MSci Project

Robert Burbidge

April 17, 2000

Abstract

The fluctuations of the Standard & Poor 500 exhibit scaling behaviour, follow a power law distribution and have a fractal temporal signature suggesting that this is an example of a self-organized critical phenomenon.

1 Introduction

1.1 What Is Self-Organized Criticality?

Science in general is the study of various aspects of the outside world. To better understand a system a mathematical model is constructed encapsulating the salient features of the system. This model can then be tested to see how well it fits the reality. If the model is a good one it can aid understanding and prediction of events. A large class of such models involve dynamical systems. That is, those that are time dependent. Ideally these can be modelled by a set of deterministic equations. Assuming knowledge of all variables, we plug them into our equations and make definite predictions. Unfortunately, either we do not know everything or the equations are very difficult to solve, as in the case of the Navier-Stokes equations for turbulence, or both. This leads us to consider the statistical behaviour of the system — is such and such an outcome more or less likely than another. In any case, there may exist stationary states that, deterministically or statistically, do not change if left alone. We can consider which stationary state a system will end up in given where it started. If the system is perturbed from such a state then it will either return to it or evolve into another state. That is the stationary state can be either an attractor or a repellor for the dynamics respectively.

A critical state of a system is one on the edge of what could be considered as uniformity and chaos [1]. In a uniform state a local perturbation does not affect anything outside its locality; predictions can be made far into the future with the accuracy not time dependent. In a chaotic state the accuracy of predictions falls off exponentially with time. In a critical state local perturbations can propagate throughout the entire system. The probability of a fluctuation of certain size falling off only algebraically with the size. This leads to "$1/f$" noise, the distribution of the lifespan of an event is proportional to $f^{-\beta}$, β near to unity. This can be thought of as the result of a superposition of cycles of all durations

— as opposed to, say, a uniform sine wave or a chaotic spectrum. A critical state also shows a similar power law distribution for the size of events. This is closely related to an underlying fractal geometry. In statistical terms this suggests that the normal (Gaussian) distribution is not an adequate model for criticality. In the Gaussian case the probability of large events falls off exponentially. Instead the Lévy distribution may be proposed. This distribution is leptokurtic, ie. the kurtosis is positive [1] where the kurtosis is given by

$$-3 + \frac{\text{square of second moment}}{\text{fourth moment}}.$$

That is the tails of the distribution are fatter than for the normal.

Since the critical state is on a knife-edge it would seem that some fine-tuning in the laboratory would be required to achieve this state. This is indeed the case for a ferromagnetic material, eg. iron. Below a certain temperature (the Curie point) the material is magnetic, all of the domains have the same direction. Above the Curie point the material is non-magnetic, the domains' alignments are random. At the Curie point the material is in a critical state. Domains with the same direction lie on fractal sets and the effect of a small perturbation (eg. a weak external magnetic field) can alter the macroscopic properties of the material when it has cooled. Other examples of fine-tuning to a critical state are phase transitions. In all cases the critical states are not just qualitatively the same but mathematically identical — depending only on the dimension and symmetries of the system. This is the principle of universality.

The point of self-organized criticality (SOC) is that certain systems actually evolve towards this critical state without any fine-tuning [2]. The critical state is an attractor for the dynamics. I outline some examples of this below.

1.2 Fractal Curves

The concept of fractals [3] is closely tied in with the idea of a fractional dimension. To understand this more clearly we review an elementary feature of the concept of dimension.

A straight line has dimension one. For any positive integer N it can be decomposed into N non-overlapping segments. Each of these segments is $r(N) = 1/N$ smaller than the whole (assuming finiteness). This is the similarity ratio. A plane has dimension two. For every perfect square N it can be decomposed into N nonoverlapping rectangles with a similarity ratio $r(N) = 1/\sqrt{N}$. Generalising, for $N^{1/D}$ a positive integer a D-dimensional hyperplane can be decomposed into N parts with a similarity ratio $r(N) = 1/N^{1/D}$. Thus

$$D = -\log N/\log r(N).$$

Now consider a von Koch curve, this is a continuous, nonrectifiable curve with no well-defined length. Its construction is outlined in figure 1. It is pos-

[1] The population moments are assumed to be infinite but it is possible to give an alternative definition

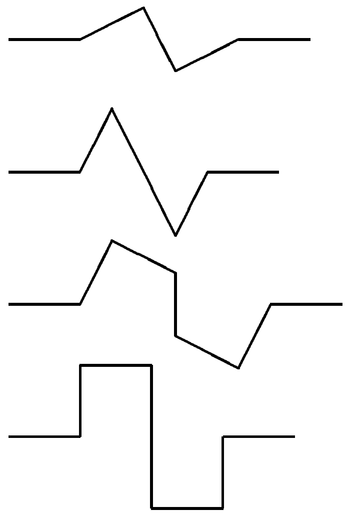

$N = 5,\ r = 1/4,\ d = \log 5/log4$

$N = 6,\ r = 1/4,\ d = \log 6/log4$

$N = 7,\ r = 1/4,\ d = \log 7/log4$

$N = 8,\ r = 1/4,\ d = \log 8/log4 = 1.5$

Figure 1: Nonrectifiable self-similar curves can be obtained as follows. Step 1: Choose any the drawings on the above. Step 2: Replace each of its N legs by a curve deduced from the whole drawing through a similarity of ratio 1/4. One is left with a curve made up of N^2 legs of length $(1/4)^2$. Step 3: Replace each leg by a curve obtained from the whole drawing through similarity of ratio $(1/4)^2$. The desired self-similar curve is approached by an infinite sequence of these steps.

sible to calculate the dimension of these curves using the method above. The dimensions turn out to be fractional, between 1 and 2.

Contrary to previous belief that such objects existed only in mathematicians' heads Mandelbrot [4] proposed that such curves occur naturally. In order to apply these ideas to natural curves, such as coastlines, frontiers etc., Mandelbrot developed a rough sketch of statistical self-similarity and stated that under wide conditions $L(G) \sim G^{1-D}$. Here the length, L, is measured by constructing a polygon of side G with each vertice on the coastline. Applying this to data collected by L. F. Richardson Mandelbrot found that for various coastlines and frontiers there is a strong linear relationship between $\log G$ and $\log L(G)$ that leads to $L(G) = MG^{1-D}$ where M is a positive constant and $D \geq 1$.

In conclusion geographical curves are random self-similar structures of fractional dimension. I use this idea later in showing that the temporal profile of the stock market is random self-similar of dimension 1.47.

1.3 Earthquakes

The Gutenberg-Richter law for earthquakes is an example of a power law. Specifically the number of earthquakes, N, of size m is proportional to $m^{-\tau}$, with $1.25 < \tau < 1.5$ depending on the location. This has been linked by Mandelbrot to the self-similar sets on which earthquakes occur. Bak and Tang [5] propose that this power law is a result of the earth being in a self-organized critical state. The fractal geometric distribution and earthquake dynamics are manifestations of SOC. They use a simple model of the earth's crust and show that it evolves to a critical state with a power law.

The model is as follows. The earth's crust is represented by a two-dimensional (2D) array of particles on a square lattice. Each particle is subject to a force. When this exceeds a critical value (the pinning force at the fault) the particle slips. The system is represented by the force at time t at each point $z(i,j)$. When $z(i,j)$ exceeds some critical value z_c the strain is released at (i,j) and transferred to its neighbours

$$\begin{aligned} z(i,j) &\rightarrow z(i,j)-4 \\ z(i\pm 1,j) &\rightarrow z(i\pm 1,j)+1 \\ z(i,j\pm 1) &\rightarrow z(i,j\pm 1)+1 \end{aligned}$$

Force is conserved except at the boundary. Starting with no force, $z=0$, the tectonic force is represented by

$$z(i,j) \rightarrow z(i,j)+1$$

at a random position. The force is applied until a slip occurs, $z > z_c$, somewhere. The force is transferred to neighbours, if the force at a neighbouring site now exceeds the critical value then further slips occur, and so on. Such a sequence of slips caused by unit increase in the tectonic force is called an 'earthquake'. When the earthquake has stopped the force is reapplied (the tectonic force operates on a geological time scale much greater than the lifetime of an earthquake). At first there are only small propagations triggered by the force. But the system evolves to a state where earthquakes of all sizes occur. The total number of particles that slip is a measure of the energy, E, of the earthquake with distribution $D(E) \approx E^{-\tau}$, $\tau \approx 1.0$. When extended to three dimensions $\tau \approx 1.35$ — closer to the observed values.

The exponents differ between models, depending on the set up and the definition of energy, but the concept of universality suggests that 'the power depends only on geometric and topological features such as the spatial dimension'. As a conclusion they contrast these results with the claim that earthquakes are chaotic phenomenon with few degrees of freedom. The criticality observed here is fundamentally different in that there are an infinity of degrees of freedom. Earthquakes exhibit power law behaviour indicative of a critical state and a simple model of earthquakes self-organizes to a critical state with the same macroscopic behaviour. In summary, the earth's crust is in a self-organized critical state. These ideas can help to construct better models of earthquakes.

1.4 Theoretical Sand Piles

In 1987, Per Bak [6] and others proposed that certain extended, dissipative, dynamical systems naturally evolve into a critical state. The characteristics of this critical state is presence of $1/f$ noise and a scale invariant (fractal) structure. The motivation for this is that $1/f$ noise has been observed in many disparate systems: light from quasars, intensity of sunspots, the current through resistors, the sand flow in an hour glass, the flow of rivers and stock exchange prices. Given the prevalence of $1/f$ noise and of fractals in nature some underlying mechanism needs to be proposed.

To show that SOC is a genuine phenomenon Bak, Tang and Wisenfeld showed that a simple model of a sand pile self-organizes to a critical state. The sand pile model in two dimensions is as follows. Construct an $N \times N$ grid and let the height of sand at each point be $h(x, y)$, put the height difference $z(x, y)$ to be

$$z(x, y) = 2h(x, y) - h(x + 1, y) - h(x, y + 1).$$

Sand is added at a random position (x, y) to give

$$\begin{aligned} z(x-1, y) &\longmapsto z(x-1, y) - 1 \\ z(x, y-1) &\longmapsto z(x, y-1) - 1 \\ z(x, y) &\longmapsto z(x, y) + 2 \end{aligned}$$

When the height difference at a point $z(x, y)$ exceeds some critical value z_c then one grain of sand tumbles in the x and y directions, i.e.

$$\begin{aligned} z(x, y) &\longmapsto z(x, y) - 4 \\ z(x, y \pm 1) &\longmapsto z(x, y \pm 1) + 1 \\ z(x \pm 1, y) &\longmapsto z(x \pm 1, y) + 1 \end{aligned}$$

(This may not be a very realistic model of a sand-pile but is to be emphasised that it is the behaviour of non-linear diffusion dynamics in which we are interested, not sand piles *per se*.)

An avalanche is the movement of two or more grains of sand, caused by the random addition of a single grain. No sand is added while the avalanche progresses. The size, s, of the avalanche is the total number of grains that move and the length, t, is the number of time steps, typically $t < s$. At first only small avalanches occur. After a while the sand pile is sloped such that the addition of a single grain can cause avalanches of all sizes — from two grains to an avalanche comparable to the size of the system. The sand pile has self-organized into the critical state.

When the simulation is performed with closed boundary conditions,

$$z(0, y) = z(x, 0) = z(N + 1, y) = z(x, N + 1) = 0,$$

starting either from a flat surface or from a nonequilibrium state then the system evolves toward the critical state. In the critical state the distribution $D(s)$

follows a power law, specifically when $N = 50$,

$$D(s) \approx s^{-\tau}, \ \tau \approx 1.0.$$

For a three dimensional (3D), $20\times20\times20$, array one finds $\tau \approx 1.37$. Furthermore the distribution of times also follows a power law

$$D(t) \approx t^{-\alpha},$$

where $\alpha \approx 0.43$ for the 2D case and $\alpha \approx 0.97$ for the 3D case. That is, there is a "1/f" spectrum. These results are robust in that if certain bonds are remove between sites (up to 25%) then a power law distribution still holds.

With open boundary conditions at $x = N, \ y = N$ the tumbling equations at the edges become

$$\begin{aligned} z(N,y) &\longmapsto z(x,y)-3 \\ z(x,N) &\longmapsto z(x,N)-3 \\ z(N,N) &\longmapsto z(N,N)-2 \end{aligned}$$

The sand flow $f(t)$ falling off the edge of the box is monitored and a power law is again observed, for a 75×75 system,

$$S(t) \approx f^{-\beta}, \ \beta \approx 0.95.$$

They also note that the clusters formed during an avalanche (ie. its domain) are fractal.

This leads to the conclusion that 1/f noise is not noise at all but 'reflects the generic dynamics of extended dynamical systems'. Furthermore they conclude that '1/f noise is intimately related to the underlying spatial organisation'. Admittedly there is no direct connection between this toy model of turbulence and, for example, the Navier-Stokes equation for turbulence, but there is a 'one-one connection for the phenomenology used to describe the two situations'. Thus, it can be seen that a model having just the same dimension and symmetries as a real life system can evolve toward a critical state without any fine tuning. Invoking the concept of universality the exponents obtained from such models as this can be taken at face value and compared to the real life data.

1.5 Real Sand Piles

In order to see whether or not real sand piles actually evolve to the critical state Glenn A. Held and colleagues [7] designed an experiment to test the predictions. A system was set up whereby sand was dropped, one grain at a time onto a balance. Sometimes the addition of a grain would cause an avalanche and no sand was added during avalanches. Over 350 000 grains of sand were dropped onto a 4cm circular plate. After the sand pile had reached a certain size avalanches of all sizes were observed, in line with the prediction that the sand pile self-organizes to a critical state. When a larger sand pile, with a diameter of 8cm, was used only large avalanches were observed. This was not understood at the time but Per Bak states that since only the sand falling off the pile was measured the smaller avalanches were not detected.

A more intricate experiment was set up by Jens Feder and Torstein Jøssang [8]. Rice was confined to a space between two glass plates and added at a slow rate (about twenty grains per minute) at the upper corner. Various spacings of the plates and various slow feed rates were used. Experiments were done on systems ranging from a few centimetres to a few metres, each lasting 42 hours. Frames were taken every 15 seconds from a high-res (2000×500) monitor. Once the pile had achieved a stationary state monitoring began. The size of an avalanche was defined to be the total downward movement between frames (proportional to the energy dissipated). The surface profile was fractal with features of all sizes. The distributions of the energy E scaled by the length L all lay on the same straight line (for $E/L \geq 10$) on a log-log plot. This indicated a power law for a real-life situation. The critical state is highly robust in that the rice pile evolved to the critical state regardless of size, thickness and feeding rate.

The critical state is not an equilibrium state since energy is not conserved but statistically speaking it is stationary. There is no characteristic energy and the distribution of energy is not changing with time. The critical state is an attractor for the dynamics of the system. Returning to sand piles the experiments show that if the system is perturbed by drying/wetting the sand or by changing boundary conditions the system reverts to the critical state. In view of the stock market, if it does exhibit SOC, then we can argue that although most aggregate fluctuations have no aggregate cause (see below) there are individual events and big players whose actions can perturb the system. In analogy to the above the system will then revert to the stationary (ie. critical) state.

1.6 A Model Economy

In economic theory the cause of aggregate fluctuations is an ill-explained phenomenon. It is unlikely that they are caused by, for example, most households consuming more or less at the same time, or most producers finding it opportune to produce more or less at the same time. Another possible source of an aggregate shock that might cause an aggregate fluctuation is changes in government policy that affect the budgets of the majority of the economic units at the same time. This has been much challenged though. By the law of large numbers one would expect that random shocks cancel out on average. To create an aggregate fluctuation would require a large number of independent events occurring simultaneously to have the same sign, the probability of this decreasing exponentially with the square of the size of the event by the central limit theorem.

Bak et al [9] proposed a simplified model for a chain of interlinked producers and consumers. In this lattice model each intermediate producer receives a demand for goods and either supplies these from stock on hand or orders from the two nearest producers further down the chain. At the top of the lattice are final consumers and at the base are initial producers (who work with raw materials). With some basic assumptions concerning storage and production costs this is a model for studying production and inventory dynamics. If the producers are arranged on a "large" square lattice then a random consumer order can cause an 'avalanche' that involves that involves all levels of the lattice (ie. from consumer to initial producers). The size of an avalanche is considered to be the number of intermediate producers that are stimulated by a single order.

The initial state of the economy at the beginning of some period t is described by the inventory holdings $x_{i,j}(t)$ for each unit (i,j). $s_{i,j}(t)$ is the number of sales by unit (i,j) in period t and $y_{i,j}(t)$ the number of units output in the same period. Then

$$x_{i,j}(t+1) = x_{i,j}(t) + y_{i,j}(t) - s_{i,j}(t).$$

Also optimal production scheduling implies that output is a function of initial inventories and orders received

$$y_{i,j}(t) = y(x_{i,j}(t), s_{i,j}(t))$$

where $y(x,s)$ is a prescribed function based on the assumptions of costing. This gives

$$x_{i,j}(t+1) = x'(x_{i,j}(t), s_{i,j}(t))$$

where $x'(x,s)$ is defined according to $y(x,s)$. The orders received by each unit with $i > 1$ are given by

$$s_{i,j}(t) = \frac{1}{2}(y_{i-1,j}(t) + y_{i-1,j-1}(t))$$

and the orders received by the first row, $s_{1,j}(t)$ are random exogenous shocks (ie. consumer orders). The aggregate demand for goods in a period t is

$$N(t) = \sum_j s_{1,j}(t)$$

and aggregate production is

$$Y(t) = \sum_{i,j} y_{i,j}(t).$$

Assuming independent exogenous shocks, put p the probability that $s_{1,j}(t)$ is one and $1-p$ the probability that it is zero. L is made large with p varying as $L^{-\gamma}$, $\frac{2}{3} < \gamma < 1$. The mean number of final goods orders per period, $N(t)$, is $p(L)L$ which grows as $L^{1-\gamma}$. The random variable $\tilde{N}(t) = N(t)/L^{1-\gamma}$ has a constant mean and as L increases $\tilde{N}$ converges in distribution to a constant. Thus there is 'no aggregate variability in the exogenous flow of final goods orders'. That is the random orders cancel each other out.

The authors then go on to show that despite this there are aggregate shocks in the production. Denoting the size of an avalanche caused by an order at the j^{th} site by $Y_j(t)$ then

$$Y(t) = \sum_{j=1}^{N(t)} Y_j(t).$$

Using the work of Dhar and Ramaswamy they then show that

$$\Pr(Y_j > y) \sim y^{-1/3}$$

for large y. 'Large avalanches are much more likely in this model than in the case of a Gaussian law.' This is highly indicative of SOC and summarises what I feel to be the underlying philosophy ...

> '... aggregate fluctuations in production continue to occur in the large economy limit, even though aggregate exogenous shocks cease to exist.'

1.7 Cotton Prices

In 1963, Mandelbrot [10] proposed that a generalisation of the Gaussian random variable was required to model variations in price $y(t)$ for certain commodities. Bachelier had proposed that successive differences are independent, normally distributed, random variables with zero mean and variance increasing with the differencing interval, $\delta y = y(t+T) - y(t) \sim N(0, \alpha T)$, α some constant. Since the standard deviation of δy is proportional to the price Mandelbrot chose to consider instead $\ln y(t+T) - \ln y(t)$. Histograms of price changes are unimodal and of Gaussian type but have many outliers. For this reason Mandelbrot chose to use a family of probability laws known as stable Paretian, described by Paul Lévy, of which the Gaussian is a limiting case. In his paper Mandelbrot describes these stable Paretian distributions using statistical methods above the level presented here.

Although ignored by economists at the time Mandelbrot proposed a new way of looking at such data which is the basis of the philosophy of SOC. One common approach is to account *a posteriori* for very large fluctuations and eliminate these before any stochastic analysis. Obviously this will make the data more Gaussian. Mandelbrot states that this is unnecessary and there need not be any observable discontinuity between the 'outliers' and the rest of the distribution.

Mandelbrot studied cotton prices and plotted frequency of price changes on a log-log plot. The daily price changes for 1900-05 and 1944-58 and the monthly price changes for 1880-1940 all illustrated the same behaviour. The log of the frequency was directly proportional to the change in the log of price. Each case had the same constant of proportionality. Although Mandelbrot interpreted these results as indicative of a Lévy distribution, it is a simple matter to extract a power law from his analysis. This power law emphasises the continuous transition from small changes to large changes.

Although this is not the limit of Mandelbrot's analysis the features that suggest SOC are that large events can be accounted for by the same dynamics as small events and the power law that this leads to.

1.8 An Economic Index

Rosario N. Mantegna and Eugene Stanley [11] observed that scaling behaviour occurs in systems that exhibit SOC and turbulence. They investigated scaling behaviour in economic systems, particularly financial markets which are subject to precise rules. They showed that the S&P 500 is non-Gaussian and follows a Lévy stable process, as did Mandelbrot for cotton prices. Denoting the successive variations as $\delta y = y(t) - y(t – \delta t)$ with δt logarithmically equally spaced from 1min to 1000min, a semi-log plot shows that δy is distributed leptokurtically. The data for $\delta t = 1$ follow a Lévy stable distribution for all but large variations.

To investigate scaling behaviour they use the 'probability of return' $P(0)$ since the central part fits a Lévy stable distribution and methods involving mainly the wings prove difficult. The data are fitted by plotting $P(0)$ versus δt on a log-log plot. This gives a slope of -0.72 ± 0.025, indicating non-normal scaling behaviour (since slope $\neq -0.5$). Using this exponent the data was scaled according to

$$Z_s = \frac{Z}{(\delta t)^{1/\alpha}}$$

where the scaling factor $\alpha = 1.40 \pm 0.05$ is obtained from the slope under the assumption that the distribution is Lévy for the central part. The scaled data collapse on the $\delta t = 1$ case. They conclude that the Lévy distribution describes the dynamics of $P(Z)$ well over three orders of magnitude. They observe that the price differences would converge to a normal for a time interval of the order of one month. They also show that α is roughly constant over the period of analysis, 1984 to 1989. In summary, the data follow a Lévy distribution which converges towards a normal distribution as δt increases. The index also exhibits scaling behaviour. Given the initial comment this is indicative of the S&P 500 being in a critical state and if so one for which no external parameters can be fine-tuned.

1.9 The Stock Market

The stock market can be thought of as an extended dynamical system. It is not in equilibrium since there is a continual input caused by traders' actions. It is dissipative in that individual actions can have a knock on effect throughout the system. So it is possible that it could naturally evolve to a self-organized critical state. Further evidence for this is the success of modelling an economic index as geometric Brownian motion with drift — a model with fractal properties. The studies outlined above also give weight to this idea. In 1995, E. Stanley and Michael Salinger showed that company growth rate depends only on the size of a company and not, as economists had previously maintained, on the type of technology used for production. It would seem that economists are rather conservative in their views and a new way of looking at the subject could lead to further developments. In analogy to the model economy of Section 1.6 the stock market is an area where there are many interlinked units (traders) each of which can affect closely associated units. There is a hierarchy of players, traders, derivatives dealers, financial institutions, banks and governments throughout which shocks may propagate. The stock market has all the ingredients for a study of SOC.

2 Some Statistical Analysis

2.1 Technical Aspects

The data used for the analysis here are from one of the largest financial markets in the world: The New York Stock Exchange. I study the daily closing prices of the S&P 500 from October 15^{th} 1931 to June 10^{th} 1993. Data were obtained from www.statlib.cmu/edu, a statistical library and analysed using personal C++ programs and the spreadsheet package Excel. It has been shown [12] that overnight price differences do not affect this kind of analysis. The analysis will involve an attempt to model the time series, an investigation of the frequency distribution, the existence of scaling behaviour and associated power laws and the temporal profile of the index.

2.2 Statistical Analysis

2.2.1 Time Series Plots

From the database I denote the value of the index as $y(t)$, this is a random process and for any fixed time t has mean $\overline{y_t}$ and variance v_t. The value of the index is plotted in a time series on page 15.

The features to note are

- The mean is exponentially rising, $\overline{y_t} \propto e^{\alpha t}$.
- The variance is increasing with t and is proportional to the mean.

These factors cause the period 1931 to 1953 to appear relatively featureless. During this period there was a recession and a world war. In a standard time series analysis these factors would be accounnted for separately. However it is in the nature of my analysis to include all data as is. For a more detailed discussion of the possible effects of the war on the analysis see the conclusion.

It is a standard technique to consider the first differences of the series, $\delta y = y(t) - y(t-\delta t)$, as proposed by Bachelier and adopted by Mantegna & Stanley. This time series is shown on page 16, the features to note here are

- The mean is constant at zero.
- The variance is increasing with time.

As we can see although the mean is zero the increasing variance still gives weight to the later values. This leads us to speculate as a model for $y(t)$ geometric Brownian motion with drift [13]

$$y(t) = y_0 e^{\mu t + \sigma W_t - \frac{1}{2}\sigma^2 t}.$$

Where μ is the drift, which accounts for the general upward trend, σ is the volatility and W_t is such that $W_0 = 0$ and W_t is normally distributed with mean 0 and variance t.[2] This can be rearranged to give

$$\ln y(t) = \ln y_0 + \mu t + \sigma W_t - \frac{1}{2}\sigma^2 t.$$

That is, $\ln y(t)$ is normally distributed with mean $\ln y_0 + \mu t$ and variance σt. In view of this we illustrate $\ln y(t)$ on a time series plot. This is on page 17, we observe

- The mean is rising linearly

  $$\overline{\ln y(t)} = \ln y_0 + \mu t$$

  where $\ln y_0 = 1.882$ and $\mu = 0.0009618$ (correlation coefficient (r) is 0.97).

- The variance is constant.

- Cycles of orders three, five and twenty years are apparent.

- Features are visible across the whole time span.

Again we take the first differences, this time of the logged values. Page 18 shows the time series of $Z(\delta t) = \ln y(t) - \ln y(t - \delta t)$. We note

- The mean is zero.

- The variance is constant, Var $= \sigma \delta t = 0.0005320$. This enables us to calculate the volatility, $\sigma = 0.0001330$.

- This series is stationary.

- Features such as the economic turbulence of the 1930s and 'Black Monday' are seen in proportion.

This is now essentially just Bachelier's model adapted for an exponentially rising mean. This is known as the log-normal model. In the following I investigate firstly if this indeed normal. Since we are interested in the fluctuations in the index we take $Z(\delta t)$ as our starting point and assemble seven datasets with δt ranging from 1 to 64 days in logarithmically equally spaced intervals. This gives datasets with from 16384 down to 256 values and a reasonable span over which to observe scaling behaviour.

[2] Also $W_d - W_c$ and $W_b - W_a$ are independent for $a < b < c < d$.

2.2.2 Probability Distribution Of Z(δt)

First we plot the probability distribution of $Z(\delta t)$ for all values of δt, this is shown on page 19.[3]

We observe that the distributions are approximately bell-shaped but highly peaked with mean, median and mode not significantly non-zero and standard deviation increasing with δt.[4]

For comparison, the probability distribution of $Z(\delta t)$ is plotted for each δt alongside the theoretical normal distribution with the same mean and standard deviation. This is shown on pages 20 and 21. This shows that the distributions are more peaked centrally than for the normal model — this is due to the presence of outliers which are most apparent for the larger values of δt. It is possible that the distributions could be Lévy but the statistics involved is beyond the level here. The kurtosis of the distributions ranges from 5 to 22, generally decreasing as the δt increases in line with the predictions that the distribution tends to a normal (zero kurtosis) as the time interval increases.

2.2.3 Scaling Behaviour

To begin with we restrict ourselves to looking at the 'probability of return' $P(0)$, where

$$P(0) = \Pr(Z = 0) = \Pr(\ln y(t) - \ln y(t - \delta t) = 0) = \Pr(\delta y = 0)$$

is the probability that the index remains unchanged after a period δt. This is plotted against δt on page 22. We find that

$$\ln P(0) = -1.403 - 0.3479 \log_2 \delta t, \; r = 0.98$$

which can be rearranged to give

$$P(0) = 0.2459 \delta t^{-0.5019}.$$

A scaling power law behaviour is observed over $1\frac{1}{2}$ orders of magnitude. The exponent of approximately 0.5 is indicative of normal scaling.

We now extend this to look for scaling behaviour over the whole range of Z, although we restrict ourselves to the case $\delta t = 1$ for simplicity since this is the largest dataset. Following Mandlebrot's analysis of cotton prices we calculate the frequency of the fractional change. That is we look at how many times the index changed by 0-5%, 5-10% etc. and plot this on a log-log graph. This is shown on 23. Putting

$$E = \left| \frac{\delta y}{y} \right|$$

and $N(E)$ the number of events of size E we find that

$$\ln N(E) = 0.2128 - 0.8563 \ln E, \; r = 0.88$$

[3]Only the range $-2.0 < Z < 2.0$ is shown although there are outliers.

[4]To be precise sd $= 0.0015 + 0.013 \log_2 \delta t, r = 0.96$

which rearranges to

$$N(E) = 1.237E^{-0.8563}.$$

That is to say that, in an analogue of the Gutenberg-Richter law, the number of index movements of size E is proportional to the size of the movement raised to some power. This is the key feature of SOC. There is no characteristic fluctuation sixe, there is a smooth transition from small events to large ones.

As a final note we confirm the fractal (scale-less) nature of the process $y(t)$ by considering it to be a path (a temporal profile) and calculate the total path length, $L(\delta t)$, given by

$$L(\delta t) = \sum_{i=1}^{2^{14}/\delta t} |y(t) - y(t - \delta t)|.$$

On page 24 is plotted $\log_2 L(\delta t)$ against $\log_2 \delta t$. We find

$$\log_2 L(\delta t) = 13.18 - 0.4691 \log_2 \delta t, \ r = 1.0$$

which can be rearranged to

$$L(\delta t) = 2^{13.18} \delta t^{-0.4691}.$$

Total length increases as we 'zoom in', that is $L(\delta t) \to \infty$ as $\delta t \to \infty$. Hence the fractal dimension of the temporal profile of $y(t)$ is 1.47, this is higher than most coastlines (eg. Britain's coast has a fractal dimension of 1.25) and is about the same as the Koch curve with $N = 8$, $r = 1/4$. This suggests that there are cycles of many orders, seasonal, annual, three yearly, five yearly, twenty yearly and longer as well as further intermediate cycles.

3 Conclusion

The variations of the S&P 500 exhibit scaling behaviour over the range 1 day to 64 days. The scaling behaviour exhibited by the probability of return suggests normal behaviour but the frequency distributions are non-Gaussian. A χ-squared test could be implemented to see whether or not the distributions are significantly non-Gaussian or the data could be modelled to a Lévy distribution. The temporal profile of the index has a fractal dimension of 1.47 indicating that this is a highly non-smooth process. The frequency of percentage changes follows a power law spectrum with a smooth transition from large to small events. In summary, the S&P 500 is in a critical state, indeed a self-organized critical state since there cannot be any fine-tuning involved.

Firstly this tells us that the basic model of geometric Brownian motion with drift may need to be modified to allow the W_t term to follow a non-Gaussian distribution such as a Lévy distribution. This model leads us to consider large fluctuations as not necessarily having specific large causes. At this point it is worth discussing certain incidents in the history of the stock market.

From the time series plots we can see that the period 1943 to 1953 has less fluctuations than other periods. There is a general drift upwards and a slight peak at the end of the war but otherwise this part of the plot is very smooth. Also, immediately apparent on the plot of first differences of the logged index, is the 'Black Monday' crash of October 1987. Another incident (not part of this analysis) is the collapse of the Eastasian tiger economies recently. I draw analogies between these three events and their counterparts on a sand pile.

The war can be considered to be an external factor affecting the system. Much like wetting the sand pile the economy slows down and less avalanches/ fluctuations occur and are generally smaller in size. At the time of Black Monday it was apparent to some that prices were artificially high and that there would soon be a crash (James Goldsmith for instance withdrew completely from the market). In the sand pile model someone standing underneath a cliff that has been building up and looks unstable would similarly get out. In the case of the collapse of the Eastasian economies this was caused by George Soros, and others, 'pummeling the currencies' (as one newspaper put it). This can be thought of as an individual dropping a lot of sand at one point until it collapses. The point of the analogy is that the system is perturbed by an external agent but returns to a critical state (with a different slope/exponent in the power law) because this is an attractor for the dynamics.

This also highlights a flaw in attempting to model systems where the individual units have free will. In most cases this makes no difference but in the case of James Goldsmith/the person under the cliff, their observations of the system can affect the system and actually alter cause a catastrophe. It does not appear that this is successfully incorporated into the model at present. This is where we can return to looking at the local dynamics. Although the theory states that in general it is not possible to predict catastrophes (only their frequency) we can consider the situation locally. If the slope of the sand pile at a particular point is greater or less than the slope in general then we can get out of the way (cash in our options) or dump more sand on (invest) respectively. We are then more likely to avoid catastrophe(cut our losses or make a profit). How we can quantify the slope of the sand pile in terms of the financial market is therefore worth investigating. One idea would be to calculate the fractional dimension of the temporal profile as a function of time and consider its correlogram. In short, more work needs to be done before we make money from this.

References

[1] Buchanan, M. *One Law To Rule Them All* New Scientist 2107 (8^{th} November 1997) 30

[2] Bak, P. *How Nature Works* Oxford University Press 1997

[3] Mandelbrot, B. *The Fractal Geometry of Nature* New York: Freeman 1983

[4] Mandelbrot, B. *How Long is the Coast of Britain?* Science 151 (1967) 637

[5] Bak, P. and Tang, C. *Earthquakes as a Self-Organized Critical Phenomenon* Journal of Geophysical Research 49 (1989) 15635

[6] Bak, P., Tang, C. and Wisenfeld, K. *Self-Organized Criticality* Physical Review Letters A 38 (1988) 364

[7] Bak, P. and Chen, K. *Self-Organized Criticality* Scientific American 264 (January 1991) 46

[8] Frette, V., Christensen, K., Malthe-Sørenssen, A., Feder, J., Jøssang, T. and Meakin, P. *Avalanche Dynamics in a Pile of Rice* Nature 379 (1995) 49

[9] Bak, P., Chen, K., Scheinkman, J.A. and Woodford, M. *Aggregate Fluctuations from Independent Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics* Richerche Economiche 47 (1993) 3

[10] Mandelbrot, B. *The Variation of Certain Speculative Prices* Journal of Business of the University of Chicago 36 (1963) 307

[11] Mantegna, R.N. and Stanley, E. *Scaling Behaviour in the Dynamics of an Economic Index* Nature 376 (July 1995) 46

[12] Brock, W.A. and Kleidon, A.W. J. Econ. Dyn. Contr. 16 (1990) 451

[13] Baxter, M. and Rennie, A. *Financial Calculus* Cambridge University Press 1996