2011-01-01

Risk Management Trends: Currency Trading using the Fractal Market Hypothesis

Jonathan Blackledge
Technological University Dublin, jonathan.blackledge@tudublin.ie

Kieran Murphy
Technological University Dublin

Currency Trading using the Fractal Market Hypothesis

Jonathan Blackledge and Kieren Murphy
Dublin Institute of Technology
Ireland

1. Introduction

We report on a research and development programme in financial modelling and economic security undertaken in the Information and Communications Security Research Group (IC-SRG, 2011) which has led to the launch of a new company - Currency Traders Ireland Limited - funded by Enterprise Ireland. Currency Traders Ireland Limited (CTI, 2011) has a fifty year exclusive license to develop a new set of indicators for analysing currency exchange rates (Forex trading). We consider the background to the approach taken and present examples of the results obtained to date. In this 'Introduction', we provide a background to and brief overview of conventional economic models and the problems associated with them.

1.1 Background to Financial Time Series Modelling

The application of mathematical, statistical and computational techniques for analysing financial time series is a well established practice. Computational finance is used every day to help traders understand the dynamic performance of the markets and to have some degree of confidence on the likely future behaviour of the markets. This includes the application of stochastic modelling methods and the use of certain partial differential equations for describing financial systems (e.g. the Black-Scholes equation for financial derivatives). Attempts to develop stochastic models for financial time series, which are essentially digital signals composed of 'tick data'[1] can be traced back to the early Twentieth Century when Louis Bachelier proposed that fluctuations in the prices of stocks and shares (which appeared to be yesterday's price plus some random change) could be viewed in terms of random walks in which price changes were entirely independent of each other. Thus, one of the simplest models for price variation is based on the sum of independent random numbers. This is the basis for Brownian motion (i.e. the random walk motion first observed by the Scottish Botanist Robert Brown) in which the random numbers are considered to conform to a normal of Gaussian distribution. For some financial signal $u(t)$ say (where u is the amplitude - the 'price' - of the signal and t is time), the magnitude of a change in price du tends to depend on the price u itself. We therefore modify the Brownian random walk model to include this observation. In this case, the logarithm of the price change (which is also assumed to conform to a normal distribution) is given by

$$\frac{du}{u} = \alpha dv + \beta dt \quad \text{or} \quad \frac{d}{dt}\ln u = \beta + \alpha \frac{dv}{dt}$$

---

[1] Data that provides traders with daily tick-by-tick data - time and sales - of trade price, trade time, and volume traded, for example, at different sampling rates.

where α is the volatility, dv is a sample from a normal distribution and β is a drift term which reflects the average rate of growth of an asset[2]. Here, the relative price change of an asset is equal to a random value plus an underlying trend component. This is the basis for a 'log-normal random walk' model (Copeland et al., 2003), (Martin et al., 1997), (Menton, 1992) and (Watsham and Parramore, 1996).

Brownian motion models have the following basic properties:

- statistical stationarity of price increments in which samples of Brownian motion taken over equal time increments can be superimposed onto each other in a statistical sense;

- scaling of price where samples of Brownian motion corresponding to different time increments can be suitably re-scaled such that they too, can be superimposed onto each other in a statistical sense.

Such models fail to predict extreme behaviour in financial time series because of the intrinsic assumption that such time series conform to a normal distribution, i.e. Gaussian processes that are stationary in which the statistics - the standard deviation, for example - do not change with time.

Random walk models, which underpin the so called Efficient Market Hypothesis (EMH) (Fama, 1965)-(Burton, 1987), have been the basis for financial time series analysis since the work of Bachelier in the late Nineteenth Century. Although the Black-Scholes equation (Black& Scholes, 1973), developed in the 1970s for valuing options, is deterministic (one of the first financial models to achieve determinism), it is still based on the EMH, i.e. stationary Gaussian statistics. The EMH is based on the principle that the current price of an asset fully reflects all available information relevant to it and that new information is immediately incorporated into the price. Thus, in an efficient market, the modelling of asset prices is concerned with modelling the arrival of new information. New information must be independent and random, otherwise it would have been anticipated and would not be new. The arrival of new information can send 'shocks' through the market (depending on the significance of the information) as people react to it and then to each other's reactions. The EMH assumes that there is a rational and unique way to use the available information and that all agents possess this knowledge. Further, the EMH assumes that this 'chain reaction' happens effectively instantaneously. These assumptions are clearly questionable at any and all levels of a complex financial system.

The EMH implies independence of price increments and is typically characterised by a normal of Gaussian Probability Density Function (PDF) which is chosen because most price movements are presumed to be an aggregation of smaller ones, the sums of independent random contributions having a Gaussian PDF. However, it has long been known that financial time series do not follow random walks. This is one of the most fundamental underlying problems associated with financial models, in general.

1.2 The Problem with Economic Models

The principal aim of a financial trader is to attempt to obtain information that can provide some confidence in the immediate future of a stock. This is often based on repeating patterns from the past, patterns that are ultimately based on the interplay between greed and fear. One of the principal components of this aim is based on the observation that there are 'waves within waves' known as Elliot Waves after Ralph Elliot who was among the first to observe this phenomenon on a qualitative basis in 1938. Elliot Waves permeate financial signals when

---

[2] Note that both α and β may very with time.

studied with sufficient detail and imagination. It is these repeating patterns that occupy both the financial investor and the financial systems modeler alike and it is clear that although economies have undergone many changes in the last one hundred years, ignoring scale, the dynamics of market behaviour does not appear to have changed significantly.

In modern economies, the distribution of stock returns and anomalies like market crashes emerge as a result of considerable complex interaction. In the analysis of financial time series it is inevitable that assumptions need to be made with regard to developing a suitable model. This is the most vulnerable stage of the process with regard to developing a financial risk management model as over simplistic assumptions lead to unrealistic solutions. However, by considering the global behaviour of the financial markets, they can be modeled statistically provided the 'macroeconomic system' is complex enough in terms of its network of interconnection and interacting components.

Market behaviour results from either a strong theoretical reasoning or from compelling experimental evidence or both. In econometrics, the processes that create time series have many component parts and the interaction of those components is so complex that a deterministic description is simply not possible. When creating models of complex systems, there is a trade-off between simplifying and deriving the statistics we want to compare with reality and simulation. Stochastic simulation allows us to investigate the effect of various traders' behaviour with regard to the global statistics of the market, an approach that provides for a natural interpretation and an understanding of how the amalgamation of certain concepts leads to these statistics and correlations in time over different scales. One cause of correlations in market price changes (and volatility) is mimetic behaviour, known as herding. In general, market crashes happen when large numbers of agents place sell orders simultaneously creating an imbalance to the extent that market makers are unable to absorb the other side without lowering prices substantially. Most of these agents do not communicate with each other, nor do they take orders from a leader. In fact, most of the time they are in disagreement, and submit roughly the same amount of buy and sell orders. This provides a diffusive economy which underlies the Efficient Market Hypothesis (EMH) and financial portfolio rationalization. The EMH is the basis for the Black-Scholes model developed for the Pricing of Options and Corporate Liabilities for which Scholes won the Nobel Prize for economics in 1997. However, there is a fundamental flaw with this model which is that it is based on a hypothesis (the EMH) that assumes price movements, in particular, the log-derivate of a price, is normally distributed and this is simply not the case. Indeed, all economic time series are characterized by long tail distributions which do not conform to Gaussian statistics thereby making financial risk management models such as the Black-Scholes equation redundant.

1.3 What is the Fractal Market Hypothesis?

The economic basis for the Fractal Market Hypothesis (FMH) is as follows:

- The market is stable when it consists of investors covering a large number of investment horizons which ensures that there is ample liquidity for traders;

- information is more related to market sentiment and technical factors in the short term than in the long term - as investment horizons increase and longer term fundamental information dominates;

- if an event occurs that puts the validity of fundamental information in question, long-term investors either withdraw completely or invest on shorter terms (i.e. when the overall investment horizon of the market shrinks to a uniform level, the market becomes unstable);

- prices reflect a combination of short-term technical and long-term fundamental valuation and thus, short-term price movements are likely to be more volatile than long-term trades - they are more likely to be the result of crowd behaviour;
- if a security has no tie to the economic cycle, then there will be no long-term trend and short-term technical information will dominate.

The model associated with the FMH considered in this is is compounded in a fractional dynamic model that is non-stationary and describes diffusive processes that have a directional bias leading to long tail (non-Gaussian) distributions. We consider a Lévy distribution and show the relation between this distribution and the fractional diffusion equation (Section 4.2). Unlike the EMH, the FMH states that information is valued according to the investment horizon of the investor. Because the different investment horizons value information differently, the diffusion of information is uneven. Unlike most complex physical systems, the agents of an economy, and perhaps to some extent the economy itself, have an extra ingredient, an extra degree of complexity. This ingredient is consciousness which is at the heart of all financial risk management strategies and is, indirectly, a governing issue with regard to the fractional dynamic model used to develop the algorithm now being used by Currency Traders Ireland Limited. By computing an index called the Lévy index, the directional bias associated with a future trend can be forecast. In principle, this can be achieved for any financial time series, providing the algorithm has been finely tuned with regard to the interpretation of a particular data stream and the parameter settings upon which the algorithm relies.

2. The Black-Scholes Model

For many years, investment advisers focused on returns with the occasional caveat 'subject to risk'. Modern Portfolio Theory (MPT) is concerned with a trade-off between risk and return. Nearly all MPT assumes the existence of a risk-free investment, e.g. the return from depositing money in a sound financial institute or investing in equities. In order to gain more profit, the investor must accept greater risk. Why should this be so? Suppose the opportunity exists to make a guaranteed return greater than that from a conventional bank deposit say; then, no (rational) investor would invest any money with the bank. Furthermore, if he/she could also borrow money at less than the return on the alternative investment, then the investor would borrow as much money as possible to invest in the higher yielding opportunity. In response to the pressure of supply and demand, the banks would raise their interest rates. This would attract money for investment with the bank and reduce the profit made by investors who have money borrowed from the bank. (Of course, if such opportunities did arise, the banks would probably be the first to invest savings in them.) There is elasticity in the argument because of various 'friction factors' such as transaction costs, differences in borrowing and lending rates, liquidity laws etc., but on the whole, the principle is sound because the market is saturated with arbitrageurs whose purpose is to seek out and exploit irregularities or miss-pricing.
The concept of successful arbitraging is of great importance in finance. Often loosely stated as, 'there's no such thing as a free lunch', it means that one cannot ever make an instantaneously risk-free profit. More precisely, such opportunities cannot exist for a significant length of time before prices move to eliminate them.

2.1 Financial Derivatives
As markets have grown and evolved, new trading contracts have emerged which use various tricks to manipulate risk. Derivatives are deals, the value of which is derived from (although

not the same as) some underlying asset or interest rate. There are many kinds of derivatives traded on the markets today. These special deals increase the number of moves that players of the economy have available to ensure that the better players have more chance of winning. To illustrate some of the implications of the introduction of derivatives to the financial markets we consider the most simple and common derivative, namely, the option.

2.1.1 Options
An option is the right (but not the obligation) to buy (call) or sell (put) a financial instrument (such as a stock or currency, known as the 'underlying') at an agreed date in the future and at an agreed price, called the strike price. For example, consider an investor who 'speculates' that the value of an asset at price S will rise. The investor could buy shares at S, and if appropriate, sell them later at a higher price. Alternatively, the investor might buy a call option, the right to buy a share at a later date. If the asset is worth more than the strike price on expiry, the holder will be content to exercise the option, immediately sell the stock at the higher price and generate an automatic profit from the difference. The catch is that if the price is less, the holder must accept the loss of the premium paid for the option (which must be paid for at the opening of the contract). If C denotes the value of a call option and E is the strike price, the option is worth $C(S, t) = \max(S - E, 0)$.
Conversely, suppose the investor speculates that an asset is going to fall, then the investor can sell shares or buy puts. If the investor speculates by selling shares that he/she does not own (which in certain circumstances is perfectly legal in many markets), then he/she is selling 'short' and will profit from a fall in the asset. (The opposite of a short position is a 'long' position.) The principal question is how much should one pay for an option? If the value of the asset rises, then so does the value of a call option and vice versa for put options. But how do we quantify exactly how much this gamble is worth? In previous times (prior to the Black-Scholes model which is discussed later) options were bought and sold for the value that individual traders thought they ought to have. The strike prices of these options were usually the 'forward price', which is just the current price adjusted for interest-rate effects. The value of options rises in active or volatile markets because options are more likely to pay out large amounts of money when they expire if market moves have been large, i.e. potential gains are higher, but loss is always limited to the cost of the premium. This gain through successful 'speculation' is not the only role that options play. Another role is Hedging.

2.1.2 Hedging
Suppose an investor already owns shares as a long-term investment, then he/she may wish to insure against a temporary fall in the share price by buying puts as well. The investor would not want to liquidate holdings only to buy them back again later, possibly at a higher price if the estimate of the share price is wrong, and certainly having incurred some transaction costs on the deals. If a temporary fall occurs, the investor has the right to sell his/her holdings for a higher than market price. The investor can then immediately buy them back for less, in this way generating a profit and long-term investment then resumes. If the investor is wrong and a temporary fall does not occur, then the premium is lost for the option but at least the stock is retained, which has continued to rise in value. Since the value of a put option rises when the underlying asset value falls, what happens to a portfolio containing both assets and puts? The answer depends on the ratio. There must exist a ratio at which a small unpredictable movement in the asset does not result in any unpredictable movement in the portfolio. This ratio is instantaneously risk free. The reduction of risk by taking advantage of correlations

between the option price and the underlying price is called 'hedging'. If a market maker can sell an option and hedge away all the risk for the rest of the options life, then a risk free profit is guaranteed.

Why write options? Options are usually sold by banks to companies to protect themselves against adverse movements in the underlying price, in the same way as holders do. In fact, writers of options are no different to holders; they expect to make a profit by taking a view of the market. The writers of calls are effectively taking a short position in the underlying behaviour of the markets. Known as 'bears', these agents believe the price will fall and are therefore also potential customers for puts. The agents taking the opposite view are called 'bulls'. There is a near balance of bears and bulls because if everyone expected the value of a particular asset to do the same thing, then its market price would stabilise (if a reasonable price were agreed on) or diverge (if everyone thought it would rise). Thus, the psychology and dynamics (which must go hand in hand) of the bear/bull cycle play an important role in financial analysis.

The risk associated with individual securities can be hedged through diversification or 'spread betting' and/or various other ways of taking advantage of correlations between different derivatives of the same underlying asset. However, not all risk can be removed by diversification. To some extent, the fortunes of all companies move with the economy. Changes in the money supply, interest rates, exchange rates, taxation, commodity prices, government spending and overseas economies tend to affect all companies in one way or another. This remaining risk is generally referred to as market risk.

2.2 Black-Scholes Analysis

The value of an option can be thought of as a function of the underlying asset price S (a Gaussian random variable) and time t denoted by $V(S,t)$. Here, V can denote a call or a put; indeed, V can be the value of a whole portfolio or different options although for simplicity we can think of it as a simple call or put. Any derivative security whose value depends only on the current value S at time t and which is paid for up front, is taken to satisfy the Black-Scholes equation given by (Black& Scholes, 1973)

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0$$

where σ is the volatility and r is the risk. As with other partial differential equations, an equation of this form may have many solutions. The value of an option should be unique; otherwise, again, arbitrage possibilities would arise. Therefore, to identify the appropriate solution, certain initial, final and boundary conditions need to be imposed. Take for example, a call; here the final condition comes from the arbitrage argument. At $t = T$

$$C(S,t) = \max(S - E, 0)$$

The spatial or asset-price boundary conditions, applied at $S = 0$ and $S \to \infty$ come from the following reasoning: If S is ever zero then dS is zero and will therefore never change. Thus, we have

$$C(0,t) = 0$$

As the asset price increases it becomes more and more likely that the option will be exercised, thus we have

$$C(S,t) \propto S, \qquad S \to \infty$$

Observe, that the Black-Sholes equation has a similarity to the diffusion equation but with additional terms. An appropriate way to solve this equation is to transform it into the diffusion equation for which the solution is well known and, with appropriate Transformations, gives the Black-Scholes formula (Black& Scholes, 1973)

$$C(S,t) = SN(d_1) - Ee^{r(T-t)}N(d_2)$$

where

$$d_1 = \frac{\log(S/E) + (r + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}},$$

$$d_2 = \frac{\log(S/E) + (r - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$

and N is the cumulative normal distribution defined by

$$N(d_1) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{d_1} e^{\frac{1}{2}s^2} ds.$$

The conceptual leap of the Black-Scholes model is to say that traders are not estimating the future price, but are guessing about how volatile the market may be in the future. The model therefore allows banks to define a fair value of an option, because it assumes that the forward price is the mean of the distribution of future market prices. However, this requires a good estimate of the future volatility σ.

The relatively simple and robust way of valuing options using Black-Scholes analysis has rapidly gained in popularity and has universal applications. Black-Scholes analysis for pricing an option is now so closely linked into the markets that the price of an option is usually quoted in option volatilities or 'vols'. However, Black-Scholes analysis is ultimately based on random walk models that assume independent and Gaussian distributed price changes and is thus, based on the EMH.

The theory of modern portfolio management is only valuable if we can be sure that it truly reflects reality for which tests are required. One of the principal issues with regard to this relates to the assumption that the markets are Gaussian distributed. However, it has long been known that financial time series do not adhere to Gaussian statistics. This is the most important of the shortcomings relating to the EMH model (i.e. the failure of the independence and Gaussian distribution of increments assumption) and is fundamental to the inability for EMH-based analysis such as the Black-Scholes equation to explain characteristics of a financial signal such as clustering, flights and failure to explain events such as 'crashesŠ leading to recession. The limitations associated with the EMH are illustrated in Figure 1 which shows a (discrete) financial signal $u(t)$, the derivative of this signal $du(t)/dt$ and a synthesised (zero-mean) Gaussian distributed random signal. There is a marked difference in the characteristics of a real financial signal and a random Gaussian signal. This simple comparison indicates a failure of the statistical independence assumption which underpins the EMH and the superior nature of the Lévy based model that underpins the Fractal Market Hypothesis.

The problems associated with financial modelling using the EMH have prompted a new class of methods for investigating time series obtained from a range of disciplines. For example, Re-scaled Range Analysis (RSRA), e.g. (Hurst, 1951), (Mandelbrot, 1969), which is essentially based on computing and analysing the Hurst exponent (Mandelbrot, 1972), is a useful tool for revealing some well disguised properties of stochastic time series such as persistence

Fig. 1. Financial time series for the Dow-Jones value (close-of-day) from 02-04-1928 to 12-12-2007 (top), the derivative of the same time series (centre) and a zero-mean Gaussian distributed random signal (bottom).

(and anti-persistence) characterized by non-periodic cycles. Non-periodic cycles correspond to trends that persist for irregular periods but with a degree of statistical regularity often associated with non-linear dynamical systems. RSRA is particularly valuable because of its robustness in the presence of noise. The principal assumption associated with RSRA is concerned with the self-affine or fractal nature of the statistical character of a time-series rather than the statistical 'signature' itself. Ralph Elliott first reported on the fractal properties of financial data in 1938. He was the first to observe that segments of financial time series data of different sizes could be scaled in such a way that they were statistically the same producing so called Elliot waves. Since then, many different self-affine models for price variation have been developed, often based on (dynamical) Iterated Function Systems (IFS). These models can capture many properties of a financial time series but are not based on any underlying causal theory.

3. Fractal Time Series and Rescaled Range Analysis

A time series is fractal if the data exhibits statistical self-affinity and has no characteristic scale. The data has no characteristic scale if it has a PDF with an infinite second moment. The data may have an infinite first moment as well; in this case, the data would have no stable mean either. One way to test the financial data for the existence of these moments is to plot them sequentially over increasing time periods to see if they converge. Figure 2 shows that the first moment, the mean, is stable, but that the second moment, the mean square, is not settled. It converges and then suddenly jumps and it is observed that although the variance is not stable, the jumps occur with some statistical regularity. Time series of this type are example of Hurst processes; time series that scale according to the power law,

$$\langle u(t) \rangle_t \propto t^H$$

where H is the Hurst exponent and $\langle u(t) \rangle_t$ denotes the mean value of $u(t)$ at a time t.


Fig. 2. The first and second moments (top and bottom) of the Dow Jones Industrial Average plotted sequentially.

H. E. Hurst (1900-1978) was an English civil engineer who built dams and worked on the Nile river dam project. He studied the Nile so extensively that some Egyptians reportedly nicknamed him 'the father of the Nile.' The Nile river posed an interesting problem for Hurst as a hydrologist. When designing a dam, hydrologists need to estimate the necessary storage capacity of the resulting reservoir. An influx of water occurs through various natural sources (rainfall, river overflows etc.) and a regulated amount needed to be released for primarily agricultural purposes. The storage capacity of a reservoir is based on the net water flow. Hydrologists usually begin by assuming that the water influx is random, a perfectly reasonable assumption when dealing with a complex ecosystem. Hurst, however, had studied the 847-year record that the Egyptians had kept of the Nile river overflows, from 622 to 1469. Hurst noticed that large overflows tended to be followed by large overflows until abruptly, the system would then change to low overflows, which also tended to be followed by low overflows. There seemed to be cycles, but with no predictable period. Standard statistical analysis revealed no significant correlations between observations, so Hurst developed his own methodology. Hurst was aware of Einstein's (1905) work on Brownian motion (the erratic path followed by a particle suspended in a fluid) who observed that the distance the particle covers increased with the square root of time, i.e.

$$R \propto \sqrt{t}$$

where R is the range covered, and t is time. This relationship results from the fact that increments are identically and independently distributed random variables. Hurst's idea was to use this property to test the Nile River's overflows for randomness. In short, his method was as follows: Begin with a time series x_i (with $i = 1, 2, ..., n$) which in Hurst's case was annual discharges of the Nile River. (For markets it might be the daily changes in the price of a stock

index.) Next, create the adjusted series, $y_i = x_i - \bar{x}$ (where $\bar{x}$ is the mean of x_i). Cumulate this time series to give

$$Y_i = \sum_{j=1}^{i} y_j$$

such that the start and end of the series are both zero and there is some curve in between. (The final value, Y_n has to be zero because the mean is zero.) Then, define the range to be the maximum minus the minimum value of this time series,

$$R_n = \max(Y) - \min(Y).$$

This adjusted range, R_n is the distance the systems travels for the time index n, i.e. the distance covered by a random walker if the data set y_i were the set of steps. If we set $n = t$ we can apply Einstein's equation provided that the time series x_i is independent for increasing values of n. However, Einstein's equation only applies to series that are in Brownian motion. Hurst's contribution was to generalize this equation to

$$(R/S)_n = cn^H$$

where S is the standard deviation for the same n observations and c is a constant. We define a Hurst process to be a process with a (fairly) constant H value and the R/S is referred to as the 'rescaled range' because it has zero mean and is expressed in terms of local standard deviations. In general, the R/S value increases according to a power law value equal to H known as the Hurst exponent. This scaling law behaviour is the first connection between Hurst processes and fractal geometry.

Rescaling the adjusted range was a major innovation. Hurst originally performed this operation to enable him to compare diverse phenomenon. Rescaling, fortunately, also allows us to compare time periods many years apart in financial time series. As discussed previously, it is the relative price change and not the change itself that is of interest. Due to inflationary growth, prices themselves are a significantly higher today than in the past, and although relative price changes may be similar, actual price changes and therefore volatility (standard deviation of returns) are significantly higher. Measuring in standard deviations (units of volatility) allows us to minimize this problem. Rescaled range analysis can also describe time series that have no characteristic scale, another characteristic of fractals. By considering the logarithmic version of Hurst's equation, i.e.

$$\log(R/S)_n = \log(c) + H\log(n)$$

it is clear that the Hurst exponent can be estimated by plotting $\log(R/S)$ against the $\log(n)$ and solving for the gradient with a least squares fit. If the system were independently distributed, then $H = 0.5$. Hurst found that the exponent for the Nile River was $H = 0.91$, i.e. the rescaled range increases at a faster rate than the square root of time. This meant that the system was covering more distance than a random process would, and therefore the annual discharges of the Nile had to be correlated.

It is important to appreciate that this method makes no prior assumptions about any underlying distributions, it simply tells us how the system is scaling with respect to time. So how do we interpret the Hurst exponent? We know that $H = 0.5$ is consistent with an independently distributed system. The range $0.5 < H \leq 1$, implies a persistent time series, and a persistent time series is characterized by positive correlations. Theoretically, what happens today will

ultimately have a lasting effect on the future. The range $0 < H \leq 0.5$ indicates anti-persistence which means that the time series covers less ground than a random process. In other words, there are negative correlations. For a system to cover less distance, it must reverse itself more often than a random process.

4. Lévy Processes

Lévy processes are random walks whose distribution has infinite moments and 'long tails'. The statistics of (conventional) physical systems are usually concerned with stochastic fields that have PDFs where (at least) the first two moments (the mean and variance) are well defined and finite. Lévy statistics is concerned with statistical systems where all the moments (starting with the mean) are infinite. Many distributions exist where the mean and variance are finite but are not representative of the process, e.g. the tail of the distribution is significant, where rare but extreme events occur. These distributions include Lévy distributions (Sclesinger et al., 1994), (Nonnenmacher, 1990). Lévy's original approach to deriving such distributions is based on the following question: Under what circumstances does the distribution associated with a random walk of a few steps look the same as the distribution after many steps (except for scaling)? This question is effectively the same as asking under what circumstances do we obtain a random walk that is statistically self-affine. The characteristic function $P(k)$ of such a distribution $p(x)$ was first shown by Lévy to be given by (for symmetric distributions only)

$$P(k) = \exp(-a \mid k \mid^\gamma), \quad 0 < \gamma \leq 2 \qquad (1)$$

where a is a constant and γ is the Lévy index. For $\gamma \geq 2$, the second moment of the Lévy distribution exists and the sums of large numbers of independent trials are Gaussian distributed. For example, if the result were a random walk with a step length distribution governed by $p(x)$, $\gamma \geq 2$, then the result would be normal (Gaussian) diffusion, i.e. a Brownian random walk process. For $\gamma < 2$ the second moment of this PDF (the mean square), diverges and the characteristic scale of the walk is lost. For values of γ between 0 and 2, Lévy's characteristic function corresponds to a PDF of the form

$$p(x) \sim \frac{1}{x^{1+\gamma}}, \quad x \to \infty$$

4.1 Long Tails
If we compare this PDF with a Gaussian distribution given by (ignoring scaling normalisation constants)

$$p(x) = \exp(-\beta x^2)$$

which is the case when $\gamma = 2$ then it is clear that a Lévy distribution has a longer tail. This is illustrated in Figure 3. The long tail Lévy distribution represents a stochastic process in which extreme events are more likely when compared to a Gaussian process. This includes fast moving trends that occur in economic time series analysis. Moreover, the length of the tails of a Lévy distribution is determined by the value of the Lévy index such that the larger the value of the index the shorter the tail becomes. Unlike the Gaussian distribution which has finite statistical moments, the Lévy distribution has infinite moments and 'long tails'.

4.2 Lévy Processes and the Fractional Diffusion Equation
Lévy processes are consistent with a fractional diffusion equation (Alea & Thurnerb, 2005) as shall now be shown. Let $p(x)$ denote the Probability Density Function (PDF) associated with

Fig. 3. Comparison between a Gaussian distribution (blue) for $\beta = 0.0001$ and a Lévy distribution (red) for $\gamma = 0.5$ and $p(0) = 1$.

the position in a one-dimensional space x where a particle can exist as a result of a 'random walk' generated by a sequence of 'elastic scattering' processes (with other like particles). Also, assume that the random walk takes place over a time scale where the random walk 'environment' does not change (i.e. the statistical processes are 'stationary' and do not change with time). Suppose we consider an infinite concentration of particles at a time $t = 0$ to be located at the origin $x = 0$ and described by a perfect spatial impulse, i.e. a delta function $\delta(x)$. Then the characteristic Impulse Response Function f of the 'random walk system' at a short time later $t = \tau$ is given by

$$f(x, \tau) = \delta(x) \otimes_x p(x) = p(x)$$

where $\otimes_x$ denotes the convolution integral over x. Thus, if $f(x, t)$ denotes a macroscopic field at a time t which describes the concentration of a canonical assemble of particles all undergoing the same random walk process, then the field at $t + \tau$ will be given by

$$f(x, t + \tau) = f(x, t) \otimes_x p(x) \tag{2}$$

In terms of the application considered in this paper $f(0, t)$ represents the time varying price difference of a financial index $u(t)$ such as a currency pair, so that, in general,

$$f(x, t) = \frac{\partial}{\partial} u(x, t) \tag{3}$$

From the convolution theorem, in Fourier space, equation (2) becomes

$$F(k, t + \tau) = F(k, t) P(k)$$

where F and P are the Fourier transforms of f and p, respectively. From equation (1), we note that

$$P(k) = 1 - a \mid k \mid^{\gamma}, \quad a \to 0$$

so that we can write

$$\frac{F(k, t+\tau) - F(k,t)}{\tau} \simeq -\frac{a}{\tau} \mid k \mid^{\gamma} F(k,t)$$

which for $\tau \to 0$ gives the fractional diffusion equation

$$\sigma \frac{\partial}{\partial t} f(x,t) = \frac{\partial^{\gamma}}{\partial x^{\gamma}} f(x,t), \quad \gamma \in (0,2]$$

where $\sigma = \tau/a$ and we have used the result

$$\frac{\partial^{\gamma}}{\partial x^{\gamma}} f(x,t) = -\frac{1}{2\pi} \int\limits_{-\infty}^{\infty} \mid k \mid^{\gamma} F(k,t) \exp(ikx) dk$$

However, from equation (3) we can consider the equation

$$\sigma \frac{\partial}{\partial t} u(x,t) = \frac{\partial^{\gamma}}{\partial x^{\gamma}} u(x,t), \quad \gamma \in (0,2] \tag{4}$$

The solution to this equation with the singular initial condition $v(x,0) = \delta(x)$ is given by

$$v(x,t) = \frac{1}{2\pi} \int\limits_{-\infty}^{\infty} \exp(ikx - t \mid k \mid^{\gamma} /\sigma) dk$$

which is itself Lévy distributed. This derivation of the fractional diffusion equation reveals its physical origin in terms of Lévy statistics.

For normalized units $\sigma = 1$ we consider equation (4) for a 'white noise' source function $n(t)$ and a spatial impulse function $-\delta(x)$ so that

$$\frac{\partial^{\gamma}}{\partial x^{\gamma}} u(x,t) - \frac{\partial}{\partial t} u(x,t) = -\delta(x)n(t), \quad \gamma \in (0,2]$$

which, ignoring (complex) scaling constants, has the Green's function solution (?)

$$u(t) = \frac{1}{t^{1-1/\gamma}} \otimes_t n(t) \tag{5}$$

where $\otimes_t$ denotes the convolution integral over t and $u(t) \equiv u(0,t)$. The function $u(t)$ has a Power Spectral Density Function (PSDF) given by (for scaling constant c)

$$\mid U(\omega) \mid^2 = \frac{c}{\mid \omega \mid^{2/\gamma}} \tag{6}$$

where

$$U(\omega) = \int\limits_{-\infty}^{\infty} u(t) \exp(-i\omega t) dt$$

and a self-affine scaling relationship

$$\mathrm{Pr}[u(at)] = a^{1/\gamma} \mathrm{Pr}[u(t)]$$

for scaling parameter $a > 0$ where $\mathrm{Pr}[u(t)]$ denotes the PDF of $u(t)$. This scaling relationship means that the statistical characteristics of $u(t)$ are invariant of time except for scaling factor

$a^{1/\gamma}$. Thus, if $u(t)$ is taken to be a financial signal as a function of time, then the statistical distribution of this function will be the same over different time scales whether, in practice, it is sampled in hours, minutes or seconds, for example.

Equation (5), provides a solution is also consistent with the solution to the fractional diffusion equation

$$\left(\frac{\partial^2}{\partial x^2} - \frac{\partial^q}{\partial t^q} \right) u(x,t) = -\delta(x)n(t)$$

where $\gamma^{-1} = q/2$ (Blackledge, 2010) and where q - the 'Fourier Dimension' - is related to the Hurst exponent by $q = 2H + 1$. Thus, the Lévy index γ, the Fourier Dimension q and the Hurst exponent H are all simply related to each other. Moreover, these parameters quantify stochastic processes that have long tails and thereby by transcend financial models based on normal distributions such as the Black-Scholes model.

4.3 Computational Methods

In this paper, we study the temporal behaviour of q focusing on its predictive power for indicating the likelihood of a future trend in a Forex time series. This is called the 'q-algorithm' and is equivalent to computing time variations in the the Lévy index or the Hurst exponent since $q = 2H + 1 = 2/\gamma$. Given equations (5), for $n(t) = \delta(t)$

$$u(t) = \frac{1}{t^{1-1/\gamma}}$$

and thus

$$\log u(t) = a + \frac{1}{\gamma} \log t$$

where $a = -\log t$. Thus, one way of computing γ is to evaluate the gradient of a plot of $\log u(t)$ against $\log t$. If this is done on a moving window basis then a time series $\gamma(t)$ can be obtained and correlations observed between the behaviour of $\gamma(t)$ and $u(t)$. However, given equation (6), we can also consider the equation

$$\log \mid U(\omega) \mid = b + \frac{1}{\gamma} \log \mid \omega \mid$$

where $b = (\log c)/2$ and evaluate the gradient of a plot of $\log \mid U(\omega) \mid$ against $\log \mid \omega \mid$. In practice this requires the application of a discrete Fourier transform on a moving window basis to compute an estimate of $\gamma(t)$. In this paper, we consider the former (temporal) solution to the problem of computing $q = 2/\gamma$.

5. Application to Forex Trading

The Forex or Foreign Exchange market is the largest and most fluid of the global markets involving trades approaching 4 Trillion per day. The market is primarily concerned with trading currency pairs but includes currency futures and options markets. It is similar to other financial markets but the volume of trade is much higher which comes from the nature of the market in terms of its short term profitability. The market determines the relative values of different currencies and most banks contribute to the market as do financial companies, institutions, individual speculators and investors and even import/export companies. The high volume of the Forex market leads to high liquidity and thereby guarantees stable spreads during a working week and contract execution with relatively small slippages even in aggressive

price movements. In a typical foreign exchange transaction, a party purchases a quantity of one currency by paying a quantity of another currency.

The Forex is a de-centralised 'over the counter market' meaning that there are no agreed centres or exchanges which an investor needs to be connected to in order to trade. It is the largest world wide network allowing customers trade 24 hours per day usually from Monday to Friday. Traders can trade on Forex without any limitations no matter where they live or the time chosen to enter a trade. The accessibility of the Forex market has made it particularly popular with traders and consequently, a range of Forex trading software has been developed for internet based trading. In this paper, we report on a new indicator based on the interpretation of q computed via the Hurst exponent H that has been designed to optimize Forex trading through integration into the MetaTrader 4 system.

6. MetaTrader 4

MetaTrader 4 is a platform for e-trading that is used by online Forex traders (Metatrader 4, 2011) and provides the user with real time internet access to most of the major currency exchange rates over a range of sampling intervals including 1 min, 5 mins, 1 hour and 1 day. The system includes a built-in editor and compiler with access to a user contributed free library of software, articles and help. The software utilizes a proprietary scripting language, MQL4 (MQL4, 2011) (based on C), which enables traders to develop Expert Advisors, custom indicators and scripts. MetaTrader's popularity largely stems from its support of algorithmic trading. This includes a range of indicators and the focus of the work reported in this paper, i.e. the incorporation of a new indicator based on the approach considered in this paper.

6.1 Basic Algorithm - The 'q-Algorithm'

Given a stream of Forex data u_n, $n = 1, 2, ..., N$ where N defines the 'look-back' window or 'period', we consider the Hurst model

$$u_n = cn^H$$

which is linearised by taking the logarithmic transform to give

$$\log(u_n) = \log(c) + H \log(n)$$

where c is a constant of proportionality

The basic algorithm is as follows:

1. For a moving window of length N (moved one element at a time) operating on an array of length L, compute $q_j = 1 + 2H_j$, $j = 1, 2, ..., L - N$ using the Orthogonal Linear Regression Algorithm (Regression, 2011) and plot the result.

2. For a moving window of length M compute the moving average of q_j denoted by $\langle q_j \rangle_i$ and plot the result in the same window as the plot of q_j.

3. Compute the gradient of $\langle q_j \rangle_i$ using a different user defined moving average window of length K and a forward differencing scheme and plot the result.

4. Compute the second gradient of $\langle q_j \rangle_i$ after applying a moving average filter using a centre differencing scheme and plot the result in the same window.

Fig. 4. MetaTrader 4 GUI for new indicators. Top window: Euro-USD exchange rate signal for 1 hour sampled data (blue) and averaged data (red); Centre window: first (red) and second (cyan) gradients of the moving average for $(N, M, K, T) = (512, 10, 100, 0)$. Bottom window: q_j (cyan) and moving average of q_j (Green).

6.2 Fundamental Observations

The gradient of $\langle q_j \rangle_i$ denoted by $\langle q_j \rangle_i'$ provides an assessment of the point in time at which a trend is likely to occur, in particular, the points in time at which $\langle q_j \rangle_i'$ crosses zero. The principal characteristic is compounded in the following observation:

$\langle q_j \rangle_i' > 0$ tends to correlates with an upward trend
$\langle q_j \rangle_i' < 0$ tends correlates with a downward trend

where a change in the polarity of $\langle q_j \rangle_i' < 0$ indicates a change in the trend subject to a given tolerance T. A tolerance zone is therefore established $| \langle q_j \rangle_i' | \in T$ such that if the signal $\langle q_j \rangle_i' > 0$ enters the tolerance zone, then a bar is plotted indicating the end of an upward trend and if $\langle q_j \rangle_i' < 0$ enters the tolerance zone then a bar is plotted indicating the end of a downward trend.

The term 'tends' used above depends on the data and the parameter settings used to process it, in particular, the length of the look-back window used to compute q_j and the size of the window used to compute the moving average. In other words the correlations that are observed are not perfect in all cases and the algorithm therefore needs to be optimised by back-testing and live trading.

Fig. 5. MetaTrader 4 GUI for new indicators. Top window: Euro-USD exchange rate signal for 1 minute sampled data (blue) and averaged data (red); Centre window: first (red) and second (cyan) gradients of the moving average for $(N, M, K, T) = (512, 10, 100, 0)$. Bottom window: q_j (cyan) and moving average of q_j (Green).

The second gradient is computed to provide an estimate of the 'acceleration' associated with the moving average characteristics of q_j denoted by $\langle q_j \rangle_i''$. This parameter tends to correlate with the direction of the trends that occur and therefore provides another indication of the direction in which the markets are moving (the position in time at which the second gradient changes direction occurs at the same point in time at which the first gradient passes through zero). Both the first and second gradients are filtered using a moving average filter to provide a smooth signal.

6.3 Examples Results

Figure 4 shows an example of the MetaTrader GUI with the new indicators included operating on the signal for the Euro-USD exchange rate with 1 hour sampled data. The vertical bars clearly indicate the change in a trend for the window of data provided in this example. The parameters settings (N, M, K, T) for this example are $(512, 10, 100, 0)$. Figure 5 shows a sample of results for the Euro-USD exchange rate for 1 minute sampled data with parameter settings using the same parameter settings In each case, a change in the gradient tends to correlates with a change in the trend of the time series in a way that is reproducible at all scales.

Figure 6 shows examples of Cumulative Profit Reports using the 'q-algorithm' based on trading with four different currencies. The profit margins range from 50%-140% which provides
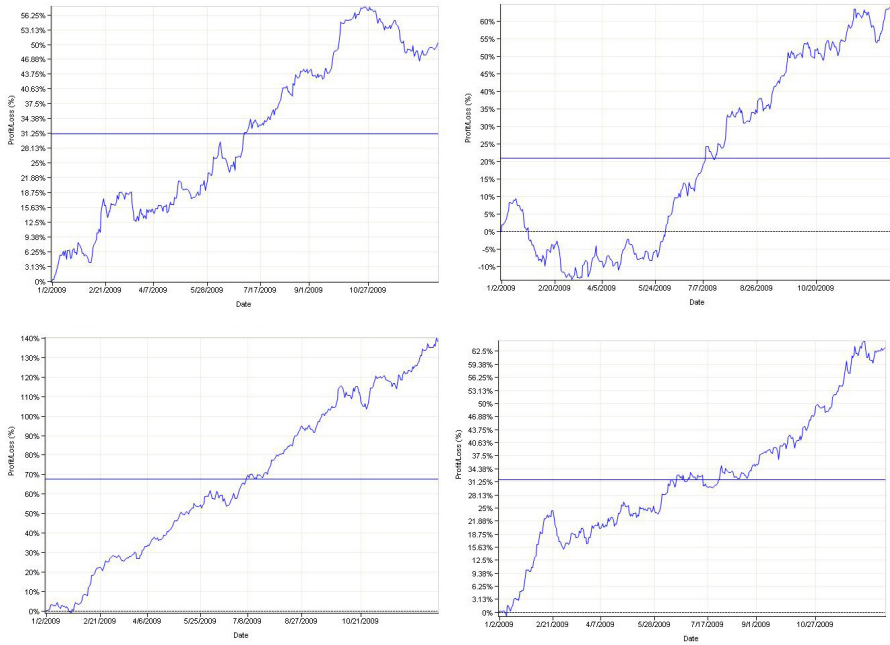
Fig. 6. Example of back-testing the 'q-algorithm'. The plots show Cumulative Profit Reports for four different currency pairs working with 1 hour sampled data from 1/1/2009 - 12/31/2009. Top-left: Euro-USD; Top-right: GGP-JPY; Bottom-left: USD-CAD; Bottom-right: UDSJPY.

evidence for the efficiency of the algorithm based on back-testing examples of this type undertaken to date.

7. Discussion

For Forex data $q(t)$ varies between 1 and 2 as does γ for q in this range since $\gamma^{-1}(t) = q(t)/2$. As the value of q increases, the Lévy index decreases and the tail of the data therefore gets longer. Thus as $q(t)$ increases, so does the likelihood of a trend occurring. In this sense, $q(t)$ provides a measure on the behaviour of an economic time series in terms of a trend (up or down) or otherwise. By applying a moving average filter to $q(t)$ to smooth the data, we obtained a signal $\langle q(t) \rangle(\tau)$ that provides an indication of whether a trend is occurring in the data over a user defined window (the period). This observation reflects a result that is a fundamental kernel of the Fractal Market Hypothesis, namely, that a change in the Lévy index precedes a change in the financial signal from which the index has been computed (from past data). In order to observe this effect more clearly, the gradient $\langle q(t) \rangle'(\tau)$ is taken. This provides the user with a clear indication of a future trend based on the following observation: if $\langle q(t) \rangle'(\tau) > 0$, the trend is positive; if $\langle q(t) \rangle'(\tau) < 0$, the trend is negative; if $\langle q(t) \rangle'(\tau)$ passes through zero a change in the trend may occur. By establishing a tolerance zone associated with a polarity change in $\langle q(t) \rangle'(\tau)$, the importance of any indication of a change of trend can

be regulated in order to optimise a buy or sell order. This is the principle basis and rationale for the 'q-algorithmŠ.

8. Conclusion

The Fractal Market Hypothesis has many conceptual and quantitative advantages over the Efficient Market Hypothesis for modelling and analysing financial data. One of the most important points is that the Fractal Market Hypothesis is consistent with an economic time series that include long tails in which rare but extreme events may occur and, more commonly, trends evolve. In this paper we have focused on the use of the Hypothesis for modelling Forex data and have shown that by computing the Hurst exponent, an algorithm can be designed that appears to accurately predict the upward and downward trends in Forex data over a range of scales subject to appropriate parameter settings and tolerances. The optimisation of these parameters can be undertaken using a range of back-testing trials to develop a strategy for optimising the profitability of Forex trading. In the trials undertaken to date, the system can generate a profitable portfolio over a range of currency exchange rates involving hundreds of Pips[3] and over a range of scales providing the data is consistent and not subject to market shocks generated by entirely unpredictable effects that have a major impact on the markets. This result must be considered in the context that the Forex markets are noisy, especially over smaller time scales, and that the behaviour of these markets can, from time to time, yield a minimal change of Pips when $\langle q(t) \rangle'(\tau)$ is within the tolerance zone establish for a given currency pair exchange rate.

The use of the indicators discussed in this paper for Forex trading is an example of a number of intensive applications and services being developed for financial time series analysis and forecasting. MetaTrader 4 is just one of a range of financial risk management systems that are being used by the wider community for de-centralised market trading, a trend that is set to increase throughout the financial services sector given the current economic environment. The current version of MetaTrader 4 described in this paper is undergoing continuous improvements and assessment, details of which can be obtained from TradersNow.com.

Acknowledgment

9. References

Information and Communications Security Research Group (2011) `http://eleceng.dit.ie/icsrg`

Currency Traders Ireland (2011), `http://www.tradersnoe.com`

Copeland, T. R., Weston, J. F. and Shastri, K. (2003), *Financial Theory and Corporate Policy*, 4th Edition, Addison Wesley.

---

[3] A Pip (Percentage in point) is the smallest price increment in Forex trading.

Martin, J. D., Cox, S. H., McMinn, R. F. and Maminn, R. D. (1997), *The Theory of Finance: Evidence and Applications*, International Thomson Publishing.

Menton, R. C. (1992), *Continuous-Time Finance*, Blackwell Publishers.

Watsham, T. J. and Parramore K., (1996), *Quantitative Methods in Finance*, Thomson Business Press.

Fama, E. (1965), *The Behavior of Stock Market Prices*, Journal of Business Vol. 38, 34-105.

Samuelson, P. (1965), *Proof That Properly Anticipated Prices Fluctuate Randomly*, Industrial Management Review Vol. 6, 41-49.

Fama, E. (1970), *Efficient Capital Markets: A Review of Theory and Empirical Work*, Journal of Finance Vol. 25, 383-417.

Burton, G. M. (1987), *Efficient Market Hypothesis*, The New Palgrave: A Dictionary of Economics, Vol. 2, 120-123.

Black, F. and Scholes M. (1973) *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy, Vol. 81, No. 3, 637-659.

Hurst, H. (1951), *Long-term Storage Capacity of Reservoirs*, Trans. of American Society of Civil Engineers, Vol. 116, 770-808.

Mandelbrot, B. B. and Wallis J. R. (1969), *Robustness of the Rescaled Range R/S in the Measurement of Noncyclic Long Run Statistical Dependence*, Water Resources Research 1969, Vol. 5, No. 5, 967-988.

Mandelbrot, B. B. (1972), *Statistical Methodology for Non-periodic Cycles: From the Covariance to R/S Analysis*, Annals of Economic and Social Measurement, Vol. 1, No. 3, 259-290.

Shlesinger, M. F., Zaslavsky, G. M. and Frisch U. (Eds.), (1994), *Lévy Flights and Related Topics in Physics*, Springer.

Nonnenmacher T. F., *Fractional Integral and Differential Equations for a Class of Lévy-type Probability Densities*, J. Phys. A: Math. Gen., Vol. 23, L697S-L700S.

Abea, S. and Thurnerb, S. (2005), *Anomalous Diffusion in View of EinsteinŠs 1905 Theory of Brownian Motion*, Physica, A(356), Elsevier, 403-407.

Blackledge, J. M., *Application of the Fractional Diffusion Equation for Predicting Market Behaviour*, IAENG International Journal of Applied Mathematics, Vol. 40, No. 3, 130 -158.

MetaTrader 4 Trading Platform, `http://www.metaquotes.net/en/metatrader4`

MQL4 Documentation, `http://docs.mql4.com/`

Nonlinear Regression and Curve Fitting: Orthogonal Regression, `http://www.nlreg.com/orthogonal.htm`

Hothouse, ICT Technologies to License,
`http://www.dit.ie/hothouse/media/dithothouse/`
`techtolicensepdf/Financial%20Risk%20Management.pdf`

Hothouse ICT Video Series,
`http://www.dit.ie/hothouse/technologiestolicence/videos/`
`ictvideos/`