Using Fractals to Improve Currency Risk Management Strategies

Michael K. Lauren
Operational Analysis Section
Defence Technology Agency
New Zealand
m.lauren@dta.mil.nz
Dr_Michael_Lauren@hotmail.com

Abstract

Most businesses in Australasia face great difficulty in planning for future requirements due to often-volatile fluctuations in currency and commodity markets. New analytical techniques based on "fractal analysis" suggest that market data exhibits temporal correlations (i.e. volatile fluctuations tend to occur around the same time, punctuating periods of relative stability), and fat-tailed probability distributions (i.e. more extreme events than might be expected from a normal "bell curve" distribution). By contrast, traditional techniques based on linear models fail to cope correctly with volatility, and may fail badly during these periods. This leads to the possibility that some companies have been making policy and optimisation decisions based on methods that may be detrimental. In this paper it is discussed how fractal methods can correctly characterise the currency market, and how this might affect trading strategies.

1 Introduction

Analysis of financial markets by many workers has revealed probability distribution functions that display "fat tails" and that follow power laws. Mandelbrot was the first to show this, using daily prices for cotton in 1963 [1]. This feature has been confirmed with more detailed data since then.

Despite this, most quantitative finance methods are based on Brownian motion, following the work of Batchelier [2] in 1900. This approach drastically underestimates large amplitude fluctuations that occur during periods of severe volatility. The widespread use of these models is attributable to the use of the Black-Scholes equation [3], which is applied to model the fair value for derivatives trading.

These risk-reducing methods rely on a number of assumptions. The first is that the movements from one day to the next are independent of each other. The second is that they follow a normal distribution with a well-defined standard deviation. This assumes large spikes in prices are extremely rare events which almost never happen.

However, in practice one observes spikes on a regular basis, often as frequently as once per month [4]. Furthermore, while the size of the daily fluctuations can remain roughly constant for long periods, the variability may suddenly jump for a briefer period. Such periods are often referred to as examples of "clustered volatility". In simple terms, clustered volatility means that the largest movements tend to occur around the same time, rather than at random intervals. Thus the size of price movements are not independent, but display a degree of correlation in time.

Here we discuss alternative approaches to estimating derivative prices based on fractal cascade methods that at least partly owe their origin to Mandelbrot. It will be seen shortly that the statistical properties of financial markets are well represented by such fractal models, particularly the property of clustered volatility. This is most obvious when examining the distribution of the gradients of the data (i.e. the size of the daily changes, for example).

1.1 Spikes, Fat-tailed Distributions and Clustered Volatility

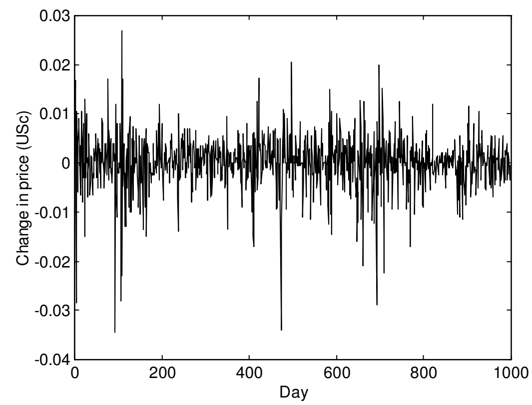

Figure 1. Daily changes in the value of the Kiwi dollar for a 1000-day period.

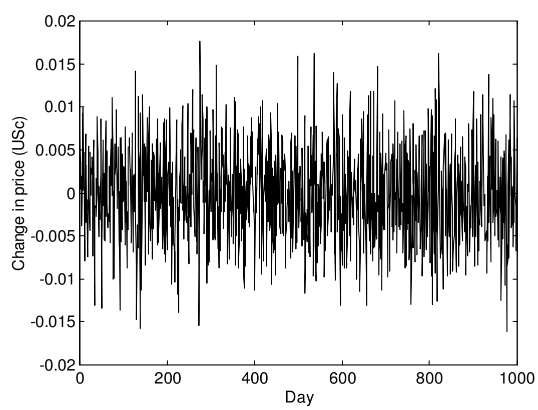

Figure 2: Simulated changes in prices using a normal distribution, with the same standard deviation as the real data above.

Figure 1 shows changes in the value of the Kiwi dollar for 1000 data points during the late 1980s and early 90s. The data shown are qualitatively quite different from those simulated in Figure 2 by taking random values from a normal distribution. One might speculate that the probability distribution for the data in Figure 1 would more closely resemble a lognormal distribution (after taking absolute values). However, the actual distribution has a greater proportion of extreme values than can be accounted for by a distribution of lognormal form. Such distributions are often said to exhibit “fat tails” or “fat wings”. The distribution is most accurately described by a Levy-Stable type distribution, of which the Cauchy distribution is an example.

To appreciate the advantages of fractal methods it is important to understand that the standard deviation of the data is not a good measure of risk. This is because the standard deviation changes dramatically from one subset of data to another. This is symptomatic of clustered volatility. Furthermore, the standard deviation of financial data can be shown to depend on the resolution at which the data is measured (which is picked arbitrarily). For example, Figure 3 shows the change of price data for the New Zealand dollar squared. From the figure, it is apparent that the largest changes are clustered together in time.

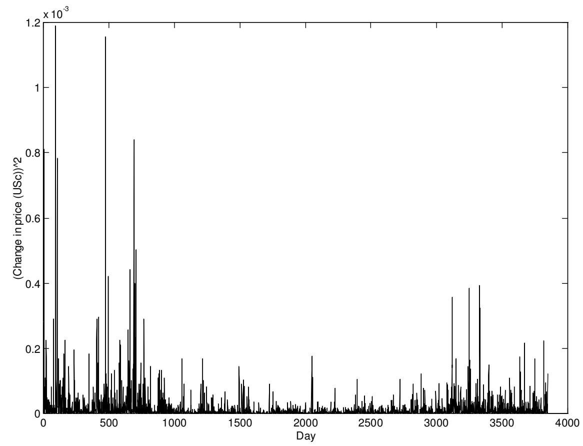

Figure 3: The square of the change in price for the Kiwi dollar highlights the degree to which the data exhibits clustered volatility.

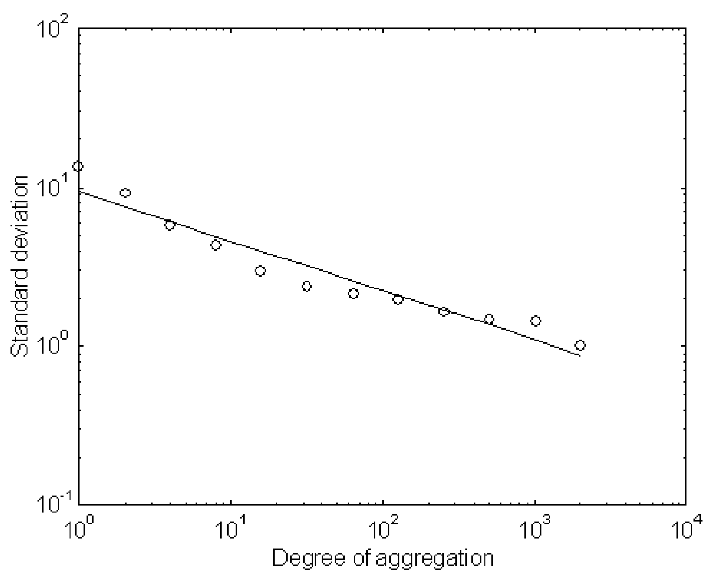

Figure 4: Dependence of the value of the standard deviation of the data shown in Figure 3 on the degree of aggregation.

If we measure the standard deviation of this data, we find that it depends on the degree of aggregation. This dependence is plotted in Figure 4. When plotted like this on a log-log scale, the standard deviation is found to obey a power-law (i.e. straight line on a log-log plot) dependence on the data resolution. The slope of the straight line is effectively the fractal dimension associated with the standard deviation (i.e. second-order statistical moment). Here the slope is –0.27.

Similar plots can be made for the other statistical moments, thus allowing a "spectrum" of fractal dimensions to be obtained, which characterise the statistical nature of the data completely. This type of characterisation is called a "multifractal" analysis. Unlike the standard deviation, this fractal dimension characterisation is scale-invariant.

2 Simulations

Obtaining such a characterisation allows a properly parameterised fractal model to generate synthetic data virtually indistinguishable from the real data, and so exhibits periods of clustered volatility. One such model is a fractal cascade. The idea of a fractal cascade is that it takes an interval over which a function is uniformly distributed, splits that interval into some fraction (say half), and reassigns new values to that section according to some kind of random generating function. This process is then reiterated on each new interval, and so on, until sufficient new iterations have occurred to generate a convincing facsimile. The process is illustrated schematically in Figure 5.

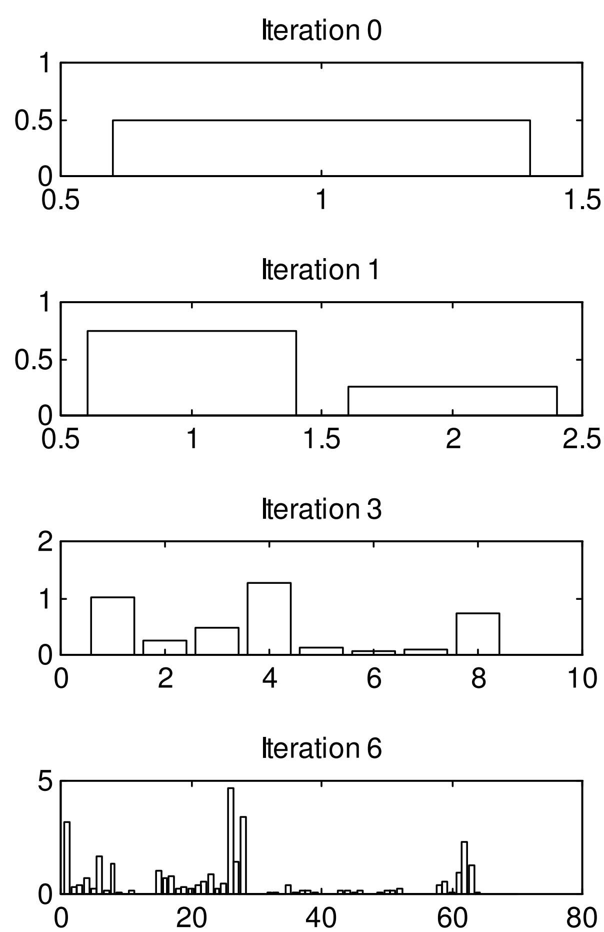

Figure 5: Schematic representation of a cascade process.

How well the simulated data display the multifractal properties of the actual data depends on the form of the generator used. Mandelbrot suggests a simple generator for financial data in [4]. The generator used to simulate the data discussed in this work is a proprietary model developed by the author.

The method generated changes in price using a similar mechanism to that shown in Figure 5, but in a way which also produces changes of negative sign as well as positive. After generating a multifractal distribution of price changes in this way, the changes must then be scaled appropriately to represent the actual market. This is done by normalising the simulated gradient data so that they have a standard deviation of 1, and then multiplying them by the standard deviation of the actual market changes. From these data it is a simple matter to construct artificial market data.

The main aspect we examine in this paper is the difference between multifractal and Brownian motion models. We can equally simulate market data using a Brownian motion approach, by multiplying the current price by fractional changes determined by random numbers taken from a lognormal distribution.

Using data representing a period of a little over 10 years, the standard deviation of the fractional changes in price was found to be 0.0065 (calculated from data expressing the value of the Kiwi in terms of US dollars). This standard deviation was used for both the multifractal and Brownian motion models.

Figure 6 shows 1000 different possible paths for the Kiwi dollar to follow, simulated using a multifractal cascade, along with the same number of paths generated using a Brownian motion approach. It also shows actual Kiwi dollar data, chosen from a volatile period of trading.

At first glance, it might appear that the fractal model has a larger standard deviation than the random walk data. Closer examination of the probability density function (shown in Figure 7) at day 120 reveals that this is not the case. In fact, both distributions have the same mean and standard deviation. We see that the PDF generated by the random walk data obeys an approximately normal distribution, while the data generated by the multifractal approach is slightly nongaussian. In particular, the fractal data displays a higher concentration around the mean value, but also a higher concentration of extreme values than the random walk data (i.e. “fat wings”). By contrast, the random walk data shown is almost exactly gaussian in distribution.

While the difference in the distributions does not look large, it will be seen below that it has implications for the theoretical value of hedging instruments.

3 Hedging

Having obtained a method for generating synthetic market data which display the correct statistical properties of the historical Kiwi dollar prices, it is possible to calculate the price that should be paid for hedging instruments, such as call options.

A call option gives the holder the right to by a unit of currency for a particular price at a particular time. If the Kiwi is currently worth 42USc, it costs an importer 238NZc to buy \$US1 worth of goods. If the importer knows that he or she will make a loss on reselling the goods if more than 250NZc per US dollar is paid for them, he or she can eliminate the risk of this happening by buying call options. For example, the importer may buy call options that allow him/her to buy US dollars at 250NZc each. If the cost of US dollars goes above this level, then the importer exercises the options, if it doesn’t, then the importer doesn’t.

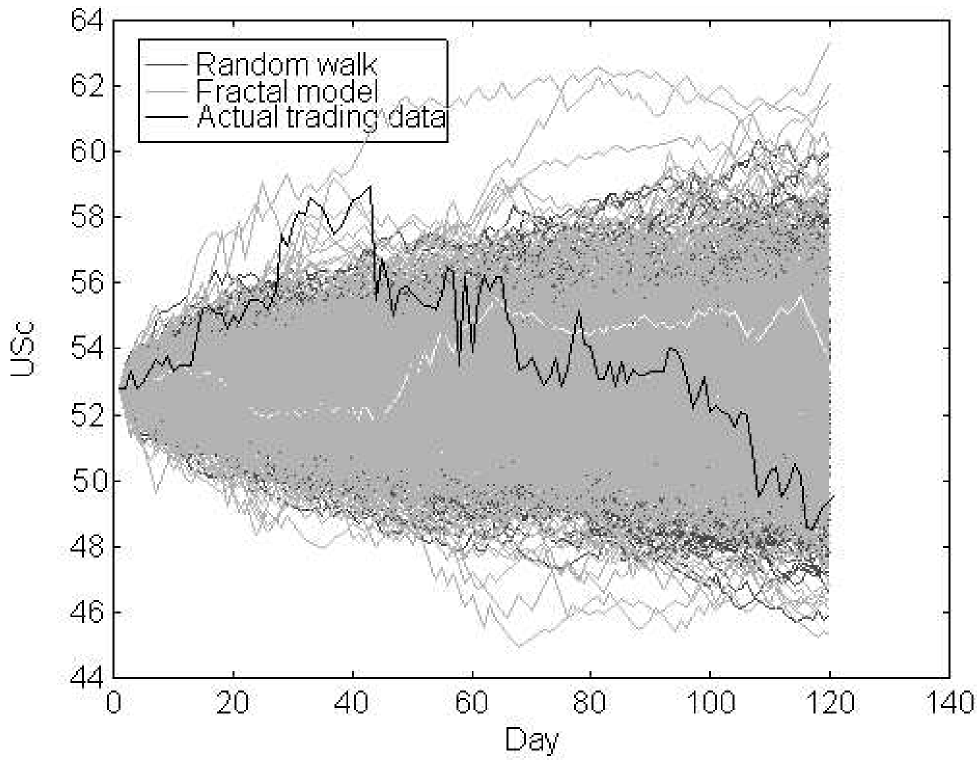

Figure 6: Simulated Kiwi dollar trading using random walk and multifractal techniques (1000 runs each). The figure also shows some actual trading data for the Kiwi.

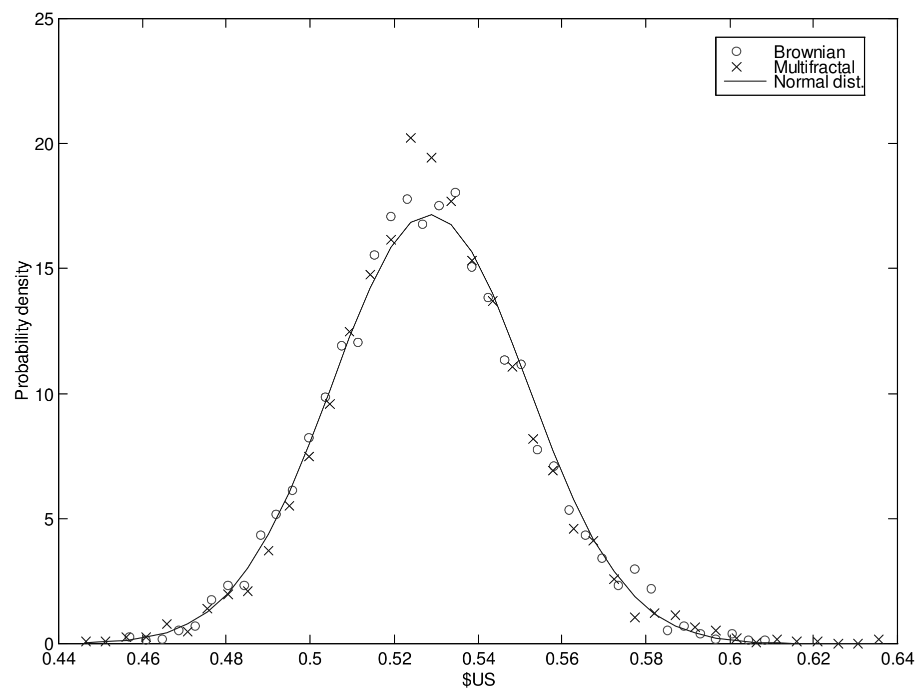

Figure 7: The PDF for the value of the Kiwi dollar 120 days into the future, starting at a value of 52Usc, calculated using multifractal and Brownian motion methods.

We can calculate the reasonable price an importer may expect to pay for the options by obtaining the expectation value for the cost of US dollars at the exercise date, and modifying it by the return that would be gained if the price of the option were instead invested in a guaranteed fixed-interest investment.

That is to say, we expect the value of the option to be:

$$w = e^{-rt} \int_{c}^{\infty} (x - c) f(x)\, dx$$

where w is the theoretical value of the option, x is the price of the US dollar, t is the time to maturity, r is the risk-free interest rate, c is the strike price, and f is a function describing the PDF of x.

The value calculated will depend on the shape of the "wings" of the PDF function, since importers are mostly interested in swings away from the mean value. As we saw in the previous section, the wings for the random walk and multifractal cases are different, so we expect different answers for each.

Suppose our maturity date is 90 days into the future, fixed interest rates are 5% per annum, and the "strike" price for the options is 250NZc (this is the price the options allow US dollars to be bought for). Using the data simulated produces estimates of 0.35NZc and 0.44NZc via the multifractal and Brownian motion approaches respectively. This suggests that the Brownian motion approach overestimates the value of the options, due to the higher concentration of the multifractal PDF about the mean.

However, the situation changes if the maturity date is much closer, say in 30 days. The estimated values of options is then 0.04NZc and 0.01NZc respectively; i.e. the Brownian case now significantly underestimates the value relative to the multifractal case. We conclude from this that the closer to the maturity date one is, the more important the "fat wings" of the multifractal distribution is.

4 Risk mitigation

In the calculations above, we have been concerned with the value of options. But this has been calculated from the point of view of a speculative trader in options. The options may be considerably more valuable to a business from the point of view of risk mitigation. As we have seen, the difference between the Monte Carlo value estimates for the multifractal and Brownian motion cases are significant but not large in absolute terms.

Given that the multifractal data display fatter-tailed distributions than the Brownian case, we might ask what is the impact of this characteristic on risk?

Note that the value of options are calculated using expectation values, but expectation values may not be a good indication of the worth of an option if a large fluctuation in currency price will put the importer out of business. As discussed, multifractal and actual financial data display large spikes in prices much more frequently than normal distributions.

We can use our multifractal characterisation of the Kiwi dollar to estimate the risk of large fluctuations in currency. Figure 9 shows the plot of 1 – (cumulative probability) estimated for the Kiwi 30 days into the future, with a starting value of 238NZc, using both the multifractal and Brownian motion models. Zooming in on the most extreme 5% of cases, we note that the multifractal case exhibits higher probabilities of greater movement in the cost of US dollars.

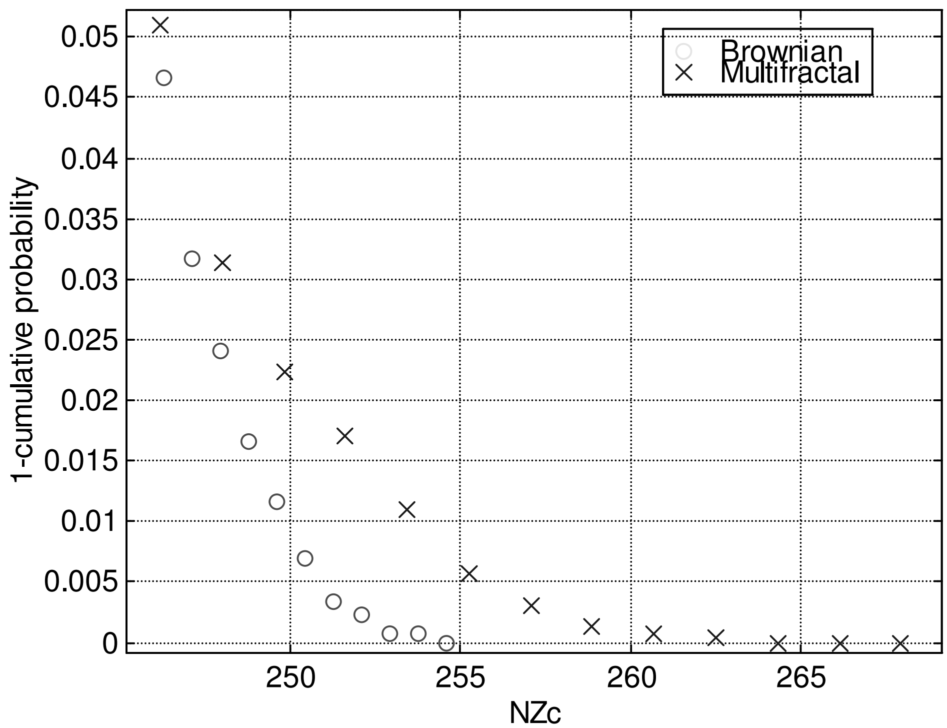

Figure 9: Probability that the currency will finish at greater than a given level (starting from 238NZc), estimated using the Brownian motion and multifractal models.

For example, the multifractal case suggests that a greater than 16NZc (6.7%) increase in the value of the US dollar will occur with 1% probability, whereas the Brownian case suggests that it should occur with only a 0.1% probability. For an increase of greater than 19NZc (7.4%), the multifractal case suggests a 0.3% probability, whereas this is outside the range of estimates for the Brownian case.

Probabilities of 1% and 0.3% may seem small, but it must be remembered that the size of the swings are large for the relatively short time to maturity. Such changes in price are the financial equivalent of a typhoon, and the multifractal model suggests that they are many times more likely to occur than the Brownian model.

5 Conclusions

It has been demonstrated how a multifractal model can be constructed to estimate currency fluctuations.

The multifractal model conforms to the statistical properties of clustered volatility, unlike a Brownian motion approach. The effect of this is that the Brownian model underestimates the possible changes in the value of the currency over short time spans, but can overestimate them over longer time spans.

This can be understood in terms of the effects of clustered volatility. A period of clustered volatility may change the price of a currency rapidly in a short period of time, but won't necessarily impact on the longer-term price.

One might suspect that traders using a Black-Scholes or Brownian motion approach might explain the failure of these approaches to explain rapid movements over short periods as being due to not estimating the volatility (standard deviation) correctly.

However, the property of clustered volatility means that the standard deviation of financial data can be highly variable, making on-the-fly estimates of volatility highly questionable.

For this study, we used the standard deviation from more than 10 years worth of data, and assumed this STD was representative of the entire period. Figure 6 suggests that this approach works well for the multifractal model, but not so well for the Brownian model.

If the multifractal model described here really does provide a more realistic picture of market risks, it would be foolhardy not to investigate them in more detail. Certainly the study presented here illustrates a number of differences between this approach and the more traditional Brownian methods.

It is worth noting that the Kiwi dollar displays relatively low volatility for a financial market, and it is not "strongly" multifractal. By this we mean that the behaviour of the currency does not deviate greatly from a Brownian motion case, whereas some multifractal financial data may deviate much more significantly, and so the consequences will be much greater.

References:

[1] Mandelbrot, B. B. 1963. "The Variation of Certain Speculative Prices", *Journal of Business,* **36**(4), 394-419.

[2] Batchelier, L. 1900. "Theorie de la speculation", *Annales Scientifiques de l'Ecole Normale Superieure*, 21-86.

[3] Black, F. and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, **81**, 637-654.

[4] Mandelbrot, B. B. 1999. "A Multifractal Walk Down Wall Street", *Scientific American*, Feb 1999, 50-53.