International Review of Financial Analysis 24 (2012) 1–11

Contents lists available at SciVerse ScienceDirect

International Review of Financial Analysis

Self-affinity in financial asset returns ☆

John Goddard [a,*], Enrico Onali [a,b]

[a] Bangor Business School, Bangor University, Bangor, Gwynedd, LL57 2DG, UK
[b] Xi'an Jiaotong-Liverpool University, Suzhou, Jiangsu, 215123, PR China

ARTICLE INFO

Article history:
Received 28 October 2011
Received in revised form 11 March 2012
Accepted 21 June 2012
Available online 28 June 2012

JEL classification:
C15
C16
G15

Keywords:
Market efficiency
Self-affinity
Fractional integration
Long memory
L-stable process

ABSTRACT

We test for departures from normal and independent and identically distributed (NIID) log returns, for log returns under the alternative hypothesis that are self-affine and either long-range dependent, or drawn randomly from an L-stable distribution with infinite higher-order moments. The finite sample performance of estimators of the two forms of self-affinity is explored in a simulation study. In contrast to rescaled range analysis and other conventional estimation methods, the variant of fluctuation analysis that considers finite sample moments only is able to identify both forms of self-affinity. When log returns are self-affine and long-range dependent under the alternative hypothesis, however, rescaled range analysis has higher power than fluctuation analysis. The techniques are illustrated by means of an analysis of the daily log returns for the indices of 11 stock markets of developed countries. Several of the smaller stock markets by capitalization exhibit evidence of long-range dependence in log returns.

1. Introduction

Long-range dependence and stable laws in asset returns have been investigated by researchers in finance for several decades. Long-range dependence implies a power–law decay of the autocovariance function in the time domain (Banerjee & Urga, 2005). Stable laws accommodate departures from normality and the central-limit theorem for independent and identically distributed returns or log returns (Levy, 1925). Following the pioneering work of Mandelbrot (1963, 1967, 1971), models that accommodate long-range dependence and stable laws have been employed to describe stock market behavior. These models represent an application of fractal mathematics to financial economics, a topic that has attracted widespread interest, and controversy, in recent years.

A fractal exhibits the properties of self-similarity or scale invariance. Stock returns may exhibit the weaker property of self-affinity. After the application of a suitable rescaling transformation, which takes the form of a single non-random contraction dependent upon the time scale only, a self-affine returns series exhibits the property of self-similarity. A self-affine returns series has the same distributional properties (after rescaling) when returns are measured at any frequency, and is said to be unifractal or monofractal.

Two classes of process, in which log returns are either non-independent or non-normal or both, embody the properties of self-affinity and unifractality (Cont & Tankov, 2004; Mandelbrot, Fisher, & Calvet, 1997). First, if log returns are fractionally integrated, the log returns series measured at any frequency exhibits the property of long-range dependence, and the log price series is characterized as Fractional Brownian Motion (FBM).[1] Second, the class of probability distributions known as Levy-stable, Pareto-Levy stable or L-stable (Levy, 1925; Mandelbrot, 1963, 1967) includes several heavy-tailed distributions with infinite higher-order moments (including the variance).[2] If log returns are either fractionally integrated and long-range dependent, or L-stable with infinite higher-order moments, then log returns are self-affine. The scaling behavior of the series with respect to variation in the frequency (the time scale over which returns are measured) is conveniently summarized by a parameter known as the Hurst exponent.

This paper contributes to two strands of literature, on fractional integration and long-range dependence, and on L-stable distributions.

☆ The authors are indebted to an anonymous referee for helpful and constructive comments on an earlier draft of this paper. We are also grateful to Ladislav Kristoufek and all participants at the 17th Conference on Computing in Economics and Finance, San Francisco, July 2011. The usual disclaimer applies.

* Corresponding author.
E-mail address: j.goddard@bangor.ac.uk (J. Goddard).

[1] FBM is a generalization of Brownian Motion, the continuous-time analog of the random walk. FBM has increments that are long-range dependent and therefore non-random (Mandelbrot & van Ness, 1968). Fractionally-integrated processes are defined in Section A.1 of the Appendix.

[2] L-stable processes are defined in Section A.2 of the Appendix.

1057-5219/$ – see front matter
doi:10.1016/j.irfa.2012.06.004

We examine the performance of estimators of the Hurst exponent, in the case where there is long-range dependence, and in the case where the distribution of log returns is L-stable (and there is no long-range dependence). Hypothesis tests for departures from the NIID case are developed, based on the application to simulated NIID log returns data of two widely-used methods for estimating the Hurst exponent: rescaled range analysis (RRA),[3] and fluctuation analysis (FA).[4]

The performance of these tests under the alternative hypothesis is examined by evaluating power functions, using simulated self-affine process characterized as either long-range dependent, or L-stable with infinite higher-order moments. Monte Carlo simulations are employed, because the asymptotic properties of the RRA and FA estimators are indeterminate (Banerjee and Urga, 2005). In addition, we draw comparisons with the performance of other tests widely employed to estimate long-range dependence (Geweke & Porter-Hudak, 1983; Robinson, 1995), and the characteristic exponent of an L-stable distribution (de Haan & Resnick, 1980; Hill, 1975; Pickands, 1975).

In much of the previous literature, researchers have reported evidence concerning the fractal properties of financial returns series in the form of point estimates of the Hurst exponent, or graphical analysis of scaling behavior. In the absence of any basis for assessing the statistical significance of possible departures from the NIID case, however, much of this evidence is at best suggestive of the possibility that models based on fractal mathematics might provide a more satisfactory representation of the behavior of returns than models embodying the NIID assumption. This paper relocates several established but informal procedures within a conventional and formal hypothesis testing framework, enabling conclusions to be drawn based on the standard criteria of statistical inference.

The principal findings are as follows. Tests for departure from the NIID case based on RRA and FA perform well when log returns are self-affine and long-range dependent under the alternative hypothesis. In this case, the test based on RRA has higher power than the tests based on the three variants of FA that are considered. However, the test based on RRA performs poorly when log returns are self-affine and L-stable with infinite higher-order moments under the alternative hypothesis. In this case, the choice of sample moments over which the FA is computed is crucial: the FA should not consider sample moments whose true values are infinite. As an estimator of the Hurst exponent, the variant of the FA that considers finite sample moments only is unique (among the estimators considered in this paper) in terms of its reliability under both of the long-range dependent and L-stable alternatives to the null hypothesis of NIID log returns.

In the empirical application, the stock market returns of several smaller markets exhibit evidence of long-range dependence. While previous studies provide evidence on the relation between market efficiency and long-range dependence in international stock markets, they do not specifically investigate the link between market size and long-range dependence. Furthermore, most previous empirical contributions neglect the impact of short-range dependence and non-normality on the estimation of the long-range dependence parameter.

For two stock markets (the US and Netherlands) we find evidence of self-affine scaling behavior that is consistent with log returns having been drawn from an L-stable distribution with infinite higher-order moments, rather than long-range dependence. This latter result is important, since it challenges the presumption that the variance of returns is finite. An infinite variance renders several of the tools of mainstream finance inapplicable, but it does not necessarily indicate a departure from the conditions for weak-form market efficiency.

[3] RRA was introduced by Hurst (1951). Refinements are suggested by Mandelbrot (1972), Mandelbrot and Taqqu (1979), and Lo (1991).

[4] This study uses the variant of FA employed by Mandelbrot et al. (1997).

The remainder of the paper is structured as follows. Section 2 outlines the implications of self-affine alternatives to NIID log returns for the Efficient Markets Hypothesis (EMH), and several other tools of mainstream finance. Section 3 describes statistical tests for an NIID null hypothesis against self-affine alternatives. Section 4 reports critical values and power functions for these tests, based on Monte Carlo simulated data. Section 5 applies the tests to log returns data for 11 developed country stock market indices for the period 1987–2011. Finally, Section 6 summarizes and concludes.

2. Forms and implications of departure from NIID financial asset log returns

Researchers have focused on questions concerning the stochastic properties of financial asset returns since the beginning of the 20th century. Bachelier (1900) is widely credited as the first to use probability theory to analyze returns in financial markets, and to formally hypothesize that returns are independent and identically distributed (IID), and therefore unforecastable. Osborne (1959) used the model of Brownian Motion to replicate the apparently random pattern of stock price movements. In contrast to Bachelier, however, Osborne and many subsequent studies model the logarithm of price, rather than the raw price.

Portfolio theory (Markowitz, 1952, 1959) examines the relationship between the expected return of a portfolio and its risk. The mean-variance rule (MV-rule), that risk-averse investors select the portfolio with the highest expected return for a given variance (or lowest variance for a given expected return), provided the basis for the CAPM, which posits a linear relationship between expected return and systematic risk. The MV-rule is rooted in the assumption that returns are normally distributed, as well as IID. If returns are other than NIID, moments higher than the first and second might enter the investor's utility function. The CAPM remains valid if returns are other than NIID, however, in the case where agents' utility functions are quadratic (Cochrane, 2001).[5]

We let p_t denote log price at time t, and $\Delta^{(n)}p_t = p_t - p_{t-n}$ denote the logarithmic return measured over the time scale n. If log returns measured in discrete time are IID, then log price follows a random walk $\Delta^{(1)}p_t = (1-L)p_t = u_t$, where L is the lag operator, and u_t is IID. According to the EMH (Fama, 1970; Fama, Fisher, Jensen, & Roll, 1969), a market is efficient when asset prices adjust continuously in response to new information, and reflect at all times the information that is relevant to their valuation. In a weak-form efficient market, an investor cannot use the information impounded in past prices to earn an abnormal return. Accordingly, the random walk hypothesis in respect of log price, with IID log returns, provides a simple model of price formation under conditions of weak-form efficiency.

If log returns (rather than raw returns) are NIID, then log price follows a Brownian Motion in continuous time. Price follows a Geometric Brownian Motion, and satisfies the non-negativity requirement. NIID log returns are accommodated by an adaptation of CAPM known as the Log Normal Asset Pricing Model (LAPM) (Cohen & Levy, 2005). NIID log returns underpin the Black and Scholes (1972, 1973) model of option pricing.

A log returns series is described as Fractional Gaussian Noise (FGN) if the decay of the autocovariance function follows a power law (see Appendix A.1). FGN is one member of a family of long-range dependent fractionally integrated processes which includes the ARFIMA(0,d,0) model, $\Delta^{(1)}p_t = (1-L)p_t = (1-L)^{-d}u_t$, where u_t is NIID. The parameter d is the order of fractional integration. Asymptotically the autocovariance function for ARFIMA(0,d,0) exhibits power–law decay. The scaling behavior of p_t is described by the Hurst exponent (Hurst, 1951), denoted H. For a fractionally-integrated process, the Hurst exponent is a simple transformation of the order of fractional integration, $H = d + 0.5$. For

[5] A quadratic utility function implies increasing absolute risk aversion, and may therefore be unrealistic (Herings & Kubler, 2007).