BANK OF ENGLAND

Financial Stability Paper No. 23 – August 2013

The Fractal Market Hypothesis and its implications for the stability of financial markets

Nicola Anderson and Joseph Noss

BANK OF ENGLAND

Financial Stability Paper No. 23 – August 2013

The Fractal Market Hypothesis and its implications for the stability of financial markets

Nicola Anderson and Joseph Noss

The authors are grateful to Evangelos Benos, James Benford, Oliver Burrows, Lewis Webber and Anne Wetherilt for their comments and suggestions. Any errors remain their own.

nicola.anderson@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

joseph.noss@bankofengland.co.uk
Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

The views expressed in this paper are those of the authors, and are not necessarily those of the Bank of England or the Bank for International Settlements. This paper was finalised on 19 August 2013.

© Bank of England 2013
ISSN 1754–4262

Contents

The Fractal Market Hypothesis and its implications for the stability of financial markets

Nicola Anderson and Joseph Noss

Time series of financial market prices appear to exhibit fractal properties: that is, under magnification, their pattern becomes increasingly complex, and seems to repeat itself, with a pattern that is qualitatively similar to that of the overall structure.

This paper examines why and how these fractal properties might arise, and considers their implications for understanding the causes of financial (in)stability. It offers a quantitative model of investor behaviour and price formation that seeks to account for fractal properties of market prices. It conjectures that the dynamic of market prices — in particular its self-similarity — might be caused by the interactions of agents with different investment horizons and differing interpretations of information. This structure appears to be associated with a special sort of stability that can be disrupted, causing prices to crash, if the normal interaction of these agents breaks down.

Introduction

Consider a tree (**Figure 1**). From a distance, it resembles the smooth symmetrical shape of a child's drawing or simple logo. But zoom in closer, and a richer, rougher pattern of branches is revealed. Underlying the solid overall structure of the tree is a sort of increasing complexity at finer levels of resolution. And within the network of branches successive generations become progressively smaller, yet their shape resembles that of the previous generation, and the overall tree. The structure of the tree exhibits 'self-similarity'.

Figure 1 A fractal tree

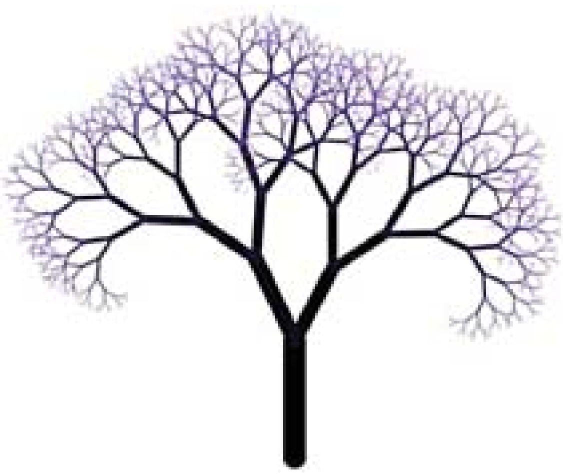

Such self-similarity — the tendency of an object to be similar to parts of itself — is a defining property of fractals, a structure found across the natural world. Such structures appear simple and smooth from a distance. But, under magnification, their pattern seems to become increasingly complex, and seems to repeat itself, with the pattern qualitatively replicating that of the overall structure. It turns out this is a property that can also be observed in time series of some financial market prices.

This paper examines why and how these fractal properties might arise, and considers their implications for understanding the causes of financial (in)stability. In particular, it revisits the 'Fractal Market Hypothesis' (FMH), a theory of market behaviour proposed by Peters (1991), which seeks to account for fractal properties of market prices. It conjectures that the dynamic of market prices — in particular its self-similarity — might be caused by the interactions of agents with different time horizons and differing interpretations of information. This structure appears to be associated with a special sort of stability that can be disrupted, causing prices to crash, if the normal interaction of these agents breaks down. Self-similarity also implies a greater degree of persistence than we might expect to observe under alternative theories of financial market pricing.

This line of research is still in its embryonic stages, but has potentially important implications for the regulation of financial markets, including of their major participants.

At the most basic level, integral to any assessment of risks to financial stability is also the ability of investors to model and manage their own risks, which relies on their understanding of the statistical properties of financial market prices. It turns out that the FMH is able to replicate many of the distributional characteristics of financial market prices where more standard models fail. While the FMH is insufficiently developed to provide a methodology upon which to base risk management practices or a rigorous assessment of the financial stability, there are nevertheless some important lessons that both investors and regulators should bear in mind, not least in developing a better understanding of the structure of the markets and their risk.

The key insight of this paper is that fractal structures typically imply some sense of stability. If we can understand precisely what market mechanics generate the fractal structure of prices, we can understand some of the defining elements of stable markets and even formulate policies to foster those elements — relating to, for example, the incentives for different agents to interact at different time horizons. Equally, we may be able to identify why financial markets often give rise to periods of illiquidity and suggest policy recommendations designed to counter them. Finally, a better understanding of how market prices behave, including what can cause their persistent deviation from fundamentals, is essential to guiding potential policies that may be targeted towards leaning against exuberance in financial markets. It may also help explain, and guide regulation designed to mitigate, the occurrence of 'flash crashes' linked to the activities of high frequency traders (see Haldane (2011)).

The contributions of this paper are broadly threefold. First, it considers how the fractal properties of market prices — and in particular the FMH that seeks to account for them — might inform an understanding of the structure of financial markets, and the causes of their stability and instability. In doing so, it reaches far beyond the traditional paradigm of market behaviour, the Efficient Market Hypothesis (EMH). Second — and, to the best of the authors' knowledge, uniquely — it offers a quantitative model of investor behaviour and price formation that formalises the qualitative conjecture of the FMH, and demonstrates its ability to match some of the observed features of financial market prices, including their fractal properties. Third, it draws from this analysis some implications that are relevant to a number of ongoing debates regarding the regulation of financial markets and of their major participants.

This text proceeds as follows. The next section recaps some salient features of financial market prices and their statistical distribution, including their self-similarity over time. It reviews the reasons why the EMH — a sometime dominant paradigm to describe market behaviour — fails to provide an adequate characterisation of the distribution of market prices. Section 2 introduces an alternative framework of markets and investor behaviour, the FMH, which better matches some of the observed characteristics we observe in prices. Section 3 provides a simple model that seeks to provide a quantitative formulation of the FMH and shows how it is able to replicate some of the observed fractal characteristics of prices. The implications for the persistence of financial market prices are examined in Section 4, while Section 5 offers some tentative policy implications. A final section concludes.

1 The characteristics of financial market prices and the limitations of the Efficient Market Hypothesis

The statistical properties of financial market prices — and, in particular, the distribution of their changes — is one of the most basic, yet important, properties of markets. To investors, it is the statistical distribution of price fluctuations that determines the return on their investments, and the risk they incur in making them. To a regulator or central bank, the distribution of prices has implications for their assessment of risks to individual institutions and to financial stability. Most obviously, the behaviour of the 'tail' of the distribution of price changes (the likelihood of a large adverse market move under circumstances of market stress) can pose threats to financial stability.

The Gaussian paradigm

Intuitive (if naïve) arguments based on the Central Limit Theorem (CLT) (Rice (1995)) suggest that prices are log-normally distributed. If $p(t)$ is the price at time t, define the return over an interval τ as:

$$r_{\tau}(t) = \log p(t+\tau) - \log p(t) \qquad (1)$$

If τ is divided into a number of smaller subintervals, then the total return $r_{\tau}(t)$ is — by definition — the sum of the returns over each of these smaller intervals. And if the log price changes in each subinterval are independent and identically distributed (iid) — and under the (seemingly benign) assumption that its variance is finite — then the CLT provides that the probability distribution of log prices over the sum of all intervals, should converge to the normal, for a large number of subintervals.

Under the normal distribution, the probability of observing an event of a given severity — that is, an outcome that lies a certain distance from the distribution's mean — decays smoothly as that distance or severity increases.(1) More precisely, the probability density of observing an outcome distance x from the mean, scales *exponentially* with x; that is:

$$P(X = x) \textit{ is proportional to } e^{\frac{-x^2}{2a}} \qquad (2)$$

where a is a constant. Under this exponential scaling, a single parameter, a — the distribution's variance, describes how its probability density changes as x (the outcome whose probability is being measured) increases. That is:

$$\frac{dp}{P(x)dx} \textit{ is proportional to } a. \qquad (3)$$

So, under the Gaussian distribution, a single parameter connects behaviour on different magnitudes — that is, the probability of observing outcomes different distances from the mean. The distribution's tails scale smoothly and are exactly described by its variance.

The Efficient Market Hypothesis

The assumption that price changes over small intervals are independently and identically distributed is consistent with the notion that markets are 'efficient'. Under the EMH, the current price of a security reflects the entirety of all available public information on the underlying source of its return (eg the profit outlook for a firm or the default characteristics of a pool of mortgages).(2) Unexpected changes in its price occur independently over time, and arise only due to the announcement of unanticipated information that affects its value. Consequently, the only remaining uncertainty — and, indeed, 'unknown' left to be modelled — is the statistical distribution of future price changes in response to information.

The EMH leaves us with a simple and intuitive theory. Markowitz (1952) further showed how, under assumed normality, standard deviation can be used as a measure of risk, and the covariance of returns could be used to explain precisely how diversification (grouping stocks whose performance are less than perfectly correlated) reduces their aggregate risk (the standard deviation of the returns of the portfolio). These are concepts that are still used widely today by firms in their risk management and by regulators in their evaluation of models such as Value-at-Risk.

Fat tails, (in)finite variance and stochastic volatility

But the ability of the Gaussian distribution to capture the likelihood of rare events has been extensively criticised.(3) A wealth of evidence suggests that the normal distribution provides a systematic underestimate of the actual probability of observing an extreme outcome, or large movement in price. **Charts 1** and **2** compare the distribution of returns of the

(1) The history and properties of the Gaussian distribution, along with the consequences of its embedded assumption for finance, is given by Haldane and Nelson (2012).
(2) Technically this is the 'weak form' of the EMH.
(3) A survey is provided by Haldane and Nelson (2012).

Dow Jones industrials stock index from 1896, over weekly and annual horizons, compared to the normal distribution. Empirical distributions, at least at these short horizons, are characterised by a high 'peak' around their mean, and 'fatter tails' than that of the normal.(1) This casts doubt on the key assumptions underlying the Gaussian paradigm.

Consider, for example, the assumption that returns are identically and independently distributed (iid). This is clearly open to empirical contention — the standard deviation of the distribution of price changes varies over time, a phenomenon commonly referred to as 'clustered volatility'. While the autocorrelation of log-returns (that is the degree to which a return reflects its previous value), is generally very small on time scales longer than a day (see Farmer (1999)), the same is not true for volatility. Volatility on successive days is positively correlated, and these correlations remain positive for weeks or months. **Chart 3** shows the volatility of the returns of the Dow Jones over rolling intervals of a year. This phenomenon of clustered volatility is consistent with the fat tails in the resulting distribution of returns.(2)

The assumption of *finite* variance — critical to the application of the CLT — also fails to hold. **Chart 4** shows the *sample* variance of daily returns of the Dow Jones index, plotted over expanding samples, between 1896 and the present day. The resulting series is clearly unsettled. Rather than converging to a single number — ie some well-defined *population* (or 'finite') variance — it jumps with some regularity. No single parameter adequately describes the second moment of the distribution as would be the case for a distribution with finite variance (as given by equation **(2)**).

The observed distribution of prices therefore fails to match that of the Gaussian distribution and that predicted by standard interpretations of the EMH. In particular it appears that changes in price *are not* independently and identically distributed over different periods, and their behaviour over events of different scales of extremity (distance from the mean) cannot be adequately described by a single characteristic scale, their variance.

Self-similarity, and a mixture of global determinism and local randomness

Another less discussed property of financial market prices is the similarity of their behaviour when viewed over multiple times scales.

From a 'distance' — that is, viewed on a long-time scale — series of financial market prices appear smooth, almost deterministic. **Chart 5** plots the daily prices of the Dow Jones from 1896. Apart for an abrupt reduction visible in the end of the series, and around the early 1930s, the overall pattern appears roughly to conform to an exponential pattern (and even the recent fall is small when compared to the growth over the century preceding it). Viewed at this level, any particular portion of the series appears relatively smooth. But zero in on a shorter-time scale and two phenomena appear:

First, under magnification, the series yields increasing complexity or 'roughness'. The smoothness of any particular portion of the series gives way to a rougher pattern at increasing orders of magnitude (**Chart 6**).

Second, under this magnification, the pattern of the prices series seems qualitatively similar to that at lower levels of resolution. **Charts 7** and **8** show the Dow Jones index between 1988 and 1997,(3) but with the price series sampled at different frequencies. Thus, for example, the annual series shows the movements in price observed over the course of a year; whereas, the 'four-yearly' series shows the movements in price over the course of four years (note that each series is constructed as an average over the relevant time period and has been rescaled to begin at 100 and have the same mean value). **Chart 9** shows the series overlaid on each other.

While the two series clearly differ, there are some marked qualitative similarities. Notice, for example, the generally high level of prices in the first half of the period — up until the summer for the annual series, and up until the second year of the four-yearly series. This is followed by a slump around two thirds of the way through — around autumn and towards the end of the third year respectively — and then a flourish towards the end of both series. For comparison, **Chart 10** shows a similar phenomenon for the FTSE All-Share index.

Price series would therefore appear to exhibit self-similarity — an increasing complexity under magnification (so that series that are smooth viewed over a long time period are less smooth over shorter time periods) and an invariance of structure to the scale on which they are viewed (so that patterns are similar despite representing movements over the course of different length time periods). Theorists have yet to agree on an exact mathematical definition of fractals,(4) but there is a broad consensus that this self-similarity — whether it be exact, or (as in this case) just a qualitative similarity — is their defining characteristic.

(1) This fat tailed property can be formally measured by using the statistic kurtosis; see Farmer (1999).
(2) For example, it is possible to show that the sum of normally distributed variables with different variances exhibits higher kurtosis than that of the normal distribution.
(3) The Dow Jones index is used here as its price history was available to the authors over the longest period.
(4) Even Mandelbrot (1982), who originally defined fractals based on their topological dimension, later rejected this definition. Others define it in more general terms as being a structure whose 'parts are related to its whole' (Peters (1991)).

Charts 1–6 The Dow Jones industrials index

Chart 1 Distributions of weekly/annual returns

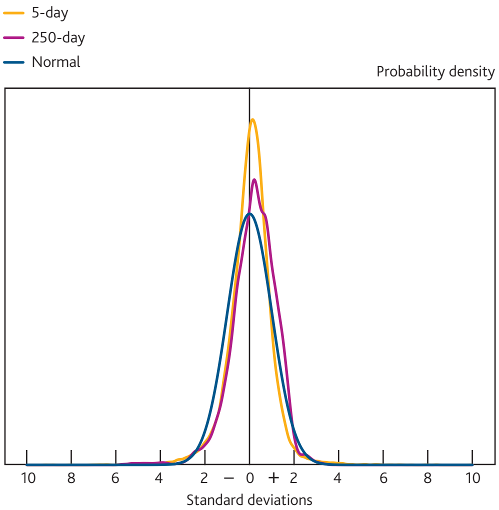

Source: Global Financial Data.

Chart 2 Distributions of weekly/annual returns (rescaled)

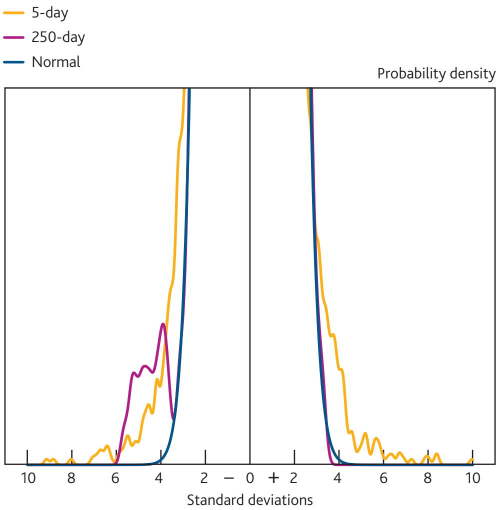

Source: Global Financial Data.

Chart 3 Volatility (over 250-day rolling windows)

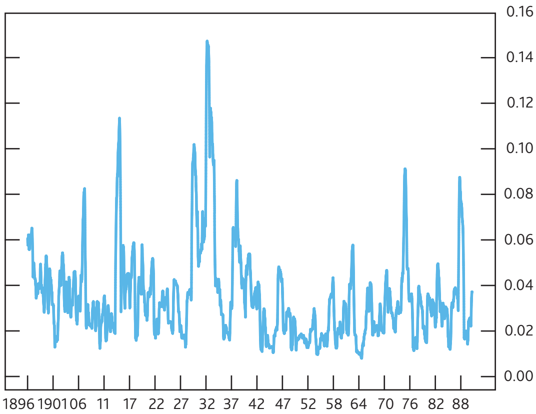

Source: Global Financial Data.

Chart 4 Variance of daily returns taken of expanding windows of daily observations (from 1 January 1896)

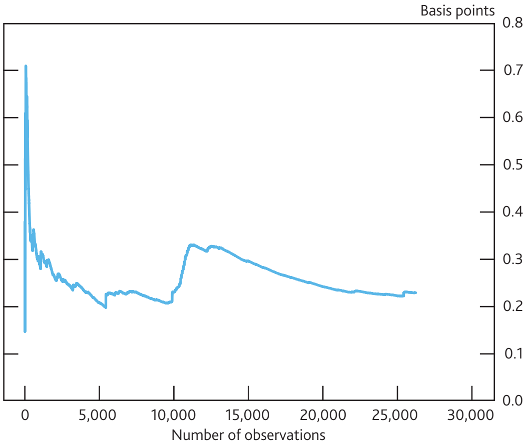

Source: Global Financial Data.

Chart 5 Prices: 1896–1991

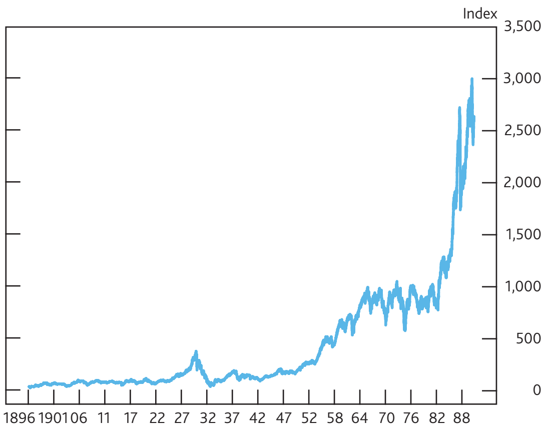

Source: Global Financial Data.

Chart 6 Prices: 1962–66

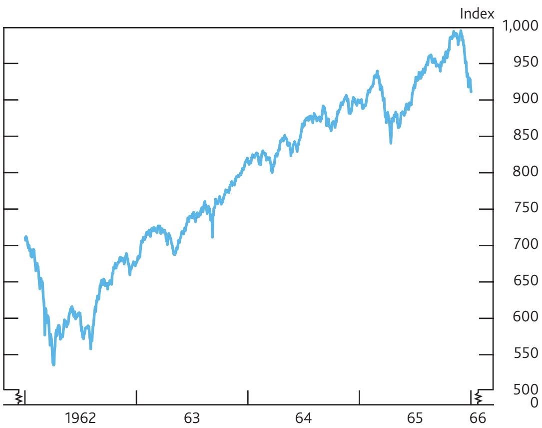

Source: Global Financial Data.

Charts 7–10 Financial market price series viewed over different horizons

Chart 7 Annual observations of the Dow Jones industrials index (1988–97)

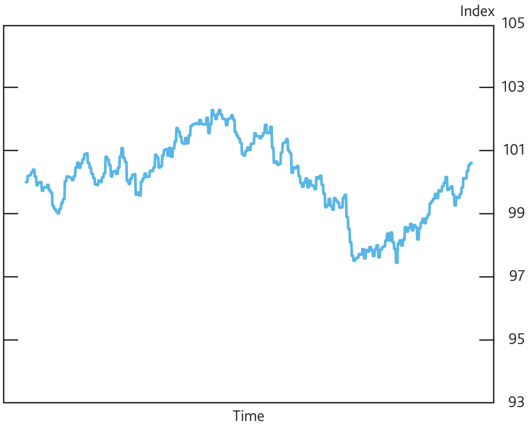

Source: Global Financial Data.

Chart 8 Four-yearly observations of the Dow Jones industrials index (1988–97)

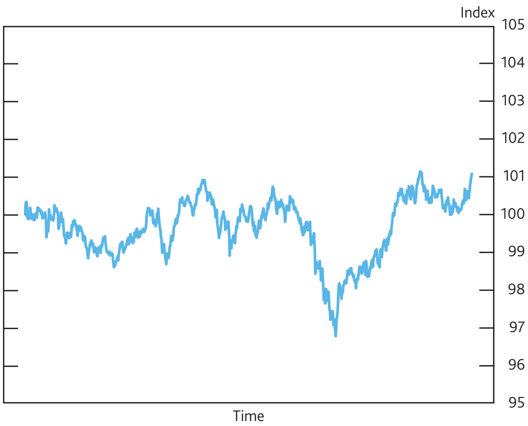

Source: Global Financial Data.

Chart 9 Annual/four-yearly observations of the Dow Jones industrials index (1988–97)

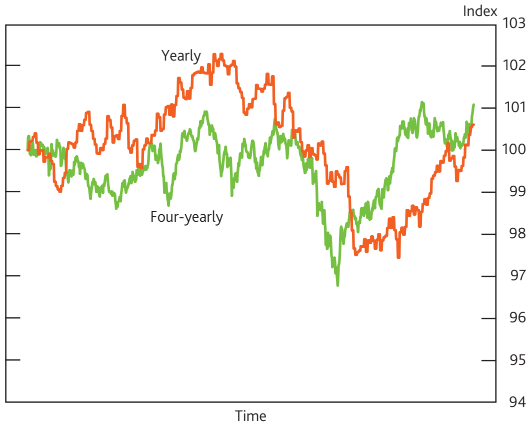

Source: Global Financial Data.

Chart 10 Annual/four-yearly observations of the FTSE All-Share index (1988–97) (rebased)

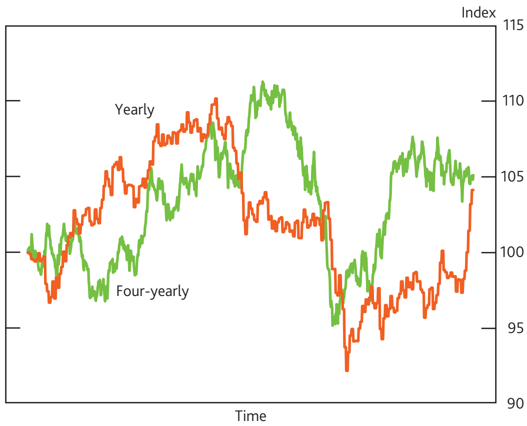

Source: Global Financial Data.

It is such insights that make fractals a promising tool with which to analyse capital markets and motivates a search for a 'better theory' of market behaviour. Desirable properties of such a theory are already clear. Besides matching the distribution of market prices (with their fat tails and time-varying volatility), it should also connect behaviours across a number of time scales; that is, it should address the 'self-similarity' property we witness in financial market prices. If we can better understand the structure of financial markets, we can improve the way in which risk is managed and regulated, and guide policies that may serve either to enhance or detract from the stability of those markets.

2 The Fractal Market Hypothesis

This section sets out the main features of the FMH, first proposed by Peters (1991). This aims to address the short comings of the EMH, and in particular understand why self-similarity exists in financial market prices.

The role of liquidity and information

At the heart of the FMH lie two elements missing from its Efficient Market predecessor: that is, a role for liquidity and (relatedly) the impact of information. Liquidity in this context is defined as 'market liquidity': that is, the relative ease with which an investor is able to buy/sell a security without their act of buying/selling having a substantial effect on its price. Loosely speaking, liquidity is generated whenever investors trade with each other.(1) For this to be the case, it is posited that two investors must have different views on a security's value. This can arise either because:

a. One investor has information on its value to which another investor does not have access (or to which he or she has yet to gain access). Abstracting from the possibility of 'insider

(1) But, importantly, market liquidity is distinct from trading volume. It is possible for a market to endure a crash — that is a sharp imbalance in supply and demand, leading to a large discontinuous fall in price — with low liquidity, but high trading volume (given the pressure to sell).

trading', where an investor trades on information not yet revealed to the market, this may arise whenever some investors receive information that affects the fundamental value of a security before others. For example, so-called 'high-frequency traders' are able to pay a premium to observe market prices — information relevant to valuing a security — at a higher frequency than the wider investment community (see Haldane (2011)).

b. Or two investors receive information simultaneously, but place a different weight on its importance, given the differing time horizons over which they invest.(1) For example, in an equity market, there are investors with short-term horizons such as hedge funds. For these investors, the daily distribution of returns may have a large bearing on their buying/selling behaviour. In particular, a fall in prices commensurate with, say, three standard deviations of daily returns, might be seen as highly adverse.

   But to a longer-term investor — for example, 'real money' investors such as pension funds — these daily high and lows may be less important. Such investors might judge their performance against the longer-term distribution of prices. As **Charts 1** and **2** show, this distribution is of a similar shape, but, under it, a three standard deviation move is commensurate with a far more extreme movement in price than under that of daily returns.

Under the FMH, market liquidity — that is investors' willingness to trade with each other — is caused by information having a different effect on different investors, either because they obtain it at different times, or because some property of their own preferences means they interpret information differently. Viewed in this context, it seems unlikely that market prices are 'entirely random', with independent increments, as predicted under the EMH. But neither should we expect them to be entirely deterministic. Rather, it would seem they lie somewhere in between. Intuitively, prices reflect a combination of short-term technical information, used by short-term traders, and long-term fundamental information of importance to long-term investors. This particular mix of local randomness over short-time scales, and global determinism over longer time scales, is a defining characteristic of 'self-similarity', as observed in fractal structures.

The Fractal Market Hypothesis: a 'special' sort of (financial) stability

Importantly, particularly from the perspective of a regulator or macroprudential policymaker, differences in the interpretation of information may also imply price stability, at least under standard market conditions. When a day-trader experiences a price move that they judge to be of a severity that causes them to sell, an investor with a longer horizon can step in and buy from them. This longer-term investor is willing and able to do so, because, with their longer-term horizon, the day-trader's *n*-sigma event is not unusual (unless, of course, *n* is *very* large): judged by the *longer-term* distribution of returns, it is closer to the mean. As long as another investor has a longer-trading horizon than the investor experiencing a crisis, and as long as the *n*-sigma event does not convey any longer-term negative news on fundamentals, the market will stabilise.

In this way, financial markets can be considered as embodying a 'special sort' of stability. Indeed, self-similar fractal structures seem to be favoured by nature as a way of ensuring the stability of systems in the natural world. In the case of the tree, for example, it is determined globally that each branch will divide to yield two or more branches, defining its overall structure. But, importantly, branches of successive generations are not identical. This ensures that if one branch of the tree were malformed, there would be other branches to compensate. The global determinism of the structure and its local randomness combine to mean the overall structure of the whole tree is not threatened by the presence of malformed branches. The fractal structure seems to embody a certain 'tolerance to error' that guarantees the stability of the system.(2)

Intuitively, it follows that financial markets can become prone to instability when this fractal structure is broken. The obvious case for a breakage to occur is when investors with a longer horizon either stop participating in the market, or become short-term investors.

Reasons for such an eventuality can be mapped to the two differing interpretation of investor horizon (a and b) described above.

- The original explanation offered by Peters (1991) in his exposition of the FMH, matches (b); in this case, an exogenous event occurs that causes, short-term investors to sell. The consequent fall in prices causes long-term investors **to doubt the validity of the longer-term information** on which they base their behaviour. This could, for example, be because it causes them to be uncertain as to their longer-term view of the longer-term payoffs of a security, formed through some measure of long-term economic fundamentals.

  Peters (1991) suggests that such an event could, for example, be some 'natural' disaster such as the events of 11 September 2001. Because the consequences of such an

(1) Intuitively, investors may have different horizons according to the structure of their funding; eg investors funding on the basis of repo borrowing may have shorter horizons than investors employing long-term savings.
(2) Goldberger and West (1987) first formalised this insight.

event for the long-term prospects for the economy were so uncertain, it caused long-term investors to 'lose faith' in their longer-term view of economic fundamentals, and subsequently either stop trading or themselves become short-term traders and begin trading on overwhelmingly negative short-term market dynamics.

- A more subtle and endogenous explanation, explored by Haldane (2011), is that, because long-term investors view the market less regularly than short-term players, they come to doubt the veracity of the information given by prices in a market where some investors view prices more frequently than others.

In either case, the removal of longer-term investors from the market, or reduction in their horizon, causes liquidity to evaporate as there is no longer heterogeneity of investor valuations to facilitate trading. This insight has potentially important policy implications, which are explored in further detail in Section 5. Before doing so, it is worth examining a quantitative model of the FMH, which should further our intuition as to the role of different horizon investors in determining the path of financial prices, including susceptibilities to loss of liquidity and associated crashes — both of which can contribute to a loss of financial stability.

3 A simple quantitative model of the Fractal Market Hypothesis

This section proposes a simple theoretical model of investor behaviour that seeks to replicate that conjectured under the FMH. This model offers a quantitative counterpart to the existing descriptions of the FMH, which, until now, have only been set out in qualitative terms, as above, and in Peters (1991). It seeks to capture how the different interpretation of information by investors of different horizons affects the stability of the resulting price series. In doing so, and through the subsequent effect that investors' buying/selling behaviour have on prices, it is able replicate many of the non-Gaussian properties of markets (including fat tails, stochastic volatility and self-similarity), described in Section 1.

A simple quantitative model

The model is devised in the spirit of interacting agents alluded to in the description of the FMH in Section 2. At its core lie two agents with different investment horizons. Both invest based on the information gleaned from the change in price they witness in the previous period, but the distribution on which they condition their behaviour varies depending on their horizon.

For the sake of simplicity, agents of each investment horizon have the same 'shape' of demand function. Under normal circumstances, they are 'mean reversionists': that is, they observe the price change in the previous period, assess that against their expectation (reflecting a 'fundamental' view) and seek to react in a way that profits from any deviation: buying if this relative price movement is negative (exhibiting positive demand), or selling if it is positive (exhibiting negative demand). Their demand is therefore the simple product of the extremity of the periods' price movement relative to its expected change and some (negative) constant that represents their readiness to buy/sell.

But if investors witness a decline in price in the previous period that is particularly extreme, their demand becomes negative and they sell in some fixed (large) quantity. This is compatible with the more 'traditional' explanation of market crashes postulated by Shiller (1987), whereby investor 'stop-loss' limits force investors to sell when prices decline by a certain magnitude, in order to limit their individual loss, and so exacerbate market-wide price falls.

Algebraically, the response to prices of the demand of short and long-term mean reversionist agents can be represented by functions h_S and h_L respectively, where:

$$h_i(x) = \begin{cases} -a_i x \text{ if } x > t_i \\ -d_i \text{ otherwise} \end{cases} \tag{4}$$

where x is the return witnessed in the previous period relative to that explained by fundamentals; t_i is the 'threshold amount' which, if the return is above, causes the investor to buy (sell) if the price change is negative (positive); d_i is the amount by which they 'force sell' otherwise; and a_i is the 'aggression' of their mean reversion (the level of buying/selling they exhibit in response to a given level of price change $x > t_i$). This function is illustrated in **Chart 11**.

Chart 11 An illustration of the function governing the demand of the mean-reversionist investors

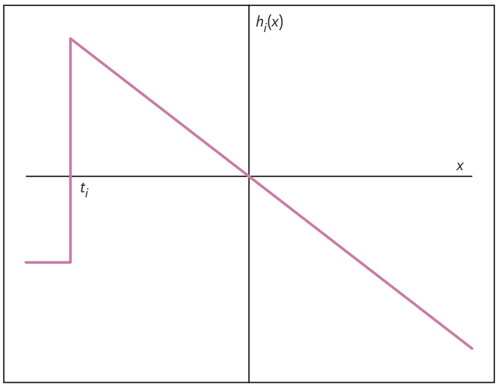

Crucially, and in keeping with the FMH, the distribution of past returns against which investors condition their behaviour — ie how extreme they regard a given change in price in the last period to be — varies depending on their horizon. In normal times, it requires a larger upward (downward) price deviation from fundamentals, in absolute terms, to provoke as strong a selling (buying) reaction from the long term investor compared to the short-term investor. This is because the variance of the long-term investors' price distribution, to which they compare any change, is far larger. And, in stressed times, the magnitude of the negative price movement that causes a long-term investor to force sell (ie their *n*-sigma selling event) is far larger than that of the short-term investor. This means that $t_L < t_S$.

Finally, there is also a 'momentum trader', who follows bull runs; that is they exert positive buying pressure when the price change in the previous period, is positive, but do not trade when it is negative (for example, due to a short-selling constraint).

$$M(x) = \begin{cases} a_M x & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

This trader serves as a natural 'antidote' to the occasional forced selling pressure exerted by the mean reversion investors, ensuring that prices do not, in the long term, fall to zero.

If markets always cleared at some 'fundamental value' of assets, these demand functions would obviously offset to have no impact on the resulting evolution of prices. But in practice, there is nothing to guarantee such market clearing. Nevertheless, we assume that prices are broadly anchored to a fundamental path, so that the three demand functions defined by equations **(4)** and **(5)** serve to create fluctuations around this path. Representing this in terms of returns, R_t:

$$R_t = r + \varepsilon_t + h_s(x) + h_l(x) + M(x) \text{ with } x = R_{t-1} - r. \quad (6)$$

Here *r* denotes the fundamental expected return on the asset, ε_t is a shock assumed to be normally distributed with mean zero variance one; together, these terms capture the drift rate in market prices alongside the arrival of hitherto unknown information (Fama (1965)). r is set to 0.05, σ to 0.05; these values are not crucial to the narrative, but they accord roughly with the long-term 'risk-free' rate of interest rates on one-year UK government bonds,[(1)] and the volatility of the Dow Jones series used in Section 1.

The model is calibrated using the methodology described in Annex 1.

Empirical results

The ability of this model and its interacting investors to match the market dynamic postulated by the FMH is clear from the resulting price time series. **Chart 12** shows the prices corresponding to one series of random innovations $\{\varepsilon_t\}$.[(2)] **Chart 13** shows a portion of same series but 'zoomed in' over a shorter-time scale. Different colours denote different combinations of investor behaviour:

- In both charts, the periods where the blue line is without coloured overlay are those in which neither investor is forced selling. That is, the price movement in the previous period has triggered neither investor's forced selling constraint.

- The light orange areas in **Chart 13** are those where the short-term investor is short selling; but, in keeping with the market dynamic postulated by the FMH, this selling pressure is compensated for immediately in the next period by the buying of the longer-term investor, who sees the resulting price fall as a 'buying opportunity' (see the qualitative description in Section 2).

- The periods in magenta in both charts are those of 'market instability'. Here both investors force sell simultaneously.

As suggested by the FMH, in standard market conditions, the interaction of short and long-term investors ensures market stability: long-term investors buy as a result of short-term investors' forced selling. But, on some rare occasions, short-term investors' forced selling in one period is — by chance — accompanied by a downward movement in the arrival of random noise (ε) in the next that is sufficient to cause long-term investors to force-sell in the period that follows. Periods of instability ensue, as this leads to a self-perpetuating spiral of forced selling by both short and long-term investors (the magenta dots in **Chart 12**). These are only reversed when information happens to arrive that is sufficient to 'reverse' the destabilising effect of long-term investors force selling, and produce a price movement that is sufficiently non-negative to avoid continued forced selling in the next period.[(3)]

The resulting distribution of price returns is able to match several of the empirical features of the Dow Jones examined in Section 1. In particular, it reflects the respect in which these observed series — and their associated distribution of price returns — deviate from that under the normal distribution and the associated naïve form of the EMH. **Chart 14** shows a quartile-quartile plot of modelled daily returns, comparing returns (on the y-axis) observed at each percentile of their distribution (on the x-axis). For observations below the median (ie the 50th percentile), the *n*th percentile of modelled prices is below that of the normal distribution, indicating that its left-hand tail is 'fat'. The combination of stability under normal market conditions with occasional bouts of instability postulated by the FMH, also combine to capture the stochastic volatility of prices witnessed in

(1) In practice, of course, we might expect fundamental returns also to include a risk premium component.
(2) For this series, t_S and t_L were calibrated to values of 0.8% and 2.5% respectively.
(3) In reality, it is also some possible that at least some (perhaps longer-term) investors are likely to place some weight on a measure of deviation of price from some measure of economic fundamentals (somehow defined); so that the (random) arrival of such positive information is not the only reason such adverse cycles of forced selling are broken.

Charts 12–16 Results of a theoretical model of the Fractal Market Hypothesis

Chart 12 A time series of modelled prices

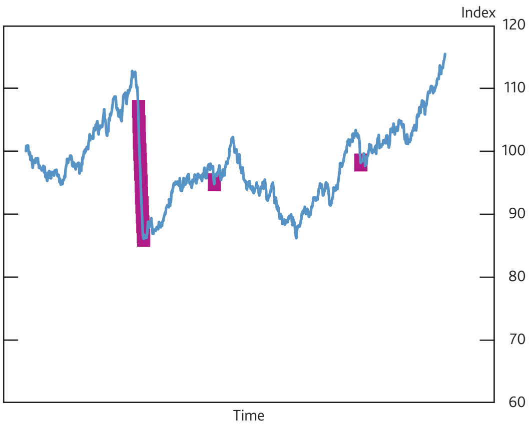

Chart 13 A time series of modelled prices (over a shorter time scale)

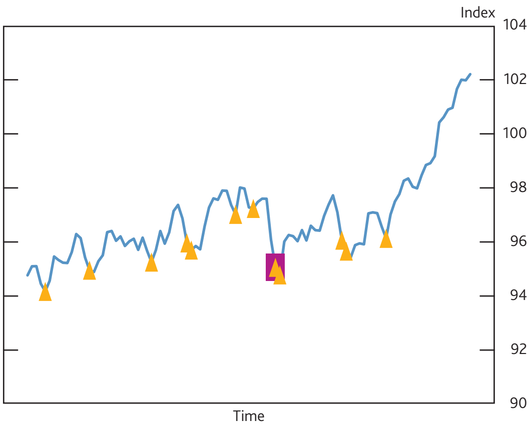

Chart 14 A percentile plot of the resulting distribution of daily returns

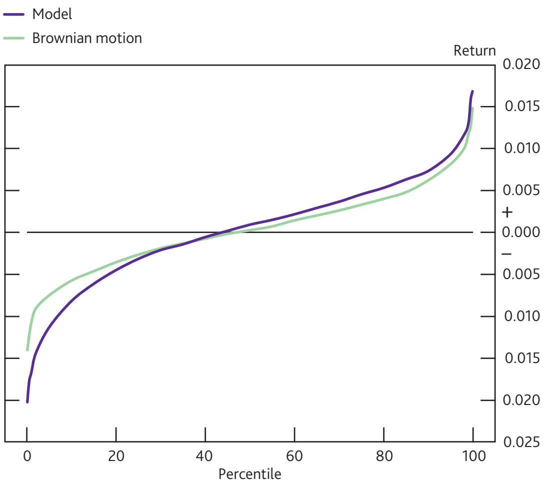

Chart 15 Volatility (over 250 period rolling windows)

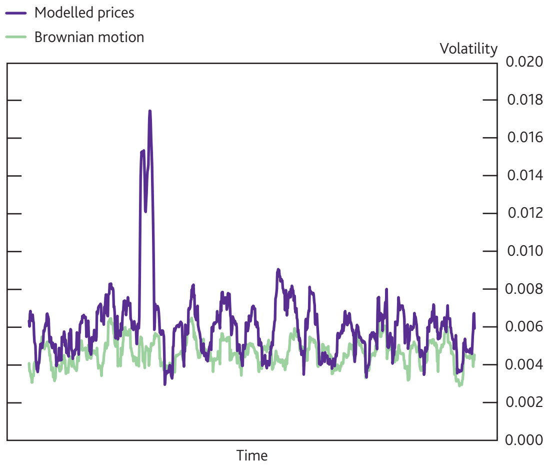

Chart 16 The price series of the theoretical model viewed over different horizons

Section 1. **Chart 15** compares the volatility of modelled prices to those of the Brownian motion. Modelled prices displays a series of 'spikes' in excess of that of the latter, in keeping with the pattern seen in **Chart 3**. Finally, the model produces a price series that are qualitatively similar when viewed over a quarterly compared to an annual horizon (**Chart 16**).

A restricted version of the model without the long-term investor

To illustrate the posited stabilising effect of investor interaction, consider a version of the model above but without the long-term investor; ie one with only the short term and momentum traders. The model is calibrated to the same information arrival process as before, but now only the volume with which the short-term investor force sells (their '*d* parameter') need be calibrated, subject to the requirement that they force sell when they witness a price move three standard deviations below the mean of the distribution of price returns at their horizon.

This is still a model with investor interaction (with the short term and momentum traders) and many of the properties of the full model are consequently preserved; for example, the resulting price series show signs of self-similarity (**Chart 20**). But we now see clear signs of a less stable market. The resulting price series is shown in **Chart 17**, where the magenta dots indicate the forced selling of the short-term trader only. Now that the long-term trader is absent, periods of instability — that is significant impact of forced selling by the remaining short-term trader — are much more frequent. There is no other trader to see extreme price movements as a buying opportunity, and support markets with their countervailing buying pressure. The price distribution also has 'fatter' tails than found under both the normal distribution, and under the full model, reflecting the more frequent occurrence of extreme price movements (**Chart 18**). And volatility (**Chart 19**) is again greater than that of the normal distribution — but now more so than under the full model. Notice, however, that the abrupt spike in volatility, experienced under the full model is now absent (compare with **Chart 15**).

Chart 17 A time series of modelled prices (restricted model)

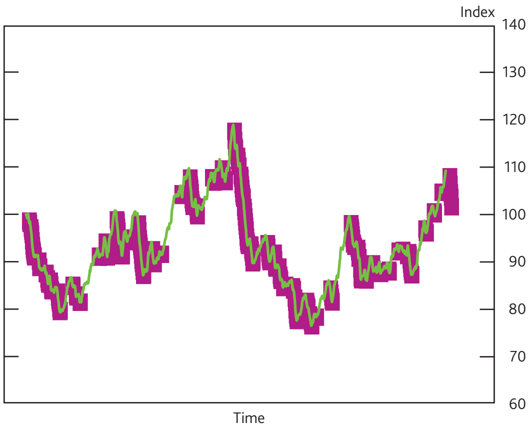

Chart 18 A percentile plot of the resulting distribution of daily returns (restricted model)

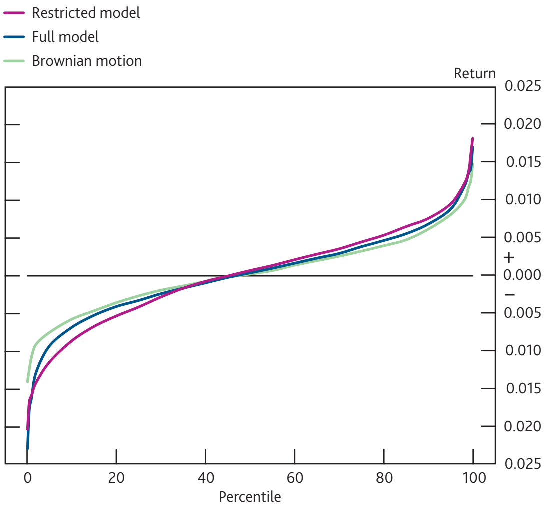

Chart 19 Volatility (over 250 period rolling windows) (restricted model)

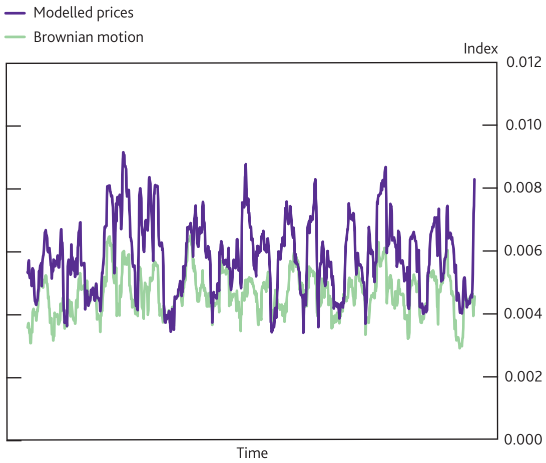

Chart 20 The price series of the restricted model viewed over different horizons

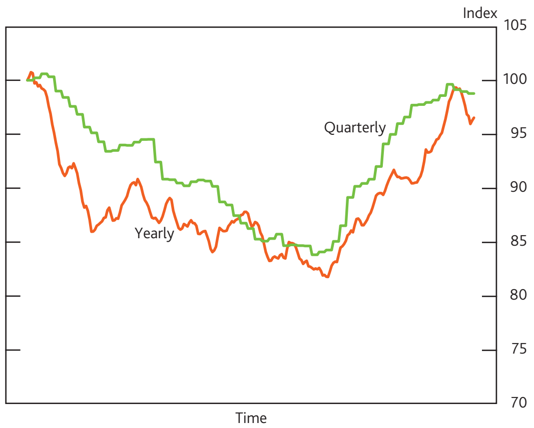

A taxonomy for price formation and the provision of liquidity

This last observation suggests that the level of heterogeneity in investor strategy seems to control a trade-off in markets — between the frequency of their periods of instability, and their severity. The full model, with its larger range of heterogeneous investors, gives rise to a price process that — under normal market conditions — is more stable: its interacting agents of multiple horizons provide consistent flows of liquidity. But it also has occasional 'spikes' of high volatility, illiquidity and discontinuous price movements when prices crash. On the other hand, the restricted model has a price series that is fundamentally less stable in that it contains frequent periods of illiquid forced selling by the short-term investor. Its volatility is therefore higher than that of the restricted model, but is less volatile. These findings are summarised in **Table A**.

Table A A taxonomy of investor interaction, and resulting distribution of returns

Model	Distribution of returns	Volatility of returns
Full model of the FMH with long and short-horizon investors.	Fat tailed.	Stochastic; low, with a spike.
Restricted investor interaction: only a short-horizon investor (and a momentum trader).	Fat tailed, and more so than full model.	Stochastic; higher without a spike.
No interacting investors.	Normally distributed.	Non-stochastic.

4 'Measuring' fractals using rescaled range analysis: persistence and fat tails

The previous section highlighted the role of long-term investors in ensuring periods of market stability. In this section, we now focus more precisely on the property of self-similarity. **Charts 16** and **20** compare the behaviour of the series produced by the full and restricted models over 'quarterly' and 'annual' horizons. As with the actual market price series observed in **Charts 7–10**, the two bear at least a qualitative similarity over the two time windows, the hallmark of fractal series discussed in Section 2. As it appears in both models, this would appear to be a property of the investor interaction required to generate market liquidity under the FMH; in the case of our model, between the short-term reversionist trader and the momentum trader.

This self-similarity is interesting because it implies a level of persistence in financial market prices that we would not expect to see if returns were independently and identically distributed, as posited under the EMH. This in turn implies that prices can travel further than they might otherwise — that we might see prolonged periods during which they deviate from some concept of their fundamentals.

But the measurement of self-similar structures and their degree of 'travel' — that is the quantification of how they fill space — is not straightforward. Take the example of a coastline. Measured with an imaginary mile-long ruler, the coast of the United Kingdom would appear a certain length. But measured with a shorter, meter-long ruler, the resulting length would grow. The same phenomenon would be true were 'rulers' of different lengths used to measure the lengths of the time series of financial prices, for example, those in **Charts 5** and **6**. This changing measurement arises typically when a regular line segment, such as a ruler, is used to approximate the length of nested self-similar structures. There is no 'true' or 'characteristic' value for their length; instead, it grows as a power function (plus a constant) of the precision of the ruler used to measure its length.(1) This lack of characteristic scale is another frequently observed property of fractals.

This difficulty in measuring self-similar structures complicates the measurement of variation in fractal time series, such as those of financial market prices. Under an assumption of normality, the incidence of market movements of a given severity, compared to the mean, decreases exponentially with variance (equation **(2)**) — its characteristic scale. But, as we have seen, the distribution of financial market prices yields no such Gaussian simplicity. An alternative measure for the 'spread' of fractal market prices needs to be 'distribution free' — and go beyond the restrictive parametric assumptions of the Gaussian.

Rescaled range analysis

A candidate solution is provided by the work of Hurst (1951). Hurst postulated that the variation of fractal time series can be described through a power law relationship, where the range increases in proportion to a power of the horizon over which the time series are viewed. Suppose we define Y_k as the sum of k small increments of a demeaned time series that extends up to n increments. We define the adjusted range to be the maximum minus the minimum of the series $\{Y_1,\ldots,Y_n\}$, or:

$$R_n = \max\left(Y_k\right) - \min\left(Y_k\right),\ 1 < k < n. \qquad (7)$$

This adjusted range can be thought of as the 'distance' the series travels over n increments of time. In the case of Y being a 'random walk' — ie a series whose increments are Gaussian — it is well known that this range increases with the standard deviation of the series multiplied by the square root of n.(2) Hurst's contribution was to generalise this relationship to:

(1) This was first observed by Mandelbrot (1967).
(2) Note that, absent the two mean-revisionists and momentum trader, the description in equation **(6)** gives rise to a price that is Gaussian distributed with mean $(r - 0.5\sigma^2)dt$ and variance $\sigma^2 dt$.

$$(R/S)_n = cn^H. \qquad (8)$$

Here S is the standard deviation for the same n observations of the demeaned time series and c is a (positive) constant. Rescaling the range of a series, by dividing by its standard deviation, allows for the measurement of time series that have no finite variance. In other words, it allows for the measurement of fractals. This method makes no assumption as to the underlying distribution of the series' increments. It measures only how it scales through time, as measure by H, known colloquially as the *Hurst exponent*.

Rescaled-range: a taxonomy of the persistence/jaggedness of fractal structures

The H statistic gives a measure of the degree of persistence in the time series that applies regardless of the time scale over which it is viewed. The greater range scale witnessed in a series with H>0.5 means that a high value in one period is likely to be followed by a higher value in a later period (the reverse applies if H<0.5). Importantly, and unlike with a series that exhibits standard serial correlation, this persistence applies regardless of the time scale on which it is viewed. All six-month periods influence all following six-month periods; all weekly periods influence subsequent weekly periods. This 'long range' form of dependence has been observed in a number of self-similar processes.(1)

Empirical results for the Dow Jones industrial index and for the models of Section 3

Measuring Hurst coefficients of actual time series involves performing the calculation in **(8)** across a range of values for n. Technical details are contained in Annex 2. **Chart 21** shows the range scales for 20-day returns (increments of 20 days) on the Dow Jones index, between 1896 and the present day. The blue line shows the ratio of the range scale, R/S to sqrt(n), for window sizes, w (so that the number of observations is $n \times w$) between 10 and 150. Also plotted is the expected value of this

Chart 21 Range-scale analysis; 20-day returns of the Dow Jones industrial average

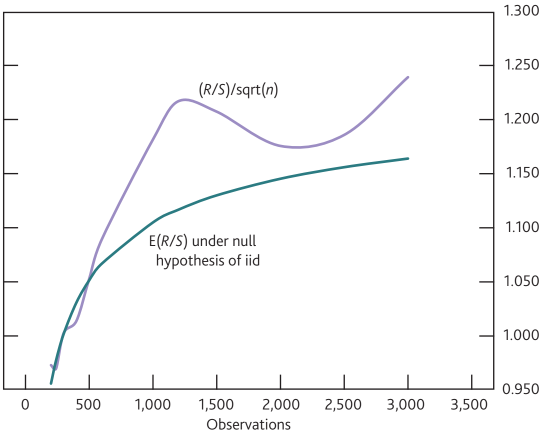

Sources: Global Financial Data and Bank calculations.

ratio, under the null hypothesis that the system is an independent process (ie with H=0.5).(2) If the series exhibits persistence (H>0.5), then the ratio between the two lines will increase with the number of observations. If it is a random walk (H=0.5), the ratio will remain constant. The ratio clearly stops growing at 1,040 trading days (corresponding to a value of w of 52). This suggests the presence of persistence (ie values of H in excess of 0.5) over observation windows shorter than this.

Table B shows the results of estimating H for values of w between 10 and 50 over these 20-day observations. The calculation yields H=0.72, compared to an expected value, E(H) (under the null, Gaussian hypothesis) of 0.62. The standard deviation of this estimate of E(H) is 0.025. Thus the estimated value for H is more than three standard deviations above its expected value — a highly significant result. From this analysis, the 20-day changes in the Dow Jones industrial index are characterised as a persistence process, with Hurst coefficient significantly in excess of that commensurate with a random walk, for periods of up to four years (1,040 trading days).

Table B Estimated Hurst coefficients

Series	Window size	Estimated value	Expected value under null hypothesis that H=0.5	Standard deviation of estimated value
Dow Jones (20-day returns)	10–50	0.72	0.62	0.025
Full model	10–50	0.66	0.53	0.032
Restricted model	10–50	0.73	0.53	0.032

Similar results are shown in **Table B** for both the full and restricted models of investor behaviour and price formation given in Section 3. Both give rise to Hurst coefficients (of 0.66 and 0.73, respectively) that are significantly larger than their expected values, suggesting that both series exhibit a degree of persistence similar to that of the Dow Jones series.

In summary, both the Dow Jones industrials index — and the price series produced by the models in the previous section — seem to be characterised by the sort of persistence present in Hurst processes. Put another way, financial markets show a systemic deviation from the dynamic predicted by the Gaussian distribution. In particular, they display persistence irrespective of the scale over which they are viewed: daily returns are correlated with future daily returns, monthly with monthly, and so on.

(1) Though, importantly, not all series exhibiting long-range dependence are fractal; Granger and Joyeux (1980) provide a summary.
(2) The expected value of the Hurst coefficient, and its standard deviation, are calculated using the methodology in Anis and Lloyd (1976).

But, at least with the Dow Jones index of equities, this 'long-term' memory is not infinite. **Chart 21** suggests that this memory effect seems to dissipate around 1,040 days (or four years). At frequencies longer than this, the Hurst coefficient returns to a value closer to 0.5. Prices behave more like those of a random walk. At these longer frequencies, prices may be tied more to the economic cycle — following random innovations in the macroeconomy — and less to the effects of investor interaction.

One way of checking the robustness of this result, and that it is not just an artefact of the time window over which the data is sampled, it is instructive to check it is independent of the time period over which returns are taken in the range-scale analysis. **Chart 22** shows the comparable estimates to **Chart 21** but using five-day returns. A break is clearly visible, and at around 1,040 trading days. This persistence for periods under four years therefore seems robust to the frequency with which the data is sampled.

Chart 22 Range-scale analysis; five-day returns of the Dow Jones industrial average

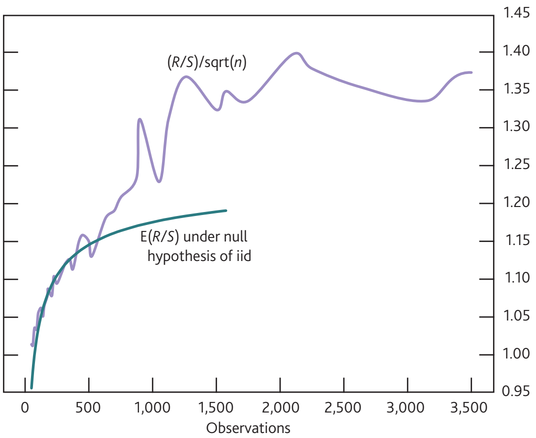

Sources: Global Financial Data and Bank calculations.

5 Possible policy implications

The model presented in this paper is purely illustrative, and is intended as a means of motivating the observed statistical distribution of financial market prices and its fractal properties. But it might shed some light on a range of important issues that regulators and policymakers could usefully bear in mind:

Managing and regulating market risk

Risk managers and regulators have long recognised that the distribution of changes in financial market prices is not necessarily Gaussian. Nevertheless, this fact is often still overlooked in certain key contexts, and there remain certain risk management systems that are still based on an underlying implicit assumption of normality. Moreover, attempts to overcome their limitations are purely statistical, and often consist of tweaking the assumptions of the normal distribution to help fit observed market data across different markets.(1)

In contrast, the model in this paper has demonstrated that — rather than being a mere statistical anomaly — the cause of the non-normality of market prices may be inherent in its underlying (fractal) structure, which may be determined by the nature of investor interaction — this will arguably differ from market to market. Moreover, different varieties of investor interaction and behaviour are likely to result in different dynamics. This suggests that, when examining the measurement and management of market risk, risk managers and regulators should consider the structure of the market, including the time horizons of its key investors and the manner in which liquidity is formed. In doing so, they should be more alert to the likelihood of sharp corrections and sudden losses in liquidity.

Unintended consequences of mark-to-market accounting

A key insight from the FMH is that, under standard market conditions, stability is likely to be enhanced when a market contains investors with different time horizons.

Recent reforms in accounting standards require banks to value their assets according to the price available for those assets in the market place. This is driven by a desire to ensure a firm's assets are valued at prices commensurate with those at which it would be possible to sell them in the market. Shin (2007) highlights the potential implications of such a regime, including greater balance sheet volatility, which may not always be justified by market 'fundamentals'. Haldane (2010), meanwhile, considers how, if the prices of an asset become misaligned, fair-value market accounting could lead their prices to 'over correct.' In the present context, the concern is that more frequent and volatile valuations may risk leading to a reduction of the investment horizon of certain sorts of investor.(2) And if the interaction of investors with a wide variety of investment horizons is, as the FMH postulates, inherent to the provision of market liquidity then this could be to the detriment of financial stability.

Enhancing the veracity of market information

A related concern is the potential for longer-term investors to leave the market suddenly — potentially causing significant price falls in otherwise relatively stable markets.

Under the FMH, one reason longer-term investors could leave the market is when they come to doubt the veracity of price information caused by the interaction of investors who viewed the market at a higher frequency than themselves (see

(1) See Noss (2010), for an example pertaining to the risk management of structured credit instruments.
(2) For example, Mayer (2001) concludes that quarterly reviews of pension fund managers encouraged them to focus on short-term performance.

Section 3). Haldane (2011) suggests that this accounted for the lack of market liquidity during the 2010 Flash Crash. With high-frequency traders operating at frequencies more minute than lower-frequency market makers were able to observe, the fractal characteristics of markets — including the risk of dispersion in future prices — emerged at shorter time scales. This dynamic aggravated the risk that lower frequency firms faced, and in the extreme meant they could not see the price at which they could trade. In other words, they lay at an acute information disadvantage in times of stress, unable to make markets without risking being arbitraged by higher-frequency traders. This may have caused them to withdraw from the market, reducing liquidity, and causing the crash.

The FMH is certainly not the only methodology available for motivating the observed statistical distribution of financial market prices. For example, it is possible to simply assert that the properties we observe reflect time variation in the arrival of fundamental information. But the key insight the FMH supports, that long-term investors are important for market stability, is an intuitive one and suggests three specific possible responses of public policymakers that could help reinforce the confidence of longer-term investors, and thereby bolster market stability:

First, market-making guidelines would seek to ensure a commitment by market makers to provide liquidity at all times. This would in effect ensure some heterogeneity of investor actions, and hence maintain investor interaction, and market liquidity — even during times of stress.[(1)] This would therefore reduce the most pernicious aspect of the FMH, by reducing the occurrence and severity of market crashes. Unsurprisingly, however, the implementation of market making guidelines is fraught with difficulty, not least around how to specify the size of the quotes at which such dedicated market makers should be required to transact (for further discussion see Benos and Wetherilt (2012)) and how all this interacts with regulatory changes, including capital requirements attached to market-making activity.

A second variety of policy response supported by the FMH is that of market 'circuit breakers'. Proposals to introduce such mechanisms date back a number of decades, including in response to the stock market crash of 1987, as outlined in the report of the so-called Brady Commission.[(2)] At their most basic, circuit breakers are simple rules under which all trading of a given product on a given exchange is halted if prices move too erratically, as judged by some pre-set criteria. Again the justification under the FMH is clear: such a break in trading would allow for the resolution of any informational asymmetries between investors who view the market at different frequencies. This would allow longer-term investors to regain confidence in market information and lessen the probability of their withdrawal from the market during times of stress.

A third type of policy proposal is that of minimum resting periods. These impose a minimum delay between the time at which a trade on a given exchange is submitted, and that at which it is executed. The justification, and its auspices under the FMH are clear: minimum resting periods serve as an *ex-ante* counterpart to market circuit breakers, that has the effect of 'slowing trading down' and maintaining longer-term investors' confidence in market information across all states of the world (rather than just those of liquidity draughts). But their drawbacks are also clear: minimum resting periods would increase the costs of transactions by high-frequency traders. Depriving the market of this class of (short-horizon) investor may have an adverse effect on liquidity, at least during normal times (Hasbrouck and Saar (2011)).

Minimum resting periods therefore involve a trade-off between lowering the risk of liquidity evaporating during times of stress, but at the same time lowering the average liquidity at other times. The FMH does, however, remain a useful prism within which to view this trade-off.

6 Conclusion

This paper examined the fractal properties of financial markets, considered how they might arise and their implications for financial stability. In particular, it revisited the FMH, under which the self-similarity of financial price series comes about due to the interaction of investors with different investment time horizons. It offered a quantitative model that formalised its predictions and was able to match some of the observed properties of financial market prices.

The FMH has potential implications for our understanding of financial markets, their dynamics, and the causes of their instability. In particular, and in contrast with more traditional paradigms of asset pricing, it highlights the role of market liquidity and heterogeneity of investors' interpretation of information as determinants of market stability. Fractal structures give rise to a sort of robustness, whereby, under normal market conditions, the differing interpretation of information by, and behaviour of, investors at different time horizons combines to ensure market liquidity and orderly price movements. But, this fractal structure also implies a certain type of fragility. Its breakage — that is the loss of long-term investors or a change in their behaviour (either becoming short horizon, or leaving the market) — can cause this liquidity to evaporate, producing panic selling and associated market crashes.

This investor interaction gives rise to persistence in the fractal structure, whereby its changes are correlated over different

(1) Venkataraman and Waisburg (2007) find evidence that dedicated market makers increase liquidity on equity markets.
(2) See US Department of the Treasury (1988).

time scales. The exact nature of this seems to depend on the exact nature of the investor behaviour and their degree of interaction. Its further investigation is, however, left as further work.

Finally, from a practical standpoint, the FMH clearly supports the crucial role of securities regulation in maintaining financial stability. The incentives and behaviour of different types of investor are highlighted as key elements in determining the stability of markets, both under normal conditions and during times of stress. Effective securities regulation is a necessary component to ensuring that — as far as possible — all types of investor are properly incentivised, or restricted, to exhibit behaviours that are in concert with a well-functioning financial system.

Annex 1
Calibration of the model in Section 3

Recall from Section 3 that the return in the *t*th period, relative R_t can be represented algebraically as:

$$R_t = r + \varepsilon_t + h_s(x) + h_l(x) + M(x) \textit{ with } x = R_{t-1} - r \qquad \textbf{(6)}$$

where ε_t is normally distributed with mean zero and variance one; r is set to 0.05, σ to 0.05.

The average price return resulting from this model **(6)** is fitted to that of the Dow Jones index examined in Section 1. This requires six parameters to be calibrated: the aggression parameters for the short horizon, long horizon and momentum traders (a_S, a_L and a_M respectively); thresholds at which the short and long-horizon investors force-sell (t_s and t_L); and a volume *d* that they force-sell that is common to both. To avoid the problem of over-identification, the aggression parameters of the short and long-horizon mean reverting investors, a_S and a_L, are set to 0.4. That of the momentum trader is set to 1.

The horizons of the short horizon and long-horizon traders are arbitrarily set to values of 1 and 10 respectively. In keeping with the spirit of the FMH, the thresholds at which the short and long-horizon investors force sell (t_s and t_L), are set to a fixed point on the distribution of price returns as *observed at their differing time horizons*: both investors force sell when they witness a price move three standard deviations below the mean of the distribution of returns at their horizon; that is

$$t_i = E\left(R_t\right) - 3 * \sqrt{\operatorname{var} R_t} \textit{ for } i = l,\ s. \qquad \textbf{(9)}$$

In summary, the calibration problem reduces to finding a volume of forced-selling, *d,* common to both investors, that best fits the average annual return of the price resulting from the model (given in **(6)**) to the 5% average return on the Dow Jones, subject to the investor behaviour specified in **(9)** being maintained. That is, it is required to find:

$$d = \arg\ \min \left| E\left(\frac{P_{t+250} - P_t}{P_t}\right) - 0.05 \right| \textit{subject to } (3). \qquad \textbf{(10)}$$

The complication in the procedure comes in how, for any random series of Gaussian disturbances $\{\varepsilon_t\}$, this distribution of prices is itself a function of the choice of parameters. This is solved by a two-step numerical procedure: for a given random series $\{\varepsilon\}$, an arbitrary d is selected. The corresponding t_s and t_L that satisfy **(9)** are then solved for numerically. The procedure is then repeated with a different value of *d*, until the minimisation above is achieved.

Under the above calibration, only the mean of the modelled returns is matched to that of observed returns. A more advanced approach might be to calibrate multiple moments of the distribution, though this would significantly add to the complexity of the procedure. It is, therefore, left as future work.

Annex 2
Estimating the Hurst exponent

The Hurst exponent H can be approximated by calculating range scales across a range of values of n. Taking logs of **(8)** it can be found easily through ordinary least squares regression as a coefficient estimate:

$$Log(R/S)_n = log(c) + H\,log(n). \tag{11}$$

The following is a step by step methodology for applying range-scale analysis to financial market price data (used in Section 4), as given in Peters (1991):

1. Begin with a time series of prices $\{P_t\}$, and convert this into a time series $\{X_t\}$ of N returns: $X_t = \log(P_t/P_{t-1})$.

2. Divide this time series (of length N) up into a series of A contiguous subperiods of length n, such that A*n = N. Label each subperiod l_a with a = 1, 2, 3,…, A. Label each element in l_a as N_k,a where k = 1, 2, 3, … ,n. For each subperiod, l_a, of length n, calculate the mean

   $$e_a = \frac{1}{n}\sum_{k=1}^{n} N_{k,a}.$$

3. Calculate the time series of accumulated departures from the mean for each subperiod, l_a, as:

   $$X_{k,a} = \sum_{i=1}^{k}\left(N_{i,a} - e_i\right)\ for\ k = 1,\ 2,\ 3,\ldots,n.$$

4. Define the range as the difference between the maximum and minimum value of $X_{k,a}$ within each subperiod l_a: R_{l_a} = max($X_{k,a}$) - min($X_{k,a}$) where 1 < k < n.

5. Define the sample standard deviation for each subperiod l_a as:

   $$S_{l_a} = \sqrt{\frac{1}{n}\sum_{k=1}^{n}(N_{k,a} - e_a^2)}.$$

6. Each range, R_{l_a} is now normalised by dividing by its corresponding S_{l_a}. Therefore the rescaled range for each subperiod l_a is equal to R_{l_a}/S_{l_a}. From step 2, this yields A contiguous subperiods of length n.

The average R/S value for length n is defined as:

$$(R/S)_n = \frac{1}{A}\sum_{a=1}^{n}\frac{R_{l_a}}{S_{l_a}}.$$

7. The length n is then increased until there are only two subperiods, ie $n = N/2$. A least squares regression is then performed with log(n) as the independent variable and log$(R/S)_n$ as the dependent variable. The slope of the equation is the estimate of the Hurst exponent, H.

References

Anis, A and Lloyd, E (1976), 'The expected value of the adjusted rescaled hurst range of independent normal summands', *Biometrika, No. 63*.

Benos, E and Wethrilt, A (2012), 'The role of designated market makers in the new trading landscape', *Bank of England Quarterly Bulletin*, Vol. 52, No. 4, pages 343–53.

Fama, E (1965), 'Random walks in stock market prices', *Financial Analysts Journal*, Vol. 21, No. 5, pages 55–59.

Farmer, J (1999), 'Physicists attempt to scale the ivory tower of finance', *Computing in Science and Engineering*, Vol. 1, No. 26.

Granger, C and Joyeux, R (1980), 'An introduction to long-memory time series models and fractional differencing', *Journal of Time Series Analysis*, Vol. 1, pages 15–30.

Goldberger, A and West, B (1987), 'Fractals in physiology and medicine', *Yale Journal of biology and medicine*, Vol. 60, No. 5, pages 421–35.

Haldane, A (2011), 'The race to zero', available at www.bankofengland.co.uk/publications/Documents/speeches/2011/speech509.pdf.

Haldane, A and Nelson, B (2012), 'Tails of the unexpected', available at www.bankofengland.co.uk/publications/Documents/speeches/2012/speech582.pdf.

Hasbrouck, J and Saar, G (2011), 'Low-latency trading', *Johnson School Research Paper Series*.

Hurst, H (1951), 'The long-term storage capacity of reservoirs', *Transactions of the American Society of civil engineers*, Vol. 116.

Markowitz, H (1952), 'Portfolio selection', *Journal of Finance*, Vol. 7.

Mandelbrot, B (1967), 'How long is the coast of Britain? Statistical self-similarity and fractional dimension', *Science*, Vol. 156, No. 3,775, pages 636–38.

Mandelbrot, B (1982), *The fractal geometry of nature*, W. H. Freeman & Co.

Mayer, C (2001), 'Institutional investment and private equity in the United Kingdom', a paper for a conference on '*Corporate Governance: Reassessing Ownership and Control*'.

Noss, J (2010), 'Modelling structured credit', *Bank of England Working Paper No. 407*.

Peters, E (1991), *Fractal Market Hypothesis*, Wiley Finance.

Rice, J (1995), *Mathematical Statistics and Data Analysis*, (Second edition) Duxbury Press.

Shiller, R (1987), 'Investor behavior in the October 1987 Stock Market Crash: survey evidence', *NBER Discussion Paper No. 2446*.

Shin, H (2007), 'Discussion of assessing the information content of mark-to-market accounting with mixed attributes: the case of cash flow hedges and market transparency and the accounting regime', *Journal of Accounting Research*, Vol. 45, No. 2, pages 277–87.

US Department of the Treasury (1988), *Report of the presidential task force on market mechanisms*, Washington D.C.

Venkataraman, K and Weisssberg, A (2007), 'The value of the designated market maker', *Journal of Financial and Quantitative Analysis*, Vol. 42, No. 3, pages 735–58.