Arbitrage in a Fractional Brownian Motion Market

Stewart Hodges

Financial Options Research Centre
University of Warwick

First draft October 1994
This version March 1995, with minor revisions May 1995.
Please do not quote.

I am grateful to Michael Dempster, Chris Rogers, and Michael Selby for help with some of the ideas in this paper, and to participants at a research seminar at the University of Essex for comments on an earlier draft.

Further acknowledgements: FORC Members and INIMS.

Arbitrage in a Fractional Brownian Motion Market

Stewart D Hodges

Abstract

Various commentators have suggested that Fractional Brownian Motion may be an appropriate model for financial markets. However, this seems a surprising choice, since arbitrage opportunities exist within this kind of structure. This paper explores the nature of the arbitrage opportunities in a Fractional Brownian Motion market. We describe why the risk neutral probabilities are degenerate and investigate the numbers of transactions required to obtain essentially riskless profits. We conclude that for a market with a Hurst exponent outside the range 0.4 to 0.6 less than 300 transactions would be required. Fractional Brownian Motion is thus a quite inappropriate model for a financial market unless that market really is grossly inefficient.

Arbitrage in a Fractional Brownian Motion Market

Stewart Hodges

Financial Options Research Centre
University of Warwick

1. Introduction

Various commentators have suggested that fractal models may apply to financial markets, and in particular that fractional Brownian motion may be an appropriate model. For example, see Peters (1994). This seems a surprising choice when we consider the probable arbitrage opportunities within this kind of structure. This short paper sets out to explore the nature of the arbitrage possibilities within a fractional Brownian motion market.

The structure of the paper is as follows. First, we consider the key properties of fractional Brownian motion and how we may generate time series with these properties. Next, we consider informally what kind of pricing operator we would expect to find in such a market, and why it must in some sense breakdown as we allow increasing numbers of transactions. Here we will need to distinguish between cases of positive serial correlation (where the Hurst exponent is greater than a half), and cases of negative correlation (where it is less than a half). With these considerations in mind we then examine the rewards to trading strategies under some alternative assumptions.

We conclude that for a fractional Brownian motion market with a Hurst exponent sensibly different from 1/2, riskless arbitrage can be accomplished by means of at most a few hundred transactions.

2. Fractal Brownian Motion

Our starting point is from two of the the key properties of fractional Brownian motion, F_t: that all its conditional distributions are Normal, and that the standard deviation increases as t^H where H is some constant between 0 and 1 which we will refer to as the Hurst exponent. The case where $H = 1/2$ corresponds to that of the usual standard Brownian Motion. Hurst *et al* (1965) found that exponents different from one half seem to occur in a variety of physical processes, such as water flows in rivers, rainfall and temperatures, thickness of tree rings and even wheat prices.

There is also some evidence that the standard deviation of the returns in financial markets do not grow exactly with the square root of time. Peters (1991) describes trending in equity markets and reports estimated Hurst exponents such as $H = 0.61$ for the pound/dollar exchange rate (daily data 1973-1989), and as high as $H = 0.78$ for the S&P 500 Index (monthly data 1950-1988). His "fractal market hypothesis" is extended and elaborated in Peters (1994) which describes fractional Brownian motion and also provides some further empirical analysis.

Peters (1994) introduces what he calls the Fractal Market Hypothesis which he claims gives an economic and mathematical structure to fractal market analysis. He starts with uncontentious assumptions: markets combine the activities of traders with very different investment horizons. The information set that is important to each investment horizon is different. We are then told that under this "fractal" structure, with no characteristic time scale, the market remains stable. Verbal analogies are made to other well known fractal strucures, such as the physical structure of the lungs, to evoke an aura of plausibility when we leap to conclusions which cannot be sustained rigorously. His Fractal Market Hypothesis remains an extremely vague creation, and the link between market structure and the fractal behaviour of prices is never properly completed.

Fractional Brownian motion may be defined in terms of the power kernel which generates it. Mandelbrot and Van Ness (1968), for example, define fractional Brownian motion in this way as "a moving average of standard Brownian increments $dB(t)$, in which past increments of B(t) are weighted by the kernel $(t\text{-}s)^{H-1/2}$." This is shown in the following equation:

$$F(t) = \int_{-\infty}^{t} (t-s)^{H-1/2} \, dB(s)$$

It is easily verified that the variance of this integrates out to the form already described.

3. Pricing Functions: Informal Arguments

Figure 1 shows a graph of simulated time series for fractional Brownian motion with Hurst exponents of 0.25, 0.5 and 0.75 respectively. These series were generated simply by applying the continuous kernel to discrete Gaussian increments, so they are only approximately fractal. Rimbaldi and Pinazza (1994) provide an alternative more efficient algorithm. The smoothness of the $H = 0.75$ series and the jaggedness of the $H = 0.25$ one are particularly striking. These degrees of positive serial correlation from a Hurst exponent greater than a half, and negative serial correlation when it is less than a half would appear to be incompatible with an efficient market. Do they provide risky opportunities for speculation or do they enable essentially riskless arbitrage to take place?

In this section we will give some heuristic arguments to suggest the nature of the pricing function in these markets. In other words, we will argue as to what form the risk neutral density function (or alternatively the Arrow-Debreu state-price density function) would have to take if it existed. We shall see that it ought to be degenerate and lead to particular sorts of arbitrages. A recent paper by Cutland, Kopp and Willinger (1993) shows that arbitrage exists in the case where $H > 1/2$, but does not elaborate on the nature or properties of the arbitrage strategies. Even more recently, Rogers (1995) provides a more rigorous (and difficult) mathematical demonstration of the existence of arbitrage, but does not consider the number of transactions involved.

The key properties are the ones already stated, and we will distinguish between the two cases $H > 1/2$ and $H < 1/2$. In each case we will assume that the standard deviation of return with a one year horizon is σ. This of course represents the situation with the best possible information based on observing the process continuously for the whole of the infinite past. Over any small time interval Δt, the standard deviation of (the unpredictable component of) return is $(\Delta t)^H$. In the "trending" case of $H > 1/2$, these are, of course, smaller than they would be in the independent and identically distributed case corresponding to $H = 1/2$. The size of these short increments would usually be interpreted as corresponding to an annualised standard deviation of $\sigma\,(\Delta t)^{H-1/2}$ and it is this figure which would be relevant for hedging the risk of an option contract with rebalancing at time intervals Δt apart. (A recent paper by Lo (1993), for example, provides a more careful justification that it is simply the unpredictable component of return over the period between revisions which matters for hedging and hence valuing options). In this "trending" case, our annualised volatility tends to zero as Δt tends to zero, and it looks as though we should price options as though the volatility of the asset is zero! In the case of the negative serial correlation when $H < 1/2$, we have the same formula for our "annualised standard deviation", but in this case it tends to infinity as Δt tends to zero.

We will now consider what the risk neutral density functions would have to be like to support such extreme option values (with respectively zero or infinite implied volatilities). We shall discover that they are quite degenerate! This is most easily seen by remembering that the risk neutral probabilities can be obtained as the second derivative of call option prices with respect to their strike price. The simpler case is for zero volatility, where we have

$$C = \mathrm{Max}\{0, S_0 - \mathrm{PV}(X)\}.$$

The pricing function is a blob of mass at $S_t = \mathrm{FV}(S_0)$, with zero elsewhere. Since we know that other values of S_t are possible we seem to have some arbitrage possibilities. The case of infinite volatility is more interesting. With infinite variance, however high the strike price, the value of the call option becomes equal to the value of the underlying asset and the value of the put option to that of the present value of the strike! This demands an even more perverse risk neutral probability distribution with most of the mass at $S_t = 0$ and the rest somewhere out at infinity. Again, this suggests that some kind of pure arbitrage must be possible. It is also clear that in neither case will a continuous change of probability measure exist (at least in the limit as Δt tends to zero).

4. Arbitrages

The arguments of the previous section suggest strongly that at least in the limit as discrete trading becomes continuous there should be some kind of risk free arbitrage under the postulated market conditions. One thought is to establish arbitrages through options related strategies. We shall discuss these possibilities first, and then look at simpler prediction based strategies in more detail. In a "trending" market with $H > 1/2$ options are extremely cheap to manufacture. In fact we would expect the naive strategy of "hedging through the cap" to actually work, (see, for example, Hull (1993) for a description of this in a Geometric Brownian Motion context). Thus if we implemented a delta hedging strategy to manufacture the outcomes from purchasing a straddle it should cost essentially nothing, and we would obtain the non-negative straddle payoff essentially free. Options strategies in the case of a negatively correlated $H < 1/2$ market seem to be a little more complicated. The risk neutral probability mass has all disappeared to the extremes, and although options are expensive to create, they are all equally expensive. We should be able to form profitable delta-hedging strategies to capture profits in the mass of the objective density function, eg. to replicate the outcomes from a butterfly spread or other similar position. However, in order to implement such a strategy we would need to choose a time interval for revisions and calculate the changing deltas of the various options involved in the position.

Although the strategies just described ought to work and generate profits, which become riskless as the frequency of trading is increased to the continuous limit (without any transactions costs!), they would not be optimal in any sense. Larger and more consistent profits for the same levels of activity should be possible based on elementary mean-variance based trading. If we know, or can estimate, the value of the Hurst exponent for a fractional Brownian Motion market we should be able to construct short term forecasts of sufficient accuracy that trading on them rapidly becomes riskless.

Treynor and Black (1973) established that when optimal investments are made using forecasts which have a correlation of ρ with their outcomes, the optimised Sharpe Ratio measure of reward for risk attributable to these investments (ie. the expected return in excess of the risk free rate divided by its standard deviation) takes the value

$$SR = \frac{\rho}{\sqrt{1-\rho^2}}$$

for a single bet, and it increases with the square root of the number of independent bets. This result is less well known than it deserves to be. A derivation of it is therefore provided in Appendix A. This principle means that by calculating the magnitudes of the variance that we can explain and also the total variance we are exposed we can quite easily estimate the profit and risk associated with different numbers of transactions under different values of the Hurst exponent.

We have computed some theoretical results for the profitability of mean-variance optimal trading strategies in a fractional Brownian motion market. (The precise integrals they were calculated from are given in Appendix B). Table 1 gives theoretical estimates of the Sharpe ratio available per trade from various values of the Hurst exponent, ie. the average profit divided by its standard deviation. As increasing numbers of trades are completed, this ratio will increase with the square root of the number of transactions. The table also shows the number of transactions required to obtain a Sharpe ratio in excess of three. Given that the distributions involved are Normal, this seems an appropriate criterion for the existence of essentially riskless arbitrage. The table gives values for two separate cases, first where the entire history of the series can be observed, and second where the length of history observed corresponds to 50 trading intervals. It is worth noting that truncating the observation period at this point makes very little difference to the Sharpe Ratio available. It has the biggest effect in the case of high *H* values (and positive serial correlation), where the kernel continues to increase without bound as we go further back in time.

Table 1

Sharpe Ratios and Numbers of Transactions Needed for Arbitrage

H Value	Sharpe Ratio Entire History	Sharpe Ratio Finite History	Number of Transactions Entire History	Number of Transactions Finite History
0.25	0.445	0445	46	46
0.30	0.354	0.354	72	72
0.35	0.265	0.265	129	129
0.40	0.177	0.177	287	288
0.45	0.089	0.089	1128	1132
0.50	0.000	0.000	n.a.	n.a.
0.55	0.093	0.092	1048	1059
0.60	0.191	0.189	247	252
0.65	0.298	0.293	102	105
0.70	0.418	0.405	52	55
0.75	0.558	0.527	29	33

We conclude that for a fractional Brownian motion market with a Hurst exponent sensibly different from 1/2, riskless arbitrage can be accomplished by means of at most a few hundred transactions. Unless we really believe that a financial market is grossly inefficient, fractional Brownian motion is a most inappropriate model indeed.

Note that since the market we are dealing with is assumed to be fractal (and has a self similarity property with respect to scaling time), the time interval between transactions is entirely irrelevant. This is interesting in that it indicates that in a fractal market transactions costs are not an obstacle to successful arbitrage. They may simply mean that we can't afford to trade very often, and have to spread the same number of transactions over a longer time period. Even so, provided we are prepared to be patient, we can still eventually become incredibly rich with almost certainty. We can do this as long as we can still observe a sufficiently long time period to make our forecasts and as long as we don't insist on completing the arbitrage before we die!

5. Conclusions

This short paper has explored the nature of the arbitrage possibilities within a fractional Brownian motion market. We have indicated why the pricing function (or risk neutral probabilities) are degenerate and how they lead to arbitrage possibilities. With finite numbers of transactions risk has to be borne, but the strategies lead to riskless profits as increasing numbers of transactions are performed, depending on the extent to which the value of the Hurst exponent deviates from 1/2. We conclude that fractional Brownian motion is a quite inappropriate model for a financial market unless the market really is grossly inefficient.

References

Cutland, N J, P E Kopp, and W Willinger, 1993, "Stock Price Returns and the Joseph Effect: Fractional Version of the Black-Scholes Model", Mathematics Research Reports, Volume VI, No.12, University of Hull.

Hull, J, 1993, *Options, Futures and other Derivative Securities*, Prentice-Hall (2nd Edition).

Hurst, H E, R P Black, and Y M Sinaika, 1965, *Long Term Storage in Reservoirs: An Experimental Study*, Constable, London.

Lo, A W and J Wang, 1993, "Implementing Option Pricing Models when Asset Returns are Predictable", Working Paper, Sloan School of Management, MIT.

Mandelbrot, B B and J W Van Ness, 1968, "Fractal Brownian Motions, Fractional Noises and Applications", *SIAM Review*, 10, 422-437.

Peters, E E, 1994, *Fractal Market Analysis*, Wiley, New York.

Peters, E E, 1991, *Chaos and Order in the Capital Markets*, Wiley, New York.

Rambaldi, S and O Pinazza, 1994, An accurate fractional Brownian motion generator", *Physica* A 208, 21-30.

Rogers, L C G, 1995, "Arbitrage with fractional Brownian Motion", Working Paper, February 1995, University of Bath..

Treynor, J and F Black, 1973, "How to Use Analysts' Forecasts", *Journal of Business*.

Appendix A

How the Sharpe Ratio Depends on the Correlation of Forecasts

In a mean-variance framework, the optimal investment x_t at time t in a risky prospect is

$k\frac{\mu_t}{\sigma_t^2}$, where

k is a constant,

μ_t is the expected (ie. forecast) excess return, and

σ_t^2 is the variance of that return.

For the forecasts to be unbiased, the excess return r_t satisfies the equation

$r_t = \mu_t + \varepsilon_t$, where

$Var[\mu_t] = \rho^2 Var[r_t]$,

$Var[\varepsilon_t] = (1-\rho^2) Var[r_t]$ and we assume

$E[\mu_t] = 0$ as well as $E[\varepsilon_t] = 0$.

We are interested in the mean and variance of the portfolio return $x_t\, r_t$.
Over sufficient time periods t,

$$E[x_t\, r_t] = k\, E\left[\frac{\mu_t^2}{\sigma_t^2}\right] = k\frac{\rho^2}{1-\rho^2},$$

$$Var[x_t\, r_t] = E[\sigma_t^2\, x_t^2] = k^2\, E\left[\frac{\mu_t^2}{\sigma_t^2}\right] = k^2\frac{\rho^2}{1-\rho^2},$$

so the Sharpe Ratio

$$SR = \frac{E[x_t\, r_t]}{Var[x_t\, r_t]} = \sqrt{\frac{\rho^2}{1-\rho^2}}.$$

This result seems to have been established first by Treynor and Black (1973).

Appendix B

Explained and Unexplained Variances Under Fractional Brownian Motion

We assume that we are at time zero and making forecasts ahead for $t = \tau$. We consider our forecasting ability when we observe either series going back T units in time, to $t = -T$, or observing it backwards forever, to $t = -$ infinity. We shall see later that because of the fractal nature of the process, we can choose any time interval for the period ahead that we are forecasting and it is only the ratio of τ/T which matters. We can therefore scale time so that $\tau = 1$.

The correlation coefficients on which the figures in Table 1 are based, are therefore calculated from the following three variances:

V_{res} : the residual variance due to shocks arising between $t = 0$ and $t = \tau$.
V_{ex}: the variance in $t = \tau$ prices attributable to shocks between $t = -T$ and $t = 0$.
V_{un}: the variance attributable to shocks in unobservable history, prior to $t = -T$.

The Sharpe Ratios are obtained from these as the square root of the ratio of explained variance over unexplained variance as:

$SR_1 = \sqrt{\dfrac{V_{ex} + V_{un}}{V_{res}}}$, for observing the infinite history, and

$SR_2 = \sqrt{\dfrac{V_{ex}}{V_{res} + V_{un}}}$, for the other case.

Finally, the three variances are found as the following integrals:

$$V_{res} = \int_0^{\tau} (\tau - s)^{2\alpha}\, ds = \frac{\tau^{2H}}{2H}.$$
$$V_{ex} = \int_0^{T} \left[(s+\tau)^{\alpha} - s^{\alpha} \right]^2 ds,$$
$$V_{un} = \int_T^{\infty} \left[(s+\tau)^{\alpha} - s^{\alpha} \right]^2 ds, \quad \text{and whe}$$
$$\alpha = H - \tfrac{1}{2}.$$

The last two integrals were evaluated numerically, taking particular care to bound their values above and below in the cases where singularities arise.

From these definitions it is straight forward to show that SR_1 is a function of H and τ/T, and that SR_2 is a function of H alone.

Figure 1: Fractional Brownian Motions

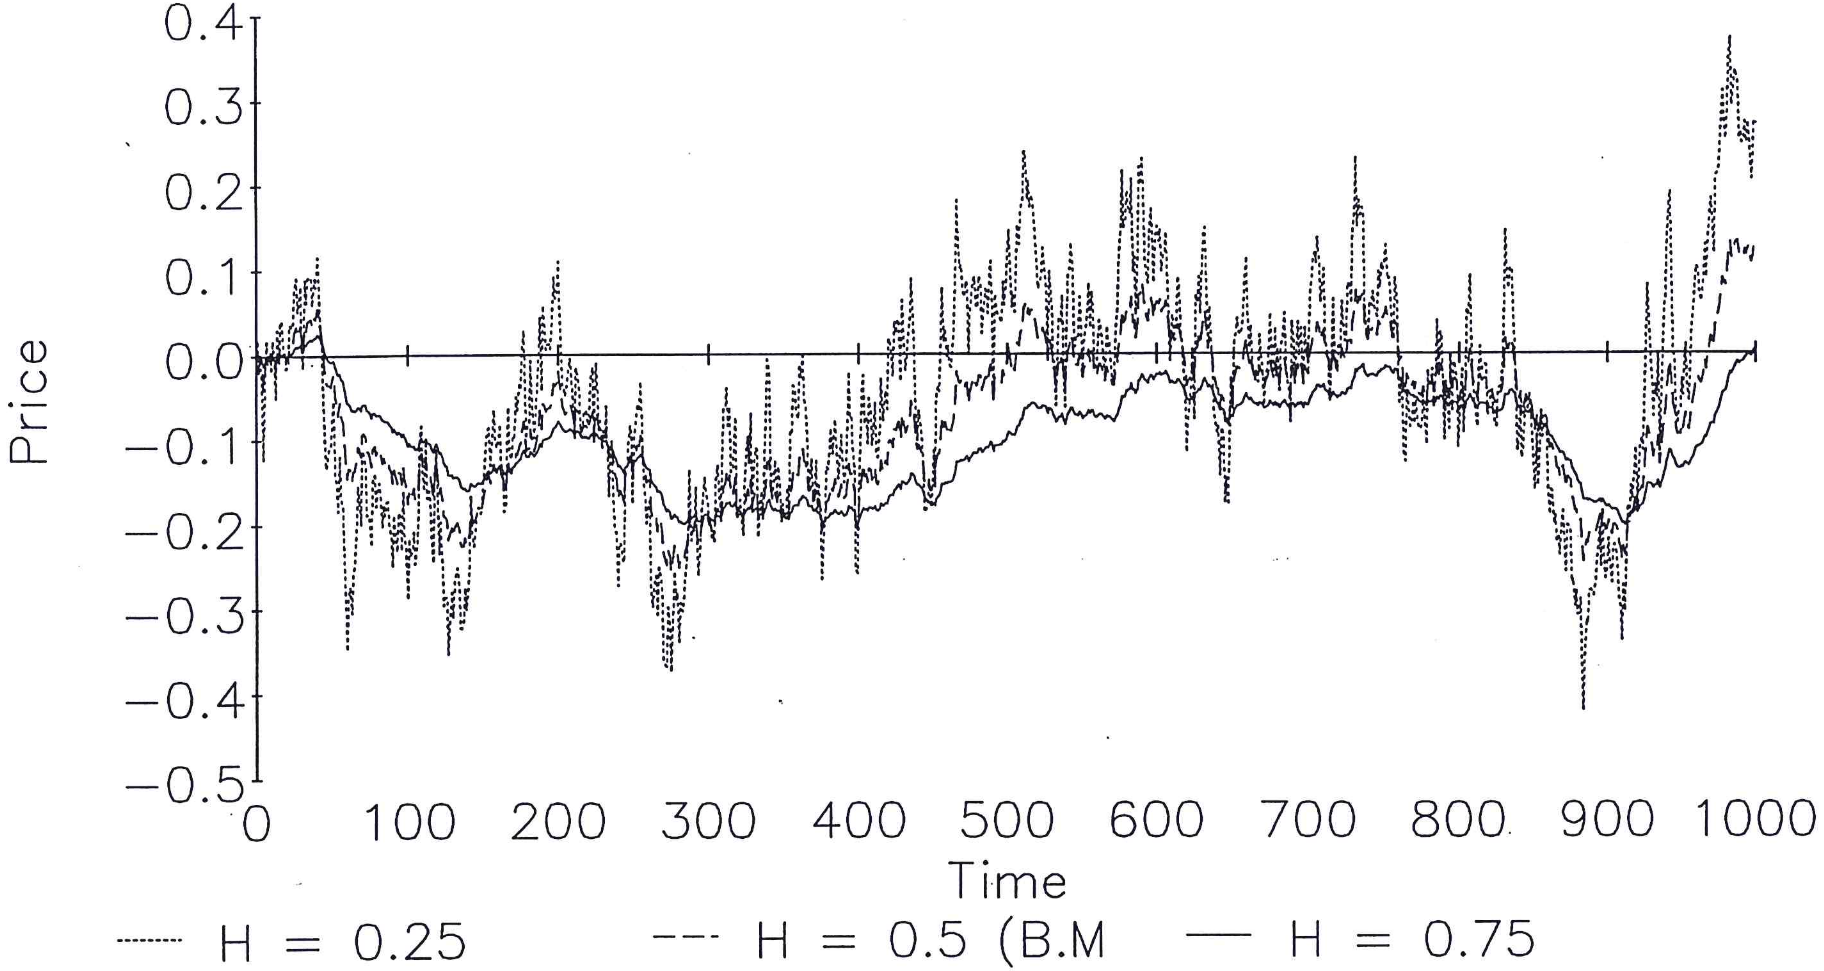

Arbitrage in a Fractal Brownian Motion Market

Stewart Hodges

Financial Options Research Centre
University of Warwick

First Draft: October 1994
This version: May 1995

Support from FORC Corporate Members, and from INIMS Cambridge is gratefully acknowledged. So too are helpful comments on an earlier version by Michael Dempster, Michael Selby, Chris Rogers, FORC colleagues, and participants at a seminar at the University of Essex.

Arbitrage in a Fractal Brownian Motion Market

Fractal Brownian Motion has been suggested as
a model for financial markets.

There are even papers which propose
how to price options on fBm assets.

What sort of arbitrage opportunities would exist?

How long would it take to be sure of a profit?

Structure of the Paper

Key properties of Fractal Brownian Motion.

How to generate fBm.

What kind of pricing operator/risk neutral probabilities would we expect to find in such a market?

How would we try to trade?
Positive vs negative serial correlation cases.

Rewards for trading in various markets.

2. Fractal Brownian Motion

Key properties:

all conditional distributions are Normal,

the standard deviation increases as t^H
where $0<H<1$ is the Hurst exponent.

($H = 1/2$ is the usual standard Brownian Motion.)

Self similarity with respect to scaling time.

Hurst *et al* (1965) found exponents different from 1/2
for flows in rivers, temperatures, thickness of tree rings etc.

Empirical Estimates from Peters

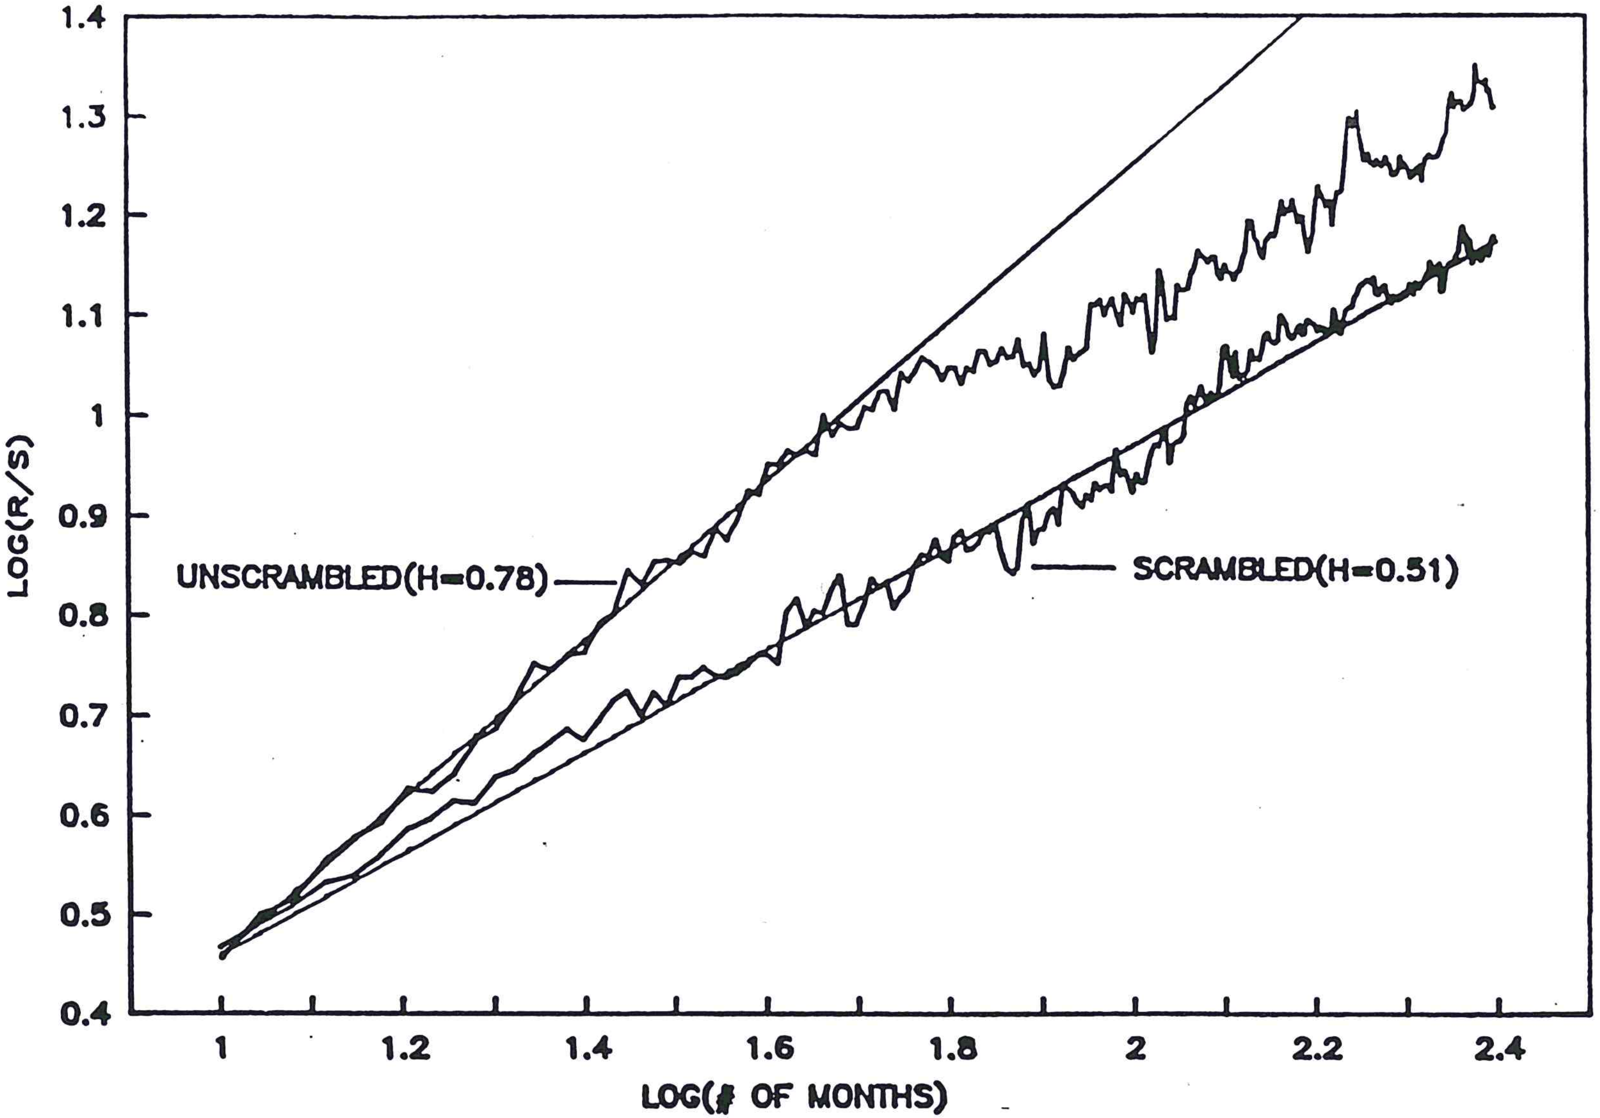

FIGURE 8.3 Scrambling test: S&P 500 monthly returns, January 1950–July 1988. Unscrambled H = 0.78; scrambled H = 0.51.

Generation of Fractional Brownian Motion

Fractional Brownian Motion may be better defined in terms of the power kernel which genererates it. (See Mandelbrot and Van Ness (1968))

$$F(t) = \int_{-\infty}^{t} (t-s)^{H-1/2}\, dB(s)$$

This integrates to give variance increasing as t^{2H}.

3. Pricing Functions: Informal Arguments

If one year returns have standard deviation σ, then over time interval Δt, the standard deviation of return is $\sigma (\Delta t)^{H}$.

"Trending case" $H > 1/2$, these are smaller than the annual standard deviation suggests.

These short increments would usually imply an annualised standard deviation of $\sigma (\Delta t)^{H-1/2}$, which tends to zero as Δt tends to zero.

Implications for Option Pricing

In this case it looks as though we should price options as though the volatility of the asset is zero!

Conversely, for $H < 1/2$ (negative serial correlation), it tends to infinity as Δt tends to zero.
Here we should price options as though the volatility is infinite!

Risk Neutral Probabilities

Risk neutral probabilities can be obtained as the second derivative of call option prices with respect to their strike price.

For the two cases we have :

$H > 1/2$, volatility = 0: $C = \text{Max}\{0, S_0 - \text{PV}(X)\}$.

which gives a unit probability at $S_T = \text{FV}(S_0)$, with zero elsewhere.

$H < 1/2$, volatility = ∞: $C = S,\ P = \text{PV}(X)$,

gives probability mass only at $S_T = 0$ and at infinity.

4. Arbitrages

Option based arbitrages:

$H > 1/2$: Hedge straddle as "through the cap"
$H < 1/2$: Delta hedge eg. butterfly spread.

Arbitrage based on portfolio selection:

Forecast instantaneous return,
Long or short position proportional to E[r].

Attainable μ/σ (Sharpe ratio) is

$$\frac{\rho}{\sqrt{1-\rho^2}}$$ and grows with $\sqrt{N}$.

How the Sharpe Ratio Depends on the Correlation of Forecasts

In a mean-variance framework, the optimal investment x_t at time t in a risky prospect is

$k\frac{\mu_t}{\sigma_t^2}$, where

k is a constant,

μ_t is the expected (ie. forecast) excess return, and

σ_t^2 is the variance of that return.

For the forecasts to be unbiased, the excess return r_t satisfies the equation

$r_t = \mu_t + \varepsilon_t$, where

$Var[\mu_t] = \rho^2 \, Var[r_t]$,

$Var[\varepsilon_t] = (1-\rho^2) Var[r_t]$ and we assume

$E[\mu_t] = 0$ as well as $E[\varepsilon_t] = 0$.

We are interested in the mean and variance of the portfolio return $x_t r_t$.
Over sufficient time periods t,

$$E[x_t r_t] = k\, E\left[\frac{\mu_t^2}{\sigma_t^2}\right] = k\frac{\rho^2}{1-\rho^2},$$

$$Var[x_t r_t] = E[\sigma_t^2 x_t^2] = k^2\, E\left[\frac{\mu_t^2}{\sigma_t^2}\right] = k^2\frac{\rho^2}{1-\rho^2},$$

so the Sharpe Ratio

$$SR = \frac{E[x_t r_t]}{Var[x_t r_t]} = \sqrt{\frac{\rho^2}{1-\rho^2}}.$$

This result seems to have been established first by Treynor and Black (1973).

Table 1
Sharpe Ratios and Transactions Needed for Arbitrage

H Value	Sharpe Ratio		Number of Transactions	
	Entire History	Finite History	Entire History	Finite History
0.25	0.445	0.445	46	46
0.30	0.354	0.354	72	72
0.35	0.265	0.265	129	129
0.40	0.177	0.177	287	288
0.45	0.089	0.089	1128	1132
0.50	0.000	0.000	n.a.	n.a.
0.55	0.093	0.092	1048	1059
0.60	0.191	0.189	247	252
0.65	0.298	0.293	102	105
0.70	0.418	0.405	52	55
0.75	0.558	0.527	29	33

5. Conclusions

We have explored the nature of arbitrage opportunities in a Fractal Brownian Motion Market.

With finite numbers of transactions, risk has to be borne, but the strategies lead to risk free profits after only a few hundred transactions.

Fractal Brownian Motion cannot be used a model for an efficient financial market.

It is particularly inappropriate as a starting point for option pricing.