Available online at www.sciencedirect.com
ScienceDirect
Physica A 377 (2007) 155–165
PHYSICA A
www.elsevier.com/locate/physa

A note on scaled variance ratio estimation of the Hurst exponent with application to agricultural commodity prices

Calum G. Turvey*,[1]

Department of Applied Economics and Management, 357 Warren Hall, Cornell University, Ithaca, NY 14853, USA

Received 13 September 2006; received in revised form 15 October 2006
Available online 5 December 2006

Abstract

The measure of long-term memory is important for the study of economic and financial time series. This paper estimates the Hurst exponent from a Scaled Variance Ratio model for 17 commodity price series under the efficient market null $H_0{:}H = 0.5$. The distribution about the estimates of H are obtained from 90%, 95% and 99% confidence intervals generated from 20,000 Monte Carlo replications of a geometric Brownian motion. The results show that the scaled variance ratio provides a very good and stable estimate of the Hurst exponent, but the estimates can be quite different from the measure obtained from rescaled range or R–S analysis. In general commodity prices are consistent with the underlying assumption of a geometric Brownian motion.

Keywords: Fractional Brownian motion; Geometric Brownian motion; Scaled variance ratios; Agricultural commodity prices

1. Introduction

This paper is concerned with the sampling space of a geometric Brownian motion (gBm)

$$dx = \alpha x\,dt + \sigma x\,dZ, \tag{1}$$

where $dZ = \varepsilon\sqrt{t}$ is a Gauss–Wiener process, and x is a time-independent random variable with instantaneous rate of change, α, and standard deviation σ. This process is fundamental to the concept of financial economics and a variety of studies have attempted to measure whether Eq. (1) holds, or whether there is persistence or memory in markets (see for example Refs. [1–19]). We treat (1) as null over the more general set of fractional Brownian motion (fBm) described by

$$dx = \alpha x\,dt + \sigma x\,dZ^H, \tag{2}$$

where dZ^H characterizes a fBm with Hurst parameter $H \in (0,1)$. Under the null, $H = 0.5$ transforms (2) to (1). A fBm is a time series sequence that is non-stationary, self affine , with covariance over a time step or duration

*Tel.: +1 607 255 5230.
E-mail address: cgt6@cornell.edu.
[1]Calum Turvey is the W.I. Myers Professor in the Department of Applied Economics and Management, Cornell University.

doi:10.1016/j.physa.2006.11.022

$k = \Delta t$ (see [20–22]) equal to

$$E([x(t) - x(0)][x(t + \Delta t) - x(t)]) = \frac{1}{2}\sigma^2\left([t + \Delta t]^{2H} - t^{2H} - \Delta t^{2H}\right), \tag{3}$$

and variance over Δt is equal to

$$E[x(t + \Delta t) - x(t)]^2 = \sigma^2(\Delta t)^{2H}. \tag{4}$$

By setting $H = 0.5$ the right-hand side of (3) collapses to zero and the independent increments and stationarity assumptions are satisfied for a gBm, and (4) reduces to the standard gBm feature that the variance is linear in time. As H approaches zero the limit of covariance approaches $-0.5\sigma^2 < 0$, and variance falls. As $H \to 1.0$, covariance approaches $\sigma^2(t\Delta t) > 0$ and variance increases. For $H < 0.5$ the covariance term decreases with increasing time steps. Hence the term 'short memory'. In contrast, the term 'long memory' comes from the results that covariance increases with increased time steps when $H > 0.5$. However, this is not to the exclusion of the possibility that $H < 0.5$ can also be a long memory process [13]. The respective Wiener processes in (1) and (2) therefore possess markedly different properties.

To place the significance of this note in the context of (financial) economics it is worthwhile to pursue the distinction between a stochastic process that is self-similar with that which is self-affine. There are other motivations, one of course being the ability or inability to arbitrage markets and confound the efficient market hypothesis, but this motivation is also linked to affine structure. The Wiener process dZ is self-similar in time, whereas dZ^H is self-affine. While conceptually similar, self-similarity and self-affinity differ in the following way (see Ref. [23] or [24]): Suppose that an initial sequence or set $\{X_1, X_2, X_3\}$ can be transformed to the set $\{r_1X_1, r_2X_2, r_3X_3\}$, then the transformation is said to be self-similar if $r_1 = r_2 = r_3$ and self-affine otherwise. If the variance obeys the power law $VAR = \sigma^2(\Delta t)^{2H}$, (i.e. Eq. (4)) it is self-affine over the entire range of H, but is self-similar only for $H = 0.5$. Therefore, and generally speaking, self-similarity is a special case of self-affinity. If indeed $r_1 = r_2 = r_3$ there is nothing in scale that would provide an advantage to X_3 over X_1, but if the scaling were self affine on the order (for example) $r_1 < r_2 < r_3$ then some advantage could be had as a result of non linearity in the scaling.

Eq. (4) is central to this analysis. Dividing the left-hand side by the 1-period variance gives the variance ratio

$$\frac{E[x(t + \Delta t) - x(t)]^2}{\sigma^2} = (\Delta t)^{2H}, \tag{5}$$

which establishes a power rule that can be used to estimate H from observed variance, and in a manner that is quite distinct, albeit consistent with, R/S analyses [25–27] and other methods including Scaled Window Variance [28] and surrogate methods (see Refs. [29,30]; see Ref. [31] for a critique). We refer to our measure as the Scaled Variance Ratio, but view it within the class of Scaled Window Variance models. In these measures the numerator discretizes the sampling space into subseries of uniform length, computes variance at each length and compares (in general) the discretized variances to the variance of the original sample. If the relationship between the scaled variance and the unscaled variance is linear in Δt then $H = 0.5$. While the principle remains the same across the various methods of calculating H, the approaches differ in how the data is scaled.

Use of the variance ratio as an estimator for determining a random walk has been well developed in the literature (e.g. [32]) and need not be repeated here. However, the relationship between the variance ratio and the spectral density in the context of long-term memory (see Refs. [33–35]), although noted in Ref. [32], has not been to our knowledge been fully exploited to provide estimates of H in the Brownian motion context of Refs. [25–27,36]. In other words the financial economics literature that uses the variance ratio selects a single level of aggregation, k, and then decides whether the ratio for a given k is statistically different from unity. Compare this to the measure in Eq. (5) that computes the variance ratio for a range of continuous k-steps and then applies the power law defined in H (hence the use of the term Scaled Variance Ratio). The key result is that if the time series evolves as a gBm and satisfies certain stationarity conditions, the variance is linear in time as depicted in Eq. (4). In other words, the variance of price changes over a 2-day period ($k = 2$) will be twice the variance of the change in 1 day or the variance over 20 days ($k = 20$) will be 20 times the variance