*RISK Magazine **7**:6 (1994) 71-76*

Creating Order Out of Chaos

Carol Alexander and Ian Giblin, University of Sussex, UK.

One of the characteristics of a chaotic dynamical system is that points on neighbouring trajectories grow apart from each other exponentially, at least in the short term. Think of these points as forecasts from mathematical models based on chaotic dynamics, and you have a fair insight about what may be happening in the new, high-tech. research into chaos in capital markets.

Opinions are mixed about the ability of mathematical models to identify a chaotic system underlying a financial time series, and/or to exploit consequent investment opportunities. Articles in the press hint at both success (*Investors Chronicle, 10 July 1992, Economist, 9 October 1993)* and failure (*Financial Times, 19 April 1993)* . Amongst books and academic articles on the subject there are those that claim to have found chaos in financial markets (Peters 1991, De Grauwe et. al. 1993), those that claim not (Hseih, 1991, Tata and Vassilicos, 1991) and those that are undecided (Scheinkman and LeBaron, 1989, Liu et. al., 1992). For the most part this research is based on derivatives of the *correlation exponent*[1] algorithm devised by Grassberger and Procaccia (1993a,b), for which there are a number of key parameters:

---

[1] The correlation dimension (or exponent) ν of a time series $\{x_t\}$ of length N is a measure of the rate of growth of the correlation integral $C(\varepsilon)$, being the number of points in a ball of radius ε, as N increases and $\varepsilon \to 0$. Thus $C(\varepsilon) \approx \varepsilon^{\nu}$ for small ε. ν is calculated by first finding a finite approximation to each of the correlation integrals

$$C(\varepsilon) = \lim \{\text{number of pairs } (i,j) \text{ such that } | x_i - x_j | < \varepsilon\} / N^2 \quad \text{as } N \to \infty$$

and then estimating $d\ln C(\varepsilon)/d\ln\varepsilon$. $C(\varepsilon)$ is approximated by embedding the time series in $\Re^m$ where m is chosen so that $m > \nu$.

- the *sample size* may need to be huge - upwards of 100,000 points - if the algorithms are to work with any degree of accuracy. Erroneous conclusions might be drawn even with quite long data sets (see Vassilicos et. al. 1992);
- the data must not contain much *noise* - these methods cannot detect underlying chaos when mixed with even a small amount of noise (see Liu et. al. 1992);
- correlation exponent algorithms depend on counting the number of points inside a small ball in m-dimensional space, and results can be quite sensitive to the choice of *diameter* for this ball (also in Liu et. al. 1992);
- the time series must be *stationary*[2] if results are to be correctly interpreted (Casdagli, 1992).

Thus very long time series are necessary for these correlation exponent methods, but on the other hand unfiltered tick-by-tick data, or even daily data, would not be suitable because of the high level of noise[3]. Since no weekly series of sufficient length are available, filtering the noise out of data is the only alternative. Also, in many cases a suitable stationarity transform needs to be employed. Not all the literature mentioned above is based on empirical methods which meet these stringent criteria - and those that are shouting 'chaos' the loudest may be the most lacking.

A new generation of algorithms

Martin Casdagli has pioneered work on the application of *radial basis functions*, which have many features in common with neural networks but take much less time to

2 A (weakly) stationary stochastic process $\{x_t\}$ is one in which the unconditional mean and variance are (finite) constants and for which the covariance of x_t and x_{t-l} is independent of time and depends only on the lag l.

3 In high frequency data noise may be accounted for by reporting errors, or other errors of measurement such as would be introduced by non-synchronous trades.

run (Casdagli, 1989). This work has given rise to a new and more efficient type of algorithm which can effectively distinguish chaos from random behaviour in time-series of 10^d data points, where d is the dimension of the attractor. Hence low-dimensional chaotic systems may be detected with only a modest amount of data, say 1000 data points. These new algorithms can also provide accurate short-term forecasts if the time series is found to be chaotic, and are therefore loosely termed *time delay prediction methods.*

The basic time delay prediction method approach to identifying chaotic series is based on a theorem of Takens (1981), which shows that it is possible to estimate the dimension of any chaotic attractor by embedding the time series in m-dimensional space as follows: map each point x_t in the time series to a point

$$\mathbf{x}_t = (x_t, x_{t-\tau}, x_{t-2\tau}, \ldots., x_{t-(m-1)\tau})$$

where the lag τ is some positive integer[4] and m is chosen to be larger than the dimension of the attractor. Having embedded the series in m-dimensional space $\Re^m$, a library of patterns in $\Lambda \in \Re^m$ is generated from the first half of the data set. Then this library is used to forecast points in the second half of the data set, r steps ahead for r=1,2,3,...

The method of forecasting varies: in Casdagli (1992) the forecasts are based on parametric linear regression; Nychka et. al. (1992) use nonparametric regression to find consistent estimates of the Liapunov exponent[5]; Sugihara and May (1990) suggest

---

[4] Sugihara and May (1990) find that results are largely independent of the choice of lag, and for the most part stick with τ=1. De Grauwe et. al. (1993) suggest that τ should be chosen a small as possible, but large enough to minimize autocorrelation between the components of **x**. This also sugests that, for our purposes, a lag of one is sufficient.

[5] The Liapunov exponent measures the average rate of exponential divergence of neigbouring trajectories along the line of greatest divergence, so it is a measure of sensitivity to initial conditions.

a non-parametric simplex projection method; Alexander and Giblin (1993) modify the Sugihara and May algorithm to use barycentric coordinates[6]. In all of these methods, if the system is chaotic the correlation between actual and forecasted values will decline as r increases, and very short term forecasts will be quite accurate. On the other hand, if the system is purely random no such decrease in prediction accuracy will be evident, and for many financial returns series even the one step ahead predictions will be uncorrelated with the actual returns.

Consider first the r-step ahead forecast of a point x_t in the second half of the series, for a fixed positive integer r. A simplex projection algorithm locates the smallest simplex σ in Λ which contain $\mathbf{x}_t$. The vertices of σ are then projected r steps ahead to get a new simplex σ^r, and a point $\mathbf{x}_{t+r}$ in σ^r which is 'equivalent' to the point $\mathbf{x}_t$ in σ is taken as the forecast of $\mathbf{x}_t$ in $\Re^m$. Now the r-step ahead forecast of x_t is the first coordinate of $\mathbf{x}_{t+r}$, x_{t+r}.

Assuming independence of initial conditions (see Wolff, 1992) the largest Liapunov exponent is given by

$$\lambda = \lim (\ln|df^r/dx|/r) \text{ as } r\to\infty,$$

where the trajectory of x is $\{x, f(x), f^2(x),.....\}$. Hence $|df^r/dx| \approx \exp(r\lambda)$ and, by Taylor linearization, this is approximately equal to the difference between x and $x+\varepsilon$ after m iterations, i.e.

$$f^r(x+\varepsilon) - f^r(x) \approx \exp(r\lambda).\varepsilon$$

A positive exponent implies that trajectories tend to separate, with a negative exponent trajectories converge.

6 The barycentric coordinates of a point $\mathbf{x}_t$ with respect to the simplex σ in $\Re^m$ are m+1 distances from the sides of the simplex, $b_1, b_2, ... b_{m+1}$, normalized so that they sum to unity. If each b_i lies between 0 and 1 then $\mathbf{x}_t$ lies within the simplex.

In chaotic systems both prediction error and standard error of prediction grow exponentially fast, since the simplex σ^r has a larger volume than σ due to separation of initially nearby trajectories. This will not be the case if the time series is random noise. Figure 1 represents our forecasting method graphically, for the case m=3.

Now that we have r-step ahead forecasts of each point in the second half of the time series, we calculate the correlation coefficient between the forecast x_{t+r} and the actual series x_{t+r}, and call this number $\rho(r)$. Then a simple plot of $\rho(r)$ against r should reveal a downward trend if the time series is deterministic chaos, and an approximately horizontal line if the time series is purely stochastic.

The barycentric coordinate method

The power of this simplex projection method to distinguish chaos, even in noisy data sets of less than 1000 data points, can be illustrated with a chaotic deterministic map into which we can mix a controlled amount of noise. We choose the tent map on 0.5, also called *white chaos* because its autocorrelation function has similar properties to a white noise stochastic process[7]. The dynamics are given by

$$x_t = \begin{cases} 2x_{t-1} & \text{if } 0 \le x_{t-1} < 0.5 \\ 2(1-x_{t-1}) & \text{if } 0.5 \le x_{t-1} \le 1. \end{cases}$$

---

[7] Sakai and Tokumaru (1980) show that the tent map has the same correlogram as an AR(1) process, and for a near to 0.5 these autocorrelations are close to those of a white noise process. For this reason the tent map with a=0.5 is called 'white chaos'. Rounding errors in binary arithmetic imply that the tent map on exactly 0.5 is difficult to generate on a computor because it rapidly converges to a boundary. The value of 0.4999 gives the required dynamics for this example.

and we can generate time series of a suitable length - say 1000 data points - from this map to demonstrate the algorithm. The correlation dimension of a tent map is known to be one, so we assume that the optimal embedding dimension is m=2. Figure 2(a) shows how the r-step ahead correlations (between actual and predicted values in the last 500 data points) decrease monotonically with respect to r, which indicates a chaotic system. Figure 2(b) repeats these calculations on the first differences of the tent map series, and it is clear that this does not alter the conclusion[8].

In a chaotic system such as this the algorithm also provides more accurate short term forecasts than standard parametric regression: the r-step ahead correlations obtained by using a first order autoregressive model based on the first 500 data points of the white chaos map is shown in figure 2(c). Since the autocorrelations of the series are similar to those of white noise, it is not surprising that r-step ahead predictions have very low correlations with the actual series. In fact autoregressive model predictions will yield correlations that are uniformly below those obtained with our prediction method, whatever the parameter in the tent map.

The 'white chaos' tent map that we've just analysed may be indistinguishable from white noise in many ways (hence its name), but prediction methods *can* distinguish the two. Figure 2(d) illustrates the correlation plot (based on m=3, τ=1) obtained from our barycentric coordinate algorithm for a Gaussian white noise process. The correlation between actual and predicted values in the last half of the data is approximately zero

---

[8] Correlation plots of other chaotic maps, such as the logistic map, or sums of several independent chaotic maps, also reveal a negative trending pattern which is characteristic of chaos. More details are given in Alexander and Giblin, 1993.

for all values of r, so the algorithm has not identified any chaotic system underlying this time series.

One of the problems with correlation exponent methods is that they fail to detect chaos when mixed with even a small amount of noise (Liu et. al., 1992). Our prediction method *can* identify chaos in noisy systems - but only if the signal to noise ratio is sufficiently large. Figure 3 illustrates correlation plots for a sequence of time series of the form

$$z_t = \alpha x_t + (1-\alpha)y_t$$

where x_t is a white chaos map and y_t is Gaussian white noise, for values of α between 1 and 0.7. In time series with more than 10% noise, no evidence of chaos is found, but the underlying chaos is clear when a small amount of noise is introduced.

Evidence of chaos.....?

The empirical literature on detecting chaos in financial data is rather ambiguous, and clearly more research needs to be done. Amongst the authors who claim to have found evidence of chaos are: Edgar Peters (1991), who finds positive fractal dimensions in four of the major world stock markets, viz. the US, UK, Germany and Japan[9];

[9] However this analysis may be open to question because the Morgan Stanley Capital Indices for the UK, Germany and Japan have been detrended with a deterministic time trend, but the detrended series may still have a stochastic trend. Indeed it is possible, even probable, that the Morgan Stanley indices are I(1) processes, rather than the I(0) + trend processes that he assumes them to be.

The S&P 500 is 'detrended' by the US consumer price index to remove any inflationary trend. Although economically meaningful, this process need not have any empirical meaning because the S&P 500 and the CPI may be *non-cointegrated* I(1) processes. Given Casdagli's remarks about searching for chaos in non-stationary series, and the different conclusions from first differences and levels of our data, we do not find it surprizing that Peters has apparently identified a fractal dimension of 2.33 for the S&P 500.

Another questionable aspect of Edgar Peters methods is whether the number of data points he uses is sufficient for the correlation dimension approach, and we would agree that less than 500 data points is unlikely to be regarded as adequate by most practitioners.

Scheinkman and LeBaron (1989), who find weak evidence of chaos in weekly S&P 500 returns with a correlation dimension of about six[10]; and finally De Grauwe et. al. (1993) find evidence of chaos with correlation dimension of about 2 in the daily Cable rate between 1973 and 1981, and in the YEN/$ rate between 1973 and 1990 with estimated correlation dimension 2.2[11]. However Liu et. al. (1992) use daily IBM S&P 500 returns from July 2 1962 to December 31 1985 and are unable to draw any firm conclusion - their results are consistent with either stochastic white noise or a true dimension of about six - and Hseih (1991) and Vassilicos, et.al. (1992) report results which suggest *no* underlying low-dimensional chaotic system in some dollar denominated currency and stock returns.

We apply the barycentric coordinate prediction algorithm to daily data on the S&P 500 from January 1 1987 to May 5 1993 (a total of 1655 observations - see figure 4(a))[12]. The data needs to be logged and then first differenced in order that it be stationary before the prediction algorithm is applied[13]. Since the basic shape of our correlation plots remains unchanged when the embedding dimension is increased from 2 up to 6, and when the lag is increased from 1, we illustrate results in figure 4(b) for m=3 and τ

---

10 However their results might suffer from insufficient data. Vassilicos et. al. (1992) have used Charles Goodhart's enormously long series of S&P 500 daily stock returns and of German Mark/ dollar FX rates, finding no evidence of an underlying chaotic system.

11 The cable results are based on only 2250 data points, which may lead to inaccuracies in correlation exponent methods, as we have seen. In none of this work are standard errors quoted for correlation dimension estimates, and so supposedly 'large' estimates of correlation dimension could, in fact, be insignificantly different from zero. Correlation dimension methods indicate no chaos in the DM/$ returns, but all three exchange rates exhibit some evidence of chaos when estimates of the Hurst coefficient are found.

12 Many thanks to Andrew Johnson of S.G. Warburgs for supplying the data on the S&P 500.

13 One of the reasons for taking logarithms is that first differences yield approximate stock returns, so that the series we are analysing has economic meaning.

=1. The correlations do not show any sign of decreasing with r, and so we conclude that there is too much noise in the data for the algorithm to identify an underlying chaotic system - if indeed such a system exists!

Weekly returns on the S&P 500 may be analysed, to see whether the consequent reduction in noise makes any difference. Data in figure 4(c) are from January 1 1969 to February 19 1993 (1260 observations), and we have used the same lag and embedding dimension as for the daily series. Figure 4(d) shows the correlation plot - and there is no evidence of chaos. Thus our findings support those of the LSE financial markets group (Vassilicos et. al. 1992) who find no evidence of chaos in the S&P 500, and contradict those of Peters (1991) who claims to have found a fractal dimension of 2.33 in weekly S&P data, and Scheinkman and Le Baron (1989) who find weak evidence of a chaotic attractor of dimension about 6 in the S&P 500.

Figure 5(a) illustrates London daily closing bid rates of the German mark/$ from January 2 1982 to December 10 1992 (excluding bank holidays, so a total of 2768 observations)[14]. As before, we take first differences of the logarithms (that is, approximate currency returns) to ensure stationarity of the data, and assume the parameters τ=1 and m=3 for the usual correlation plot of $\rho(r)$ against r[15]. Figure 5(b) shows no notable downwards trend, and so the algorithm has not detected any low-dimensional chaotic attractor underlying the data.

---

[14] Thanks to State Street Global Advisors, London, for providing these data, and similar data on Cable.

[15] De Grauwe et. al. (1993) find evidence of chaotic attractors with dimensions between 0.8 and 2.2 for two of the three polar exchange rates.

Figure 5(c) shows a correlation plot (with the same parameters as above) which would be obtained if we were to apply the algorithm to the *levels* of the log DM/$ rate, but it would be erroneous to conclude 'chaos' from this. If a non-linear chaotic system *had* generated the data, the application of a linear filter such as a first difference would not have had much effect on the correlation plot (see Broomhead et. al. (1992)). But earlier we showed that when taking first differences yields a marked change in correlation plots, such as in figures 5(b) and 5(c), the relevant plot is the one which applies to *stationary* data, i.e. figure 5(b)[16].

In some ways (for example, in GARCH modelling) the Cable rate in figure 5(d) behaves very differently from other dollar rates[17] and it is possible that this could be due to differences in some deterministic system which underlies these rates. Certainly de Grauwe et. al. (1993) find stronger evidence of chaos in the Cable rate than in the Yen or German Mark dollar rates. We therefore apply the same prediction method to a daily $/£ rate over the same period. Figures 5(e) and 5(f) are the correlation plots for first differences and levels of the Cable rate. Again there is again no evidence of chaos (note that the levels plot is not valid because data are non-stationary)[18].

---

[16] We have used parameters m=3 and τ=1for all plots. With a longer lag length, the monotonicity of the levels correlation plots will be less evident, so this should be taken into account when interpreting the correlation plot.

[17] For this reason Gallant, Hseih and Tauchen (1991) have termed it a 'recalcitrant' series.

[18] Similar remarks hold for any non-stationary transform of the S&P 500. Although the levels S&P 500 correlation plots are not shown, they also display the distinctive shape of chaotic maps, but nothing should be inferred from this.

Concluding Remarks

Previous empirical research into chaos in financial markets, which has mostly been based on correlation exponent algorithms, has suffered from a number of sample size and/or noise problems - and in some cases there are also methodological defects. The results are conflicting and a consensus opinion seems a long way off! To add to the increasing body of evidence *against* chaos, we have used a barycentric coordinate version of the Sugihara and May simplex projection method for detecting chaos in reasonably short noisy time series. We have found no evidence of chaos in daily or weekly data on the S&P 500, in the German Mark dollar rate, or the Cable rate.

The two most important advantages of prediction methods over correlation exponent tests are that they need far less data, and are able to distinguish chaos in data with a small amount of noise. Weekly series on markets with a relatively low trading volume are the most likely data to reveal any chaotic deterministic system underlying the noise. Such series will not, as present, be long enough for the correlation exponent approach, and we maintain that prediction methods will therefore provide a way forward in this research.

Time delay prediction methods such as the one that we have described have been the subject of some interest in the recent financial press. It is the opinion of the authors that these methods will prove far more useful in the provision of simple non-parametric forecasts than for the detection of chaos in financial markets.

References

Alexander, C.O. and I. Giblin (1993) 'Searching for chaos in financial markets' *University of Sussex Mathematics Research Report, 93-24.*

Broomhead, D.S., J.P. Huke and M.R. Muldoon (1992) 'Linear filters and non-linear systems' *Jour. Royal Statistical Society, Series B.* 54, 373-382.

Casdagli, M. (1989) 'Non-linear prediction of chaotic time series' *Physica D 35 335-356*

Casdagli, M. (1992) 'Chaos and deterministic vs stochastic non-linear modelling' *JRSS series B 54 303-328.*

Gallant, A. R., D. Hseih and G. Tauchen (1991) 'On fitting a recalcitrant series: The pound/dollar exchange rate 1974-83' in *Nonparametric and semiparametric methods in econometrics and statistics*, W. Barnett, J. Powell and G.Tauchen (eds) CUP.

Grassberger, P. and I. Procaccia (1983a) 'Characterisation of strange attractors' *Phys. Rev. Letters. 50 346-349.*

Grassberger, P and I. Procaccia (1983b) 'Measuring the strangeness of strange attractors' *Physica 9D 189-208.*

de Grauwe, P, H. Dewachter and M. Embrechts (1993) *Exchange rate theory: Chaotic models of foreign exchange markets* Blackwell.

Hseih, D.A. (1991) 'Chaos and non-linear dynamics; application to financial markets' *Jour. Finance 46 1839-1877.*

Liu, T. C.W.J. Granger and W.P. Heller (1992) 'Using the correlation exponent to decide whether an economic series is chaotic' *Jour. App. Econometrics 7 25-39.*

Nychka, D., S. Ellner, A.R. Gallant and D. McCaffrey (1992) 'Finding chaos in noisy systems' *Journal of the Royal Statistical Society, Series B* **54** 399-426.

Peters, E.E. (1991) *Chaos and order in the capital markets* John Wiley and Sons.

Ruelle, D. (1990) 'Deterministic chaos: the science and the fiction' *Proceedings of the Royal Society of London, A.* **427**, 241.

Scheinkman, J.A. and B. LeBaron (1989) 'Non-linear dynamics and stock returns' *Jour. Business 62 311-328.*

Sakai, H., and H. Tokumaru (1980) 'Autocorrelations of a certain chaos' *IEEE Transactions on Acoustics, Speech and Signal Processing* **ASSP-28** (5), 588-590.

Smith, R.L. (1992) 'Estimating dimension in noisy chaotic time series' *JRSS series B 54 329-352*

Sugihara, G. and R.M. May (1990) 'Non-linear forecasting as a way of distinguishing chaos from measurement error in a data series' *Nature 344 734-741*.

Takens, F. (1981) 'Detecting strange attractors in fluid turbulence' *Dynamical systems and Turbulence* (eds. D. Rand and L-S Young) Spinger.

Tata, F. and C. Vassilicos (1991) 'Is there chaos in economic time series? A study of the stock and foreign exchange markets' *LSE financial markets group DP 120*

Vassilicos, J.C., A.Demos and F.Tata (1992) 'No evidence of chaos but some evidence of multifractals in the foreign exchange and the stock markets' *LSE financial markets group, DP 143*.

Wolff, R.C.L. (1992) 'Local Liapunov exponents; looking closely at chaos' *JRSS series B 54 353-372*

