imply the acf in economic time. In particular, given some directing process, W , that associates clock time with economic time

$$W : t \in \mathbb{N} \longrightarrow W_t \in R^+$$

and that the process of interest is evolving in economic time as

$$X^* : w \in R^+ \longrightarrow X_w^* \in R^+$$

we may induce the process observed in clock time $t \in \mathbb{N}$ by considering

$$X_t = X^* \circ W_t = X_{w_t}^*$$

Given an assumption regarding the nature of the time deformation process, W , for instance if volume measuring information flow, determines the way in which the natural time scale can be stretched and contracted, we can consider two autocorrelograms in terms of clock time

and an irregular spaced intrinsic time. The autocorrelogram in an irregularly spaced data set may then be computed through a kernel smoothing procedure. Two different kernels were used by Ghysels, Gouriéroux and Jasiak to estimate the acf's for daily returns on the NYSE with trading volume used as a directing process.

lag	Standard acf	Gaussian kernel	Bounded kernel
0	1.00	1.0	1.0
0.25	-	.96	.989
0.75	-	.71	.02
1	-0.07	0.556	-0.0566
1.5	-	0.29	-0.02
2	-0.03	0.126	-0.02

Table 2: Jasiak, Ghysels and Gouriéroux Time deformed kernel acf's

The lags for the standard acf are based on the daily observations and for the kernel autocorrelations lag x corresponds to $w = xm_v$, where m_v is the average daily trading volume. This approach works in the opposite direction from that we are interested in, ie. it attempts to move from an aggregated daily acf to the underlying irregularly spaced autocorrelation structure. Never-the-less it is striking from the table (2) how the daily acf can give quite a different impression of the underlying autocovariance structure in real time from the time deformed acf's with the daily acf indicating relatively little time dependence while both time deformed acf's show significant autoregressive structure.

It is clearly important, given these results, to recognise the significant effect temporal aggregation potentially has on the measurement of the autoregressive structure. The critical uncertainty in this process lies in the choice of time deformation transformation. It is also quite apparent that the two different kernel choices deliver quite different estimated autocorrelograms at specific discrete lags.

3.2 Regularly Sampled Autocorrelation

Let us now consider the issue of which value to take as representative for an aggregated interval. The different first order autoregressive coefficients that could be drawn from our DM2000-2 data set with transactions aggregated into 20 second intervals by selecting different representative values is shown in the following table. The two figures below Fig(1) and Fig(2) show how the choice affects the surrogate time path created for the aggregated series. If the right hand value is used then that value will be the value assumed for the aggregated series, shown with the circles, until the next transaction occurs, similarly with the left hand value.




Figure 1:

The solid line in the graphs provides the linear interpolation value which assumes that there will be transactions continuously over any interval between actual transactions.

The different time paths implied affect the estimated first order autocorrelation coefficients as shown in table (3)

Figures

Time Period	LH INT	RH INT	LIN INT	DeJ-N
0-24	-0.22	-0.24	0.03	
18-6	-0.28	-0.41	0.12	0.29
6-8	-0.09	-0.10	0.01	-0.15
8-10	-0.18	-0.12	-0.19	-0.14
10-12	-0.15	-0.07	-0.03	-0.38
12-14	-0.18	-0.12	-0.11	0.03
14-16	-0.18	-0.11	-0.08	-0.00
16-18	-0.08	-0.06	0.23	0.05

Table 3: First Order ACF Coefficients for Various Interpolation Methods and De Jong and Nijman

This table shows clearly how inference regarding the existence of autoregressive structure could be seriously affected by a particular choice of representative element in the aggregated interval⁴. We also note the

⁴We believe that the results reported in Danielsson and Payne(1999) and shown




Figure 2:

radically different behaviour displayed by the DeJN method in this comparison and shows how sensitive even a consistent interpolation method will be to these assumptions on selecting representative values if we seek to aggregate the data into regular intervals.

Given these ambiguities and the different estimates of autoregressive structure that could be generated by different choices of time deformation transformation, interpolation method or representative value in aggregated interval we have chosen not to focus on the autocorrelation function in our analysis. Instead we have decided to explore the direct use of the variogram to measure autoregressive structure in the irregularly spaced sample of transaction on the Dollar-Dmark exchange rate.

4 The Variogram and Cross Variogram with Irregularly Observed Data

Under the intrinsic stationarity assumption the classic variogram estimator (defined here for the moment on regularly spaced data) is given by

$$\hat{\gamma}(h) = \frac{1}{2}(n-h)^{-1} \sum_{i=1}^{n-h} (x_{i+h} - x_i)^2$$

above in table (1) used the right hand value as our acf results are similar to their own in this case.

and is unbiased. The unbiasedness follows since critically this estimator only uses differences of the data and since it does not require estimates of the mean unlike the autocorrelation. We will use a smoothed version of the variogram estimated at different lags h

$$\hat{\gamma}(h) = \frac{1}{2} \frac{1}{|N(h)|} \sum_{N(h)} (x(s_i) - x(s_j))^2 \quad (6)$$

where the sum is over $N(h) \equiv \{(i, j) : s_i - s_j = h\}$ and $|N(h)|$ is the number of distinct elements in $N(h)$. In principle we could simply calculate the value of $\hat{\gamma}(h)$ for every h that occurs in the data⁵, but we choose to use h 's defined on a regular grid. First we construct a tolerance window around each increment, h , so that all pairs satisfying

$$|s_i - s_j - h| \leq tol$$

are found and

$$N(h) \equiv \{(i, j) : |s_i - s_j - h| \leq tol\}$$

Clearly we can only reliably estimate the variogram at lags for which we have enough observations and so we use a cut off rule, similar to that suggested in Cressie(1991), that $N(h) \geq 30$. We have also experimented with different smoothing rules such as using a fixed number of neighbours and the results are not dramatically different from those reported below. More sophisticated variable bandwidth kernel and robust estimators of the variogram could be considered, which we will explore in the future.

As an example figure(3) shows the irregular variogram estimated on the DM2000-2 data over one day between 16.00 and 18.00 using steps of 20 seconds and 1 second and a tolerance of 0.02 seconds for both. The horizontal axis shows the lag interval in seconds. We can see that as the interval becomes large there are less observations in both variograms but more easily in the 20 second variogram. Note that the transactions in DM2000-2 are recorded on a one hundredth of a second time scale. There is no interpolation or choice of representative value involved in these calculations beyond the averaging implied in using a form of smoothing estimator. The average is typically taken over several thousand observations at each lag. Where insufficient observations exist the variogram is simply not computed at that value. Basically the same pattern is seen for the two variograms

⁵In the geostatistics literature this sort of plot is called the variogram cloud.




Figure 3:

4.1 Discriminating Processes with the Variogram

The basic variogram we will work with given in equation 6, and we will compare the estimated version with the distributions of the variogram under different processes. There is some asymptotic theory available. In particular Lo and MacKinlay (1988) derive the distribution for a scaled version of the variogram, $R(h) = V(h)/V(1)$ where $V(h) = \frac{n}{h(n-h)}\hat{\gamma}(h)$. Under standard normalization the resulting test statistic is distributed standard normal. However, as Kormendi and Meguire (19) demonstrate the finite sample distribution of such variograms can be seriously skewed, making the symmetric asymptotic distribution less useful. We therefore proceed by using simulated distributions. In order to facilitate easy comparison of the plots we use a simple scaled variogram $\hat{\gamma}(h)/\hat{\gamma}(h_1)$. h_1 is the minimum lag we distance we will estimate the variogram for. So, for example if we use a 20 second interval, then h_1 is 20.

In the following plot, figure(4) we provide a simulation⁶ where the dashed lines are estimated normalised variograms for two autoregressive processes together with the simulated empirical distributions (97.5 and 2.5 percentiles) for the unit root. We can see that the case with an AR parameter ρ of 0.95 cannot be distinguished from a unit root on this basis at the 95% level but $\rho = 0.8$ is clearly seen to be significantly different from a unit root. The empirical distribution is seen to be skewed, demonstrating the usefulness of using simulated distributions with limited sample sizes. Figure 5 shows the empirical distribution from

⁶We used a sample size of 200, and 100 replications of the experiment.




Figure 4: The empirical distribution for the variogram for a unit root process. The two dotted lines show the mean variogram over a simulation for AR(1) processes with parameter 0.95 (upper line) and 0.8.

another simulation study where the data is simply the differences from the unit root process, i.e. i.i.d. normal errors. We will use these types of simulated distributions to test hypotheses about the FX data later.

The Cross- variogram, which is useful in exploring cross temporal dependence, is defined for a vector of intrinsically stationary processes $\{\mathbf{Z}(s); (Z_1(s), Z_2(s), \dots, Z_k(s))\}$ defined on the continuous observation space

as

$$2\gamma_{ij}(h) = \text{var}(Z_i(s+h) - Z_j(s)) \quad h \in R^d$$

Similar estimators for the cross variogram as described for the variogram above have been used in the calculations reported below. Figure(6) shows the cross variogram between two series generated from the following simulated processes;

$$p(t) \sim N(0, 1)$$

$$q(t) = -0.8p(t - 10) + e(t)$$

$$e(t) \sim N(0, 1)$$

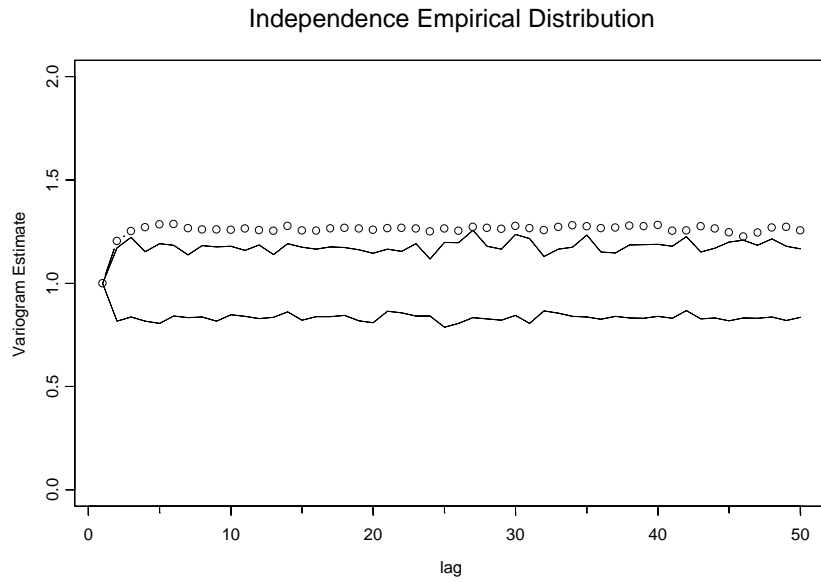


Figure 5: The empirical distribution for the variogram for i.i.d $N(0,1)$ errors. The dotted line show the mean variogram over a simulation for AR(1) processes with parameter 0.2

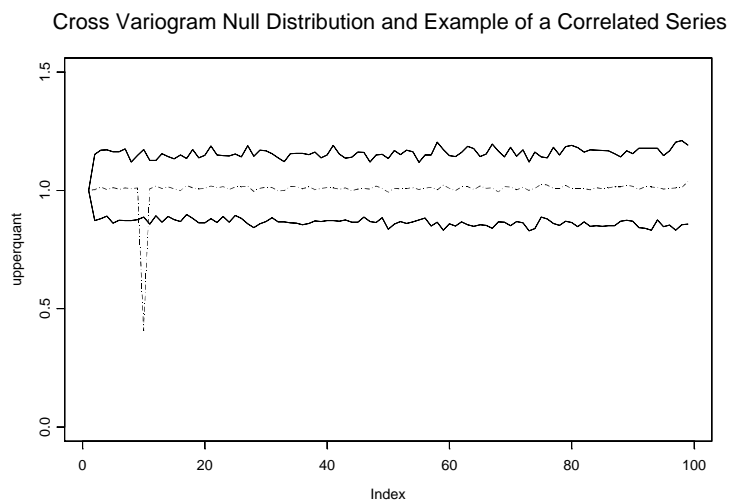


Figure 6:

The calculated cross-variogram is normalised as above to take the value of unity on zero cross-covariance. The 95% confidence intervals in this case are calculated by bootstrap methods under independent resampling yielding an approximate method for indicating interdependence. Note that this empirical distribution is the same as that of the variogram measuring dependence within a single i.i.d. process.

4.2 Sampling Properties of the Variogram and the Acf.

The sample autocovariance can be hopelessly biased when the process is not mean stationary and the autocovariance itself is not even defined when the process is not variance stationary. The standard autocovariance estimate is given by

$$\hat{\kappa}(h) = (n - h)^{-1} \sum_{i=1}^{n-h} (x_{i+h} - \bar{x})(x_i - \bar{x})$$

and the major problem with the use of the acf in practice is the need to estimate the mean of the process. It is well known that the variance of the sample mean for correlated data is not σ^2/n and in consequence $\hat{\kappa}(h)$ will be biased in finite samples, see Fuller (1976, p 236). Even if not explicitly realised a trend (\bar{x}) is implicitly estimated through the mean and the acf is calculated using the effective residuals $(x_i - \bar{x})$. Since the residuals are linearly related through $\sum_i (x_i - \bar{x}) = 0$, even if the $\{x_i\}$ are independent the residuals $\{x_i - \bar{x}\}$ will be negatively correlated and a finite sample bias induced in the estimated acf. Indeed the more pronounced the autocovariation the more pronounced the bias as shown for instance by Newbold and Agiakloglou(1993) in the case of fractional noise.

The sampling properties of $\hat{\kappa}(h)$ have been discussed and compared with those of the variogram by Cressie and Grondona(1992) and Haslett (1997). These results show that the classical variogram given above, is unbiased when the process is only mean stationary and a natural alternative estimator only has a small bias when the process even when the process is neither mean nor variance stationary. The alternative estimator, which is the sample variance of the differences where $\hat{\gamma}(h)$ is the average of their squares, is given by Haslett as

$$\tilde{\gamma}(h) = \frac{1}{2}(n - h - 1)^{-1} \sum_{i=1}^{n-h} (d_{hi} - \bar{d}_h)^2$$

where $d_{hi} = x_{i+h} - x_i$

5 The Data Set: Reuters DM2000-2

The DM2000-2 electronic dealing system provides a continuous auction market for dealers from major banks in which *market* and *limit* orders input by participants are matched automatically by the system with the orders of others. The data set was made available to us from Reuters via the FMG and provides a direct feed from the system showing the placement of orders, their price and volume and enabling the structure of the order book (which is not available in real time to the dealers using the system) to be calculated. Dealers using the system can observe only the best current bid-ask spread and volumes. A fuller analysis of how the raw data was transformed into a usable data set for analysis is described in appendix 1. The data used covers all Dollar/DM trades on DM2000-2 over the week 6th-10th October 1997. Unlike the Olsen FXDX data set which is drawn from the Reuters FX page and shows indicative quotes (bid-ask advertising) rather than transactions the data we have represents real transactions on what now represents some 40 to 50% of the market.

Some 30,000 transactions have been drawn from approximately 130,000 entries into the system which may either be a *market* order or a *limit* order, on either the bid or offer side of the market. We also have the entry and exit times of the order, the bid and ask prices, the actual transactions prices and the reasons for the withdrawal of the order. We do not have information on the identity of the banks making the order although this is known to the traders dealing on the system. Over this period the traded price ranges from 1.73DM/\$ to 1.77DM/\$. The average volume of the orders was 2.283 million and the volume of trades ranged from 1 million to 30 with a mean of 1.816. Total volume of transactions was about \$60 billion, 70,406 orders were cancelled before being acted upon and 38,239 were removed partially filled. A market emulation programme was written to extract the transactions data from the full data set⁷. Figure 7 shows the movement in the exchange rate over the entire week with the horizontal scale in seconds. The panels in this graph indicate the period 6.00pm to 6.00am GMT each day and although the system is open 24 hours a day obviously different behaviour applies during the overnight. This data has had some 20 observations (out of 30,000) removed. Several significant outliers could be seen in the original data and these required some explanation. Reuter's staff suggest that these represent real transactions since they are recorded on the

⁷A number of difficulties were found when extracting the transactions data from the continuous record and we are grateful for the help of Reuters personnel in resolving these problems.




Figure 7:

system but offer the following explanations for their presence; they may arise when a trader has left a limit order in the system and forgotten about it and it subsequently gets hit, or alternatively a trader simply needs to trade as a customer demands it or near the end of the day a trader has to close out their position and just takes what ever price he can get immediately. Since we are not interested in such behaviour at the moment we simply took these observations out of the series.

What is much more important for our purposes is to recognise the impact of macro-economic news on the market during the week. The general situation during the week can be described as uncertainty relating to tension in the Gulf and the prospects for European Monetary Union. The Italian government fell on Thursday (9/10/97) of this week as it failed to secure a compromise with its communist coalition partner on the proposed budget— aimed at satisfying the Maastricht criteria. Rather more significant for the evolution of the Dollar Dmark exchange rate was the uncertainty surrounding the Bundesbank's repo rate. On the same day as the Italian Government fell this central rate was increased by 30 basis points for the first time in five years. The timing was a surprise and the effect was dramatic, at least on the time scale of figure(7). The Financial Times of that week indicates substantial speculation surrounding such an increase and when it would occur. In fact speculation apparently existed that such an announcement would be made on the Tuesday (7/10/97) but the Bundesbank left the rate unchanged on that day. There is a clear appreciation of the DM in an-

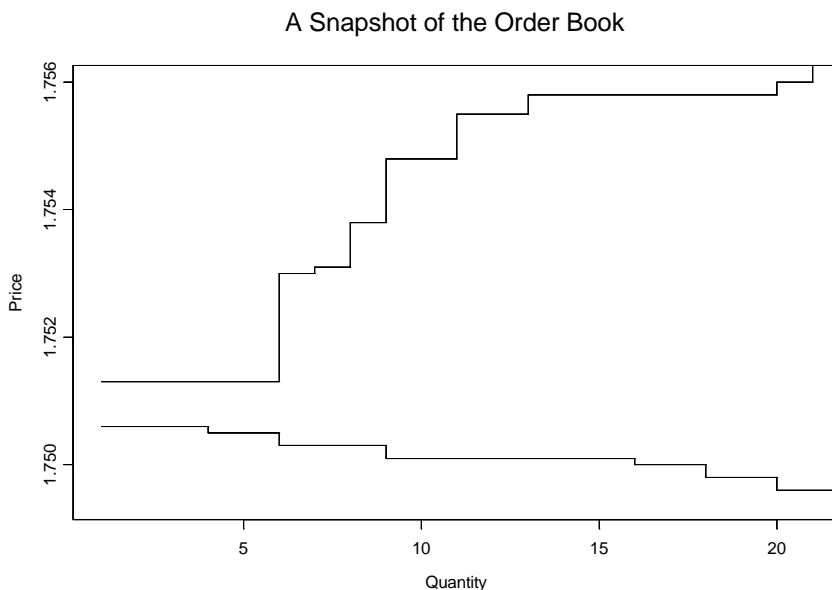


Figure 8: Sample Bid and Offer Curves

icipation of a higher repo rate up to the time of announcement on the 7th which is unwound slowly during the afternoon trading. We analyse the market around these two events in the last section of the paper.

Some stylised facts of the transactions prices and returns are reported in the following figures. Typical bid and offer curves are shown in figure(8) with limit orders placed in the system at particular prices and for particular quantities shown on the vertical and horizontal axes respectively. The inside spread in this particular case is relatively narrow but it can vary significantly during the day as shown in figure(9) where the horizontal axis is in 30 minute intervals from 6.00am. There is a spike at the middle of the day, though we suggest later this is in fact driven by the events on one day, the 9th October.

We have also constructed some measures of depth and liquidity in the market. The depth measure is simply the sum of all orders in the order book on both the demand and supply side. This gives us the total volume available on each side of the market. As a measure of liquidity we have estimated the slopes of both the bid and the offer curves locally to the market price and used their average as a measure of liquidity. The curves are in fact often nonlinear and also we find that sometimes there are a few orders sitting in the system that are at quite off market prices. We simply consider a linear approximation to the first part of

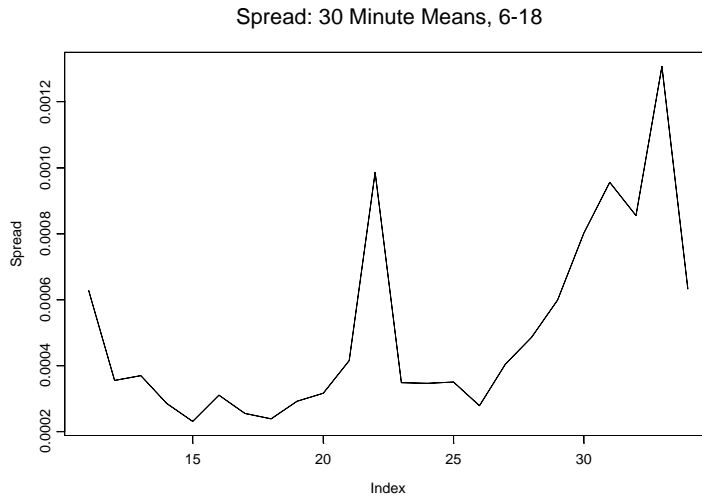


Figure 9:

the curve, looking only at the curve up to the first 5 different prices⁸. If we denote the different prices in the bid side as say bp_1, bp_2, bp_3 , etc., where $bp_n < bp_{n-1}$, and bp_1 is the best bid, with corresponding quantities bq_n , then our slope is $(bp_1 - bp_5)/(bq_5 - bq_1)$. This accords with a notion of liquidity which implies that in a liquid market you will be able to transact without substantial changes in the market price. We are also interested in other shape measures of the bid and offer curves such as the curvature but for the moment just concentrate on this simple measure. Figure(10) provides some indication of how the average slope (over the week) of the bid curve changes through the day (in 30 minute intervals). Not surprisingly there seems to be a U shape with the demand getting more elastic (shallower slope) during the busy part of the day. The offer curve has a similar intra-day pattern. As a measure of volatility we show the absolute return during the day in figure(11). Other seasonal (or diurnal) patterns have been documented by Daniellson and Payne (1999). These patterns confirm general patterns regularly observed in such transactions data.

6 Empirical Results

6.1 Variogram

⁸We also tried up to 10 different prices, but it made very little differences. We are currently considering more sophisticated ways to examine measures of market liquidity over time.



Figure 10:

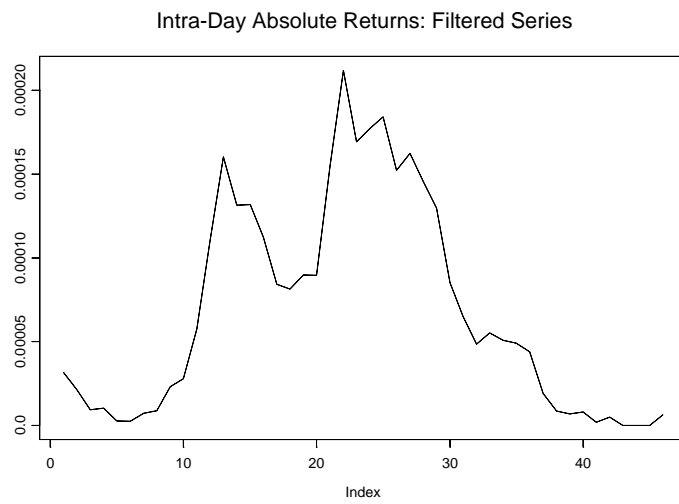


Figure 11:

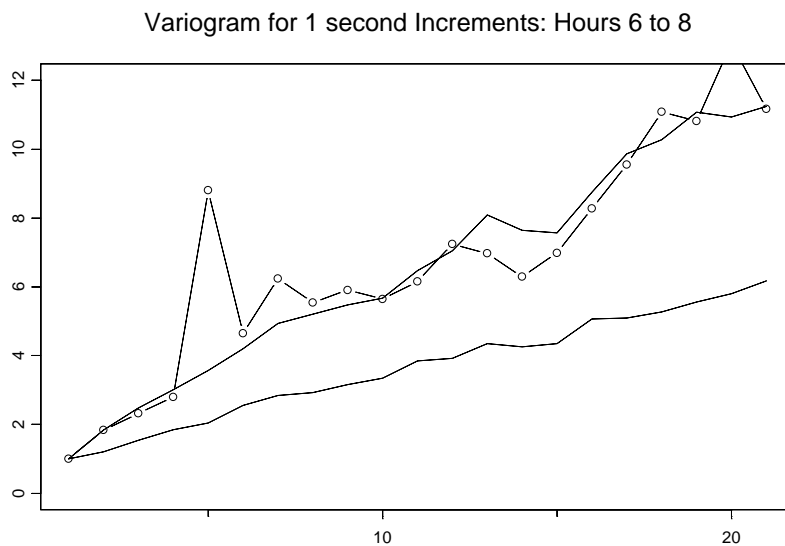


Figure 12:

We applied the scaled variogram to the DM2000-2 data set by splitting the day into 2 hour intervals as we might expect some differing patterns of dependence through the day. We used a lag interval of 1 seconds and a tolerance window of 0.1 seconds. Because of the density of the data, this means that even with such a small tolerance window we are using thousands of pairs of observations for each lag interval. We report a selection of results⁹ where we show 95% confidence intervals constructed in the following way. We retain the same time-scale as in the actual i.e. $s_1, s_2, s_3, \dots, s_N$, but our surrogate price series are generated by simulating a unit root process with *i.i.d.* errors¹⁰. We then plot the 2.5 and 97.5 quantiles from a simulation of 100 series all with the same sample size as the original series. This is not less than 5,000 for any series. Plots 12, 13 and 14 show the results for three typical sub-samples. We can see that the actual data appears to be inconsistent with that generated by a unit root process, with most of the variogram lying outside the confidence bands. The fact that the empirical variogram for hours 10 to 12 lies above the bands indicates that the sample is displaying explosive

⁹Results for other two hour segments are available from the authors. These results were however typical, in all cases we rejected the null of a unit root in the levels, and the null of independence in the returns.

¹⁰We have chosen to use simulated distributions for inference for the reasons explained earlier.

Variogram for 1 second Increments: Hours 8 to 10




Figure 13:

Variogram for 1 second Increments: Hours 10 to 12




Figure 14:




Figure 15:

dynamics¹¹. Next we considered what kind of inferences we might make if we looked at an autocovariance based measure of dependence instead of the variogram. We estimated the sample autocovariance in the same way as the variogram (using a tolerance window around each lag increment h , and normalizing by $k(1)$), and also generated confidence bands in the same way. Figures 15, 16 and 17 show the respective plots. Interestingly we do not see the rejection of the unit root hypothesis via the autocovariance, suggesting the variogram is in fact uncovering some deviation from unit root behaviour that the autocovariance is not picking up¹². These variograms are looking at very short term dependence, it would be of great interest to look at longer term dependencies and in particular how the apparent dependency patterns vary under aggregation. We know for example that it is hard to reject a unit root in the exchange rate at the daily level, and we conjecture that the temporal dependence (as measured by both the autocovariance and the variogram) varies as we aggregate up to lower frequencies.

¹¹We found with some exploratory simulations that a process with an AR(1) process with a parameter of 1.01 is enough to generate the same sort of divergence from the confidence band as in the data.

¹²We are currently pursuing further comparisons between the two measures for investigating dependence using both real data and simulated processes.

Autocovariance for 1 second increments: Hours 8 to 10




Figure 16:

Autocovariance for 1 second increments: Hours 10 to 12




Figure 17:




Figure 18:

6.2 Cross Variogram

Figures 18 to 20 reports the cross variogram between three micro variables and returns. The returns series is simply the differences in the consecutive (log) prices, so these do not correspond to the return over the same time interval, but instead are best thought of as the incremental price process. These cross variograms are estimated on the main day segment, between 6AM and 18PM. The confidence bands are generated by sampling from replacement from the original series, therefore providing a null of independence between the two series. For all three variables there seems to be evidence that each variables affects future returns for a short period. This we used a lag interval used for the estimation was 20 seconds, and the x-axis on the graph gives the lags in seconds. We can see that volume appears to have an effect on returns for up to 100 seconds ahead. Liquidity appears most weakly related, only just breaking out of the 95% confidence bands for about 60 seconds.

These variograms and cross variogram results seem to offer clear indication of temporal structure and hence predictability in DM2000-2, though clearly we are talking about very short term dependence. Whether a trader could exploit these dependencies is uncertain. Certainly a trader could not calculate the microstructural variables in real time as we have, but to the extent that liquidity is deterministically

Spread(t) to Returns(t+h)




Figure 19:

Liquidity(t) to Returns(t+h)




Figure 20:

predictable (on time of day for example) there may be some scope for exploiting the fact that short term dependencies may vary through the day. One might also try and construct real time variables that might give some indication of the depth of the market, like the number of transactions in the last n seconds, or the volume available at the spread prices. Of course traders themselves have a feel for the depth of the market, and may in fact throughout the day enter limit orders at off-market prices in order to gauge the depth.

7 Regimes of Predictability associated with Macro Events?

Finally we return to the issue raised in the introduction which was the question of how micro and macro models of the exchange rate may be rationalised. In particular we want to consider what impact macro events have had on the micro structure of the market, what persistence exists and whether a greater degree of predictability can be found around such events. In our one week sample we are fortunate to find two such events; on the 7/10/97 the anticipated but unfilled market speculation of a repo rate increase by the Bundesbank and then on the 9/10/97 when the repo rate was in fact increased by 30 basis points for the first time in five years. Regimes of predictability have been found in a number of earlier studies including our own previous work, Guarda and Salmon(1997) and Hillman(1998) and we believe the same ideas will hold on the transactions time scale. In this earlier work we have emphasised that the development of several local models with distinct regimes of behaviour is a more profitable route to follow than any attempt to construct a complex global model that attempts to account for several different modes of behaviour. Generally the study of such transitions is confused by temporal aggregation of the data but with the current data set we have market information down to the hundredth of a second.

Figure(21) shows the movement the exchange rate over the day of 7/10/97, the spread over the same period and the bid and offer slopes. We expect these slopes to become very sensitive as traders form differing views regarding the future equilibrium exchange rate. Figure (21) shows that the appreciation and subsequent depreciation of the exchange rate takes place more gradually in this time scale than appeared from figure(7). In the period prior to the usual announcement time we can see that the offer slope in particular becomes much steeper as the exchange rate is appreciating and indicating we suggest substantially different views emerging in the market regarding the probability of an interest rate increase and ultimately the value of a new equilibrium exchange rate. It is noticeable that at the time of non-announcement the slope

of the bid curve increases while the offer curve flattens out. There is considerably more uncertainty on the offer side of the market prior to the event than immediately afterwards.

The increase of the repo rate on 9/10/97 clearly had an immediate impact on the market which appears to have been largely unexpected as can be seen from figure(22). It also seems to have taken the market much longer, several hours in fact to digest the news and to find a new equilibrium than suggested in the survey by Cheung, Chin and Marsh(1999). It would seem to be clear from the subpanels in Figure(22) that the repo announcement led to an immediate appreciation of the DMark and an immediate substantial widening of the spread and again a dramatic increase at the time of the announcement on the slope of the offer curve with little or no change in the bid slope. The question now arises of whether or not there is any increased predictability, to be found immediately after the event. Figures(23),(25) and (24) show the cross variogram results for 9/10/97 between the micro variables and future returns. We can see immediately that there is much stronger rejection of independence between these variables on this day, then there was over the whole week. In particular there is now clear evidence of dependence between the liquidity variable and future returns. Apart from the short term dependence that last from between 1 and 2 minutes, there also appears to be some evidence of more long term dependence (the spread in particular) with longer lag estimates of the cross variogram falling outside the confidence bands.

So, it appears that on the 9th there is somewhat more dependence between the micro variables and future returns. Finally we ask a simple question of the data. Is this dependence observed throughout the day, or as we suspect, after the Bundesbank announcement takes place and the market adjusts to the new equilibrium exchange rate? We split the day into a pre and a post announcement segment, and reestimated the cross variograms. Figures 26 to 31 confirm our intuition. We can't reject the null of independence pre-announcement, but clearly can do so post-announcement.

8 Conclusion

On the basis of the results presented above we find it very difficult to believe that there is no predictability in the Dollar/Dmark transactions data in DM2000-2. We feel the fact that we find this result is in part due to the choice of statistical technique and this experience leads us to suggest that the variogram may provide a better tool to explore the issue of finding structure in irregularly spaced data rather than standard autocorrelogram analysis. The next stage of this research may be to

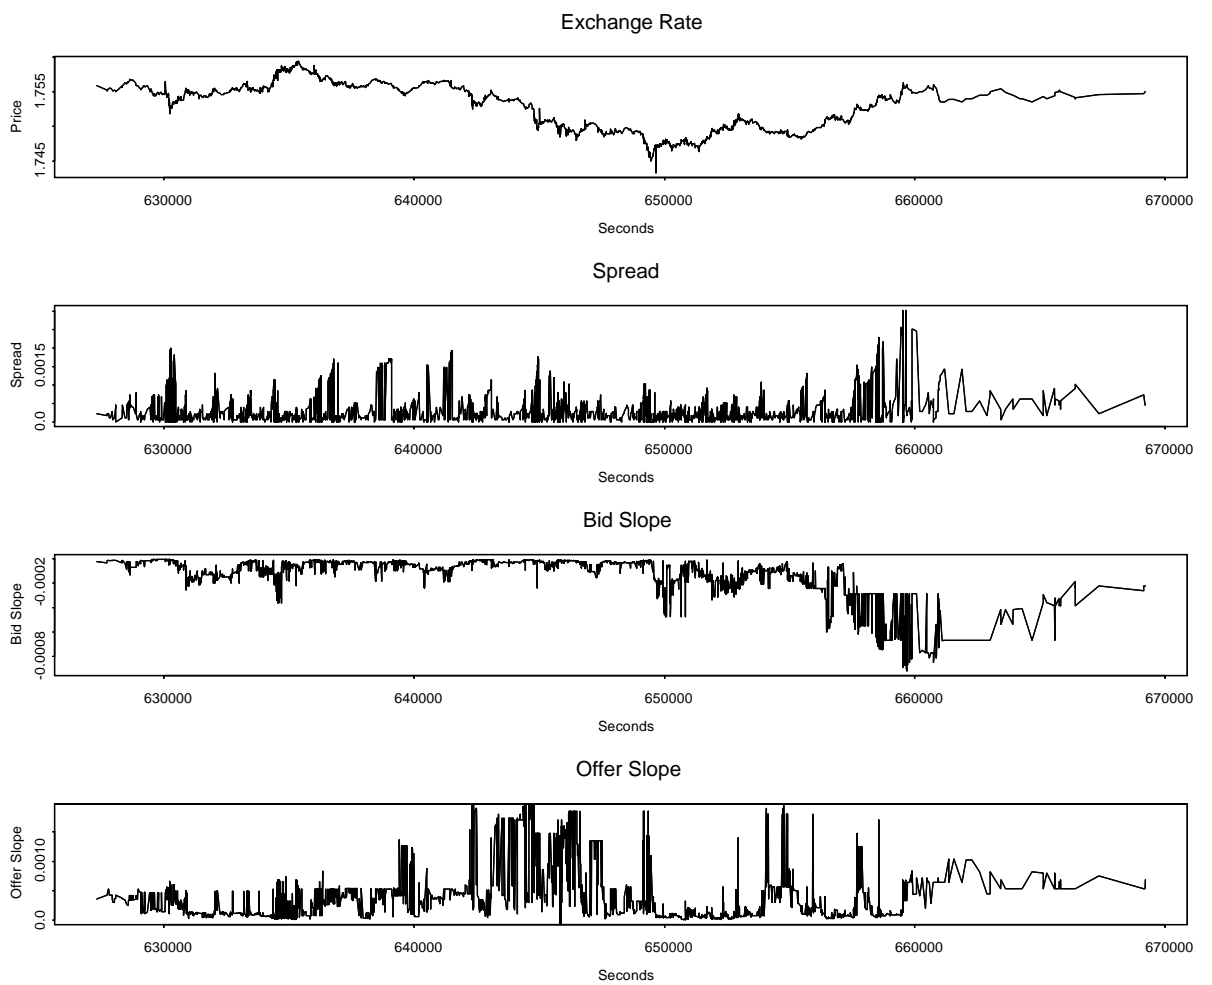


Figure 21: The Micro Variables on the 7th October 1997

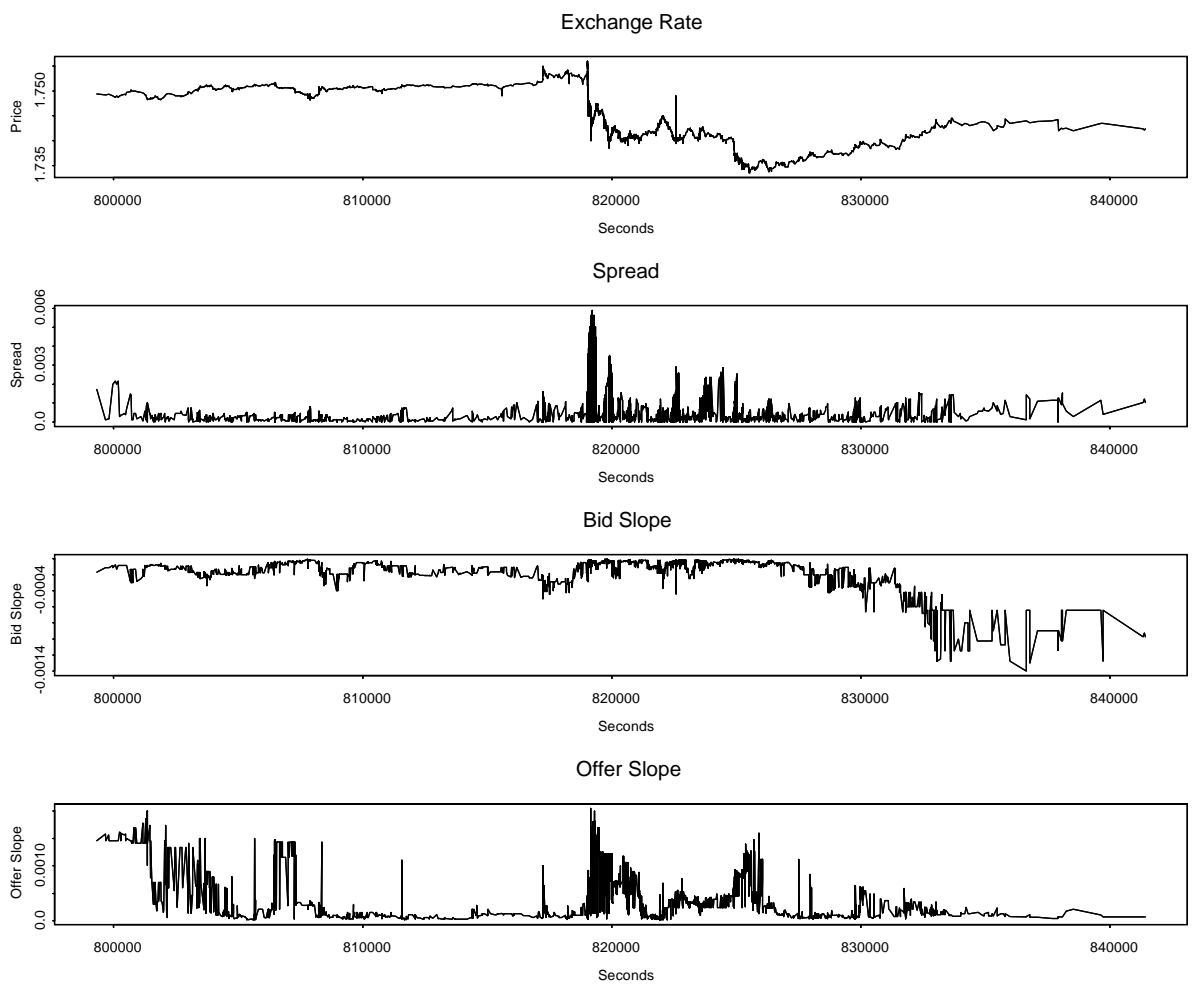


Figure 22: The Micro Variables on the 9th October 1997

Liquidity(t) to Returns(t+h) 9th October 1997




Figure 23:

Spread(t) to Returns(t+h) 9th October 1997




Figure 24:

Volume(t) to Returns(t+h) 9th October 1997




Figure 25:

Volume(t) to Returns(t+h): Pre-Announcement




Figure 26:




Figure 27:

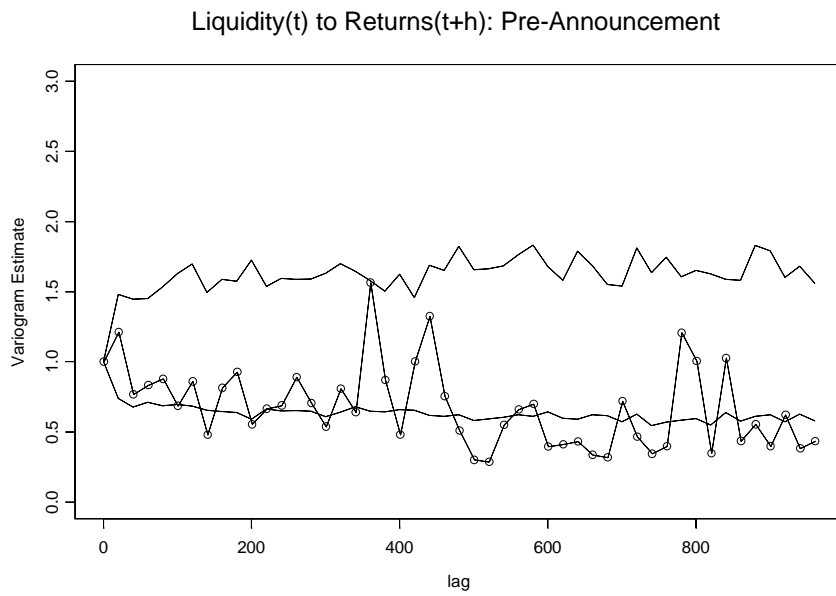


Figure 28:

Liquidity(t) to Returns(t+h): Post-Announcement




Figure 29:

Spread(t) to Returns(t+h): Pre-Announcement




Figure 30:

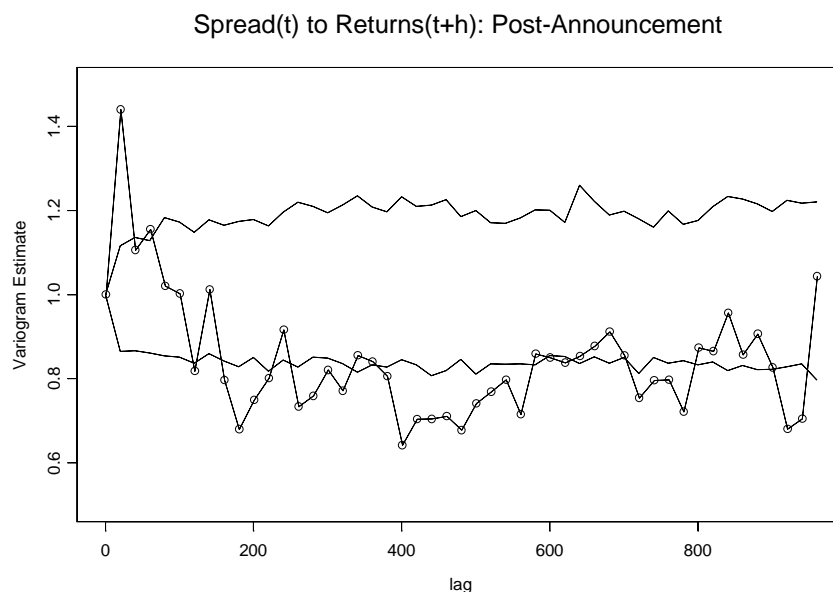


Figure 31:

develop kriging procedures which develop MMSE forecasts point process directly from the second order properties of the process as captured in a fitted variogram model (see appendix 3). This geostatistical approach differs substantially from standard time series method which attempts to match linear time series models with implied acfs to the observed acf and then forecast from these time series models. We also intend to explore the use of the cross variogram in macro event study analysis more deeply using more finely defined micro market based data and to explore the potential of cokriging as a prediction procedure in high frequency financial data. Given what we have learnt from the current research we are also moving to a better position from which to develop a structural micro market model probably as a non-homogeneous Cox process that accounts for the temporal dependence and both the macro and micro influences described found above. At present the model that suggests itself to us is one in which the market periodically, on the advent of macro news, sets a new reference level or equilibrium view of the exchange rate and in between these readjustments normal micro-trading incentives apply. From this perspective it may not in fact be too difficult to find an intellectually coherent transition from macro to micro theories of exchange rate determination. We find it difficult to infer any conclusion regarding the *efficiency* of the DM2000-2 market

from this work at this stage.

Index	Entry Time	Exit time	Type	Result	Price	Q. rd	Q. t
807273	00:18:06.450	00:19:02.280	EnterOffer	EntryRemoved	1.7545	1	1
NA	00:19:02.240	00:19:02.280	EnterTake	EntryTaken	1.7545	2	1
807606	00:54:50.920	00:54:53.180	EnterBid	EntryRemoved	1.7541	1	1
NA	00:54:53.080	00:54:53.180	EnterHit	EntryHit	1.7541	1	1

Table 4: Some Typical Entries in the Data Set

9 Appendix 1:Data Preparation

Obtaining the transaction series used in the current paper takes some effort. The original D2000-2 data set is a list of entries into the electronic trading system. The entries are in consecutive order, beginning at the end of the 5th October 1997 and ending around mid-afternoon (GMT) on the 10th October 1997. There are two types of entry, either a market order (to buy or sell at the current best price) or a limit order (in which the order is logged in the system to be transacted at some later date or withdrawn) . In Table 4 we give examples of four typical entries. rule.

Understanding the entries in this table will explain much of the structure of the data. The first and third entries are limit orders. In the first entry someone has made an offer (an ‘ask’) to sell 1 million dollars at 1.7545DM/Dollar. The entry was entered into the system at 18 minutes past midnight. Within this entry we can also see when the order exited the system and what happened to it. The entry was ‘Removed’ at 19 minutes past midnight, and 1 million was traded. As the quantity requested was 1 million the order was completely filled.

Inspection of the next entry shows that a market order to buy (which does not have a system index number, hence the “NA”) was entered 40 hundredths of a second before the limit order was removed. In fact as we can see that both the market order and the limit order leave the system at precisely the same time, it’s fairly clear that it was in fact this market order that filled the preceding limit order. The market order was actually for 2 million, but we can see that only 1 million was traded, the remaining 1 million does not enter the system.

The next two entries are a similar pair, except this time the limit order is to buy (bid), and the corresponding market order is to sell (a hit). So, in this example we can identify two transactions very easily. The entry of the market order causes the limit order to be removed instantly, and at exactly the same time as the market order itself. This is a simple way to identify transactions, simply matching pairs. There are other type of similar trades that can be similarly identified. For example when a market order pairs with 2 limit orders, causing 3 en-

tries to leave the system at the same time. However, only about 15,000 transactions can be identified this way. More difficult is to identify the automatic ‘crossing’ of limit orders. This happens when the bid and ask (offer) curves cross, and the system automatically trades. Some of these seem easy to recognise, for example when an EnterBid for 1 million is removed at exactly the same time as an EnterOffer for the same quantity is removed. Other limit orders are not so easy to identify however, and to extract these transactions we need to emulate the D2000-2 system. This means maintaining the whole order book over time. In principle this is not difficult, and we have written code in S-Plus that mimics the D2000-2 entry processing. This way the program identifies when limit crosses occur, and given we also have the corresponding removal entries in the data set, we can fairly easily process the data set as if it were being entered into the D2000-2 system itself. There is one particular difficulty mentioned by Danielson and Payne(1999) which we have had to tackle. This is when the emulation program signals a limit cross should occur, but it doesn’t in the data set. This is due to the fact that the two parties on either side of the transaction may not have agreed credit between themselves. The difficulty from our point of view is that there is nothing to signal this ought to happen unless the order book becomes unbalanced and successive trades fail. In some cases this is not too hard to identify. For example suppose there is 2 million at 1.745 on the bid side. Next someone enters a offer limit order for 2 million at 1.745. This ought to cross and thus remove all quantities available at 1.745. Suppose however the next entry in the data file is an EntryHit for 2 million at 1.745, and this is recorded as succeeding immediately. We know then that the limit offer must have failed, else the market hit would not have succeeded. In this way we can simply recognise the inconsistency when it occurs, restore the emulated order book to it’s state before the previous limit order, block the next limit order, and restart the processing. In fact we find that after all the matching pairs have been taken from the data this method allows us to process the rest of the data robustly. Then we can combine the transactions series and form our final transactions series. Performing the processing on the whole data set (including all the matching pairs) is more difficult as there are far more limit orders in the system, and so proportionately more failed limit crosses that need to be corrected. However the general principle of recognising ex-post that a limit order must have been crossed wrongly, then backing up the system and blocking past limit orders seems to work. The advantage from processing the whole data set is that we can build up the entire order book and thus retrieve micro-structural information like the spread, the offer and bid curve shapes and so on.

10 Appendix 2: Measuring Predictability with different information sets and models

Granger and Newbold (1986) emphasise that our understanding of *deterministic* and hence *predictability* is *conditional* on the information set used given that a random variable may be perfectly predictable given one information set but not with another. In particular a series $\{X_n\}$ is stb *Deterministic* if it can be forecast without error or with zero cost given some information set, I_n , eg.

$$\lim_{N \rightarrow \infty} E[C(X_{n+h} - f_{n,h}) | \{I_n : (x_{n-j})_{j=0}^N\}] = 0$$

If the limiting expected cost resulting from a forecast is less than that from using a purely random or white noise process as the forecast then we could say that the series contains a degree of *predictability* P , (conditional on the information set).

$$P = 1 - \frac{\lim_{N \rightarrow \infty} E[C(X_{n+h} - f_{n,h}) | \{I_n : (x_{n-j})_{j=0}^N\}]}{E[C(X_{n+h} - \varepsilon)]}$$

alternatively a no-change comparison in which $f_{n,h} = x_n$ or a “certainty equivalent” measure which might imply $f_{n,h} = 0$ could be used. In the latter case the measure is similar to that proposed by Granger and Newbold (1986, page 310) in which they use the forecast error variance to represent the cost. cf. Messe and Rogoff (1971)

Measures of predictability such as those suggested by **Diebold and Kilian** (1997) which rest on comparisons of forecast performance over different horizons with a fixed information set

$$1 - \frac{E[C(X_{n+h} - f_{n,h}) | \{I_n : (x_{n-j})_{j=0}^N\}]}{E[C(X_{n+k} - f_{n,k}) | \{I_n : (x_{n-j})_{j=0}^N\}]} \quad k < h$$

may be inapplicable as the relevant information set, economic theory and therefore model *change* as the forecast horizon changes.

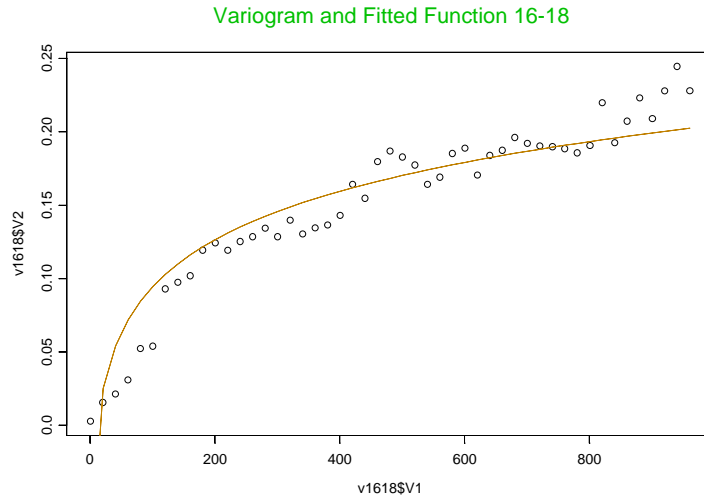


Figure 32:

11 Appendix 3:Kriging

In Figure (32) we show a variogram and a power function fitted to this data. A number of different options as to the choice of model exist and some systematic model selection procedures need to be employed. However with an estimated variogram function we may move directly to develop MMSE predictions. It is also possible to use the estimated variogram to determine the presence or not of long memory, Deibold (1989) and Cressie (1991).

The following algebra is for the simplest case of a constant mean.
Given a model

$$Z(s) = \mu(s) + \varepsilon(s)$$

where $\varepsilon(s)$ is an intrinsically stationary process we seek to predict the value of $Z(s_0)$ at some point, s_0 , in the potentially continuous real line given the values of observations on a sample $\{Z(s_1), \dots, Z(s_n)\}$. The ordinary kriging predictor given by Matheron(1971), is linear, uniformly unbiased and minimises the mean square prediction error (BLUP) and is given by

$$\hat{Z}(s_0) = \left(\gamma + \mathbf{1} \frac{(1 - \mathbf{1}'\Gamma^{-1}\gamma)}{\mathbf{1}'\Gamma^{-1}\mathbf{1}} \right)' \Gamma^{-1} \mathbf{Z}$$

where $\mathbf{Z} \equiv (Z(s_1), \dots, Z(s_n))'$

$\gamma \equiv (\gamma(s_0 - s_1), \dots, \gamma(s_0 - s_n))'$

and Γ is an $n \times n$ matrix whose (i, j) 'th element is given by $\gamma(s_i - s_j)$

The minimum mean-square prediction error can be simple derived as

$$\sigma^2(s_0) = \gamma' \Gamma^{-1} \gamma - (\mathbf{1}' \Gamma^{-1} \gamma - 1)^2 / (\mathbf{1}' \Gamma^{-1} \mathbf{1})$$

and then we form the usual prediction (for example 95%) intervals

$$A \equiv (\hat{Z}(s_0) - 1.96\sigma^2(s_0), \hat{Z}(s_0) + 1.96\sigma^2(s_0))$$

This formula which gives us the optimal predictor for irregularly spaced data for intrinsically stationary processes can be seen to depend simply on the variogram of the process. We have no need to pass through the intermediate step of fitting say ARIMA models to generate forecasts given the estimated variogram. Second order prediction theory can obviously be generated directly from the autocorrelogram when it exists but the real advantage in our case is that the class of models we need to consider when modelling irregular data generated from some point process are much less well established than the ARIMA class for linear time series with covariance stationarity and regular observations.

In terms of time series notation we can apply the above formulas substituting t_1, t_2, \dots, t_n for s_0, s_1, \dots, s_n , and the prediction point s_0 is now t_{n+h} where h is the forecast horizon.

In figure 33 we give an example of the kriging forecast where the simulated process is a simple AR(1) model with a coefficient of 0.8. We also plot the 95% confidence intervals. We use the data up to point 100, to estimate the variogram, and then go direct from the variogram to forecasting the next 50-steps out.

Plot 34 gives an example of out-of-sample forecasts on a sub-sample of the D2000-2 series. We can see that some of the short-horizon forecasts are somewhat erratic, and as the horizon increases we see a smoother adjustment of the forecast. We believe the erratic early forecasts are due to the method of fitting the variogram. In particular the estimates of the variogram at short lags are dependent on the bias or offset of the variogram (i.e. where it meets the origin). In the kriging literature there is a considerable attention paid to this issue, which is generally called the 'nugget effect', because of the spatial mining context. We are currently pursuing the usefulness of kriging forecasts further.

References

- [1] Andersen T. and Bollerslev T.,(1997), *Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns*, **Journal of Finance**, Vol 52,No.3,pp 975-1005

Kriging Forecasts with an AR(0.8) Simulated Process




Figure 33: Out-of-Sample h-step Forecasts from the Kriging Model, and 95% confidence interval.

Kriging Forecast For a Sub-Sample of the D2000-2 Data




Figure 34: Kriging Forecasts For a Sub-Sample of the D2000-2 Series

- [2] Andersen T. and Bollerslev T.,(1998), *DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements and Longer Run Dependencies*, **Journal of Finance**, Vol 53, No.1,pp 219-265
- [3] Bauwens L. and Giot P.,(1998),*Asymmetric ACD Models: Introducing Price Information in ACD Models with a Two State Transition model*, **CORE Discussion Paper 98XX, June 1998.**
- [4] Beran Jan,(1994), **Statistics for Long Memory Processes**, Chapman Hall, Monographs on Statistics and Applied Probability.
- [5] Brillinger D.A., (1972), *The spectral analysis of stationary interval functions*, **Sixth Berkeley Symposium on Probability and Mathematical Statistics,I.** University of California Press, Berkeley.483-513.
- [6] Cheung Y-W, Chin M.D. and Marsh I.,(1999), *How do foreign exchange dealers think their market operates?*, **Mimeo CUBS.**
- [7] Clark P.K.,(1975), *The use of operational time to correct for sampling interval misspecification*, **Review of Economics and Statistics**, 225-230
- [8] Cochrane J.H.,(1988), *How big is the random walk in GNP?*, **Journal of Political Economy** ,96 ,no.5,893-920.
- [9] Cressie N.,(1986), *Kriging non-stationary data*, **Journal of the American Statistical Association**, 81, 625-634.
- [10] Cressie N., (1988), *A graphical procedure for determining non-stationarity in time series*,**Journal of the American Statistical Association**,83,1108-1116. [Correction JASA(1990), 85,272]
- [11] Cressie N., (1991) , **Statistics for Spatial Data**, John Wiley.
- [12] Cressie N. and Grondona M.O.,(1992), *A comparison of variogram estimation with covariogram estimation*, in **The Art of Statistical Science**, K.V.Mardia, Wiley.
- [13] Dacco R. and S.Satchell, (1998)*High Frequency Switching Regimes: A Continuous-time Threshold Process*, in **Nonlinear Modelling of High Frequency Financial Time Series**, edited by Dunis and Zhou, Wiley.
- [14] Danielsson J. and Payne R.,(1999), *Real Trading Patterns and Prices in Spot Foreign Exchange Markets*, **mimeo, FMG, LSE.**
- [15] Darolles S., Gouriéroux C. and Le Fol G., *Intraday Transaction Price Dynamics*, **Discussion Paper CREST.**
- [16] DeJong F. and T. Nijman, (1995), *High frequency analysis of lead-lag relationships between financial markets*, **Journal of Empirical Finance**,4,259-277.
- [17] Diebold F.,(1989), *Random walks versus fractional integration: power comparisons of scalar and joint tests of the variance time function*,in **Advances in Econometrics and Modelling**, ed

- .B.Raj, Kluwer.
- [18] Diebold F. and L. Kilian (1997), *Measuring Predictability: Theory and Macroeconomic Applications*, mimeo, University of Pennsylvania.
 - [19] Engle R. and Russell J.R., (1998), *The Autoregressive Conditional Duration Model*, **Econometrica**, 66, 1127-1163.
 - [20] Fuller W.A., (1976), **An Introduction to Statistical Time Series**, John Wiley, New York.
 - [21] Gel'fand I.M. and Vilenkin N.Y., (1964), **Generalised Functions: Applications to Harmonic Analysis, vol(4)**, New York, Academic Press.
 - [22] Ghysels E., Gouriéroux C. and Jasiak J., (1996), *Kernal Autocorrelation for Time Deformed Processes*, **CIRANO working paper** 96-19.
 - [23] Guarda P. and Salmon M., (1997), *On the Detection of Nonlinearity in Foreign Exchange Data*, in **Nonlinear Dynamics in Economics**, Cambridge University Press, edited by William Barnett, Alan Kirman and Mark Salmon
 - [24] Granger C. and Newbold P., (1986) *Forecasting Economic Time Series*, 2nd, Edition, Academic Press.
 - [25] Haslett J., (1997), *On the sample variogram and the sample autocovariance for non-stationary time series*, **The Statistician** 46, 4, 475-485.
 - [26] Hillman R., (1998), *Econometric Modelling of Nonlinearity and Non-Stationarity in the Foreign Exchange Market*, **PhD Thesis**, University of Southampton, 1998.
 - [27] Jasiak J., (1999), *Persistence in Intertrade Durations*, **working paper, York University**, Canada.
 - [28] Kolmogorov A.N., (1941), *The local structure of turbulence in an incompressible fluid at very large Reynolds numbers*, **Doklady Akademii Nauk SSSR**, 30, 229-303.
 - [29] Lo A. and Mac Kinley, (1988), *Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test*, **The Review of Financial Studies**, Vol 1, No.1, pp 41-66.
 - [30] Lundin M., Dacorogna M., and Muller U.A., (1998), *Correlation of high frequency financial time series*, in **The Financial markets Tick by Tick**, P. Lequeux ed., John Wiley and Sons, London.
 - [31] Matheron G., (1973), *The intrinsic random functions and their applications*, **Advances in Applied Probability**, 5, 439-468.
 - [32] McDonald R., (1998) *What do we really know about real exchange rates*, **Oesterreichische Nationalbank, Working Paper** no.28.
 - [33] Mussa M., (1990), *Exchange rates in theory and reality*, **Princeton**

Essays in International Finance, No.179, Dec 1990.

- [34] Newbold P. and Agiakloglou C.,(1993), *Bias in the sample autocorrelations of fractional noise*, **Biometrika**, 80,3,698-702.
- [35] Payne R., (1999),*Informed Trade in Spot Foreign Exchange Markets: an Empirical Investigation*, **Mimeo, FMG, LSE**.
- [36] Poterba J.M. and Summers L.H., (1988), *Mean reversion in stock prices*,**Journal of Financial Economics**, 22,27-59.
- [37] Quenouille M.H.,(1958), Discrete autoregressive schemes with varying time intervals, **Metrika**, 1,21-27.
- [38] Richardson M. and Stock J.,(1989), *Drawing inferences from statistics based on multiyear asset returns*, **Journal of Financial Economics**, 25, 323-348.
- [39] Von Neumann J., (1941), *Distribution of the ratio of the mean square successive difference to the variance*, **Annals of Mathematical Statistics**, 12,367-395.
- [40] Whittle P.,(1954), *On stationary processes in the plane*, **Biometrika**, 41,434-449.
- [41] Whittle P.,(1983), **Prediction and Regulation by Linear Least Squares Methods**, second edition, Blackwell.
- [42] Yaglom A.M.,(1955),*Correlation theory of processes with random stationary n 'th increments*. **Matematicheskii Sbornik**, 37,141-196, translation in **American Mathematical Society Translations, series 2**, vol 8, (1958), American Mathematical Society , Providence, RI, 87-141.