European Journal of Operational Research 227 (2013) 122–132

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Patterns in stock market movements tested as random number generators

John R. Doyle [a,*], Catherine H. Chen [b,1]

[a] Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, Cardiff CF10 3EU, United Kingdom
[b] Department of Accounting and Finance, Middlesex University, The Burroughs, Hendon, London NW4 4BT, United Kingdom

ARTICLE INFO

Article history:
Received 24 September 2012
Accepted 30 November 2012
Available online 8 December 2012

Keywords:
Stock market time series
Financial data mining
Forecasting
Finance
Overlapping serial test

ABSTRACT

This paper shows that tests of Random Number Generators (RNGs) may be used to test the Efficient Market Hypothesis (EMH). It uses the Overlapping Serial Test (OST), a standard test in RNG research, to detect anomalous patterns in the distribution of sequences of stock market movements *up* and *down*. Our results show that most stock markets exhibit idiosyncratic recurrent patterns, contrary to the efficient market hypothesis; also that OST detects a different kind of non-randomness to standard econometric long- and short-memory tests. Exposure of these anomalies should contribute to making markets more efficient.

1. Introduction

Within disciplines that make practical use of simulation, such as management science, researchers must take an interest in Random Number Generators (RNGs). There is always a need to ensure that the results of simulation models are not contaminated by properties of the RNG used to actualize those models. Consequently, researchers have devised many ways of detecting whether pattern exists in the output of RNGs, in contradiction of its claim to randomness. If flaws are exposed, researchers eventually devise new RNGs, which then spawn new tests, in an endless evolutionary cycle. In particular, new RNG tests become ever more subtle in detecting new sources of non-randomness, and are particularly valued for doing so.

There are parallel enterprises in science, since it too searches for patterns. Sometimes the patterns may be obviously present (yet still in need of explanation). But sometimes the issue is whether patterns exist at all – is there signal in the noise? For instance, are galaxies uniformly distributed throughout the cosmos, as required by certain theories of physics, or do they tend to clump together? In the biological sciences, researchers are interested in the degree of randomness in EEG, EMG, or cardiovascular readings; and in the mathematical/computational sciences researchers may even compare the apparent randomness of consecutive digits in irrational numbers such as π, e, $\sqrt{2}$, or $\sqrt{3}$. Similarly, in the social sciences, do fluctuations in internet activity represent national security threats, emerging fashions, political shifts; or are they just variations in background chatter? For these problems, the knowledge gained in testing RNGs may prove useful in deciding whether patterns exist in the apparent randomness of phenomena.

Inspired by this philosophy, this paper borrows the Overlapping Serial Test (OST) from RNG research to test the widely debated and researched Efficient Market Hypothesis (EMH) in Finance, treating markets as if RNGs. By examining whether the stock market exhibits its anomalous patterns of movements, and using 76 daily closing price indices from around the world, we show that O-S tests are able to detect anomalies in the distribution of their patterns of movements in most indices. That is, most stock exchanges are not efficient markets, as judged by their principal price indices. Also, we find large variations in the degree and types of market inefficiency between different indices, even between indices quoted on the same stock exchange.

This paper therefore makes several contributions. First, we demonstrate the value of borrowing tests from different domains. Good's (1953) overlapping serial test falls within the remit of OR,[2] as it is commonly used in testing RNGs, which obviously underpin all simulations; but it is not part of the standard toolbox that econometricians use to test for analogous patterns in market movements. Not only does OST detect significant departures from randomness, but it detects departures that are different from (orthogonal to) those detected by standard EMH tests. Another contribution is to set these findings in a broader context. Although the application in this paper is about stock market movements, the approach may be applied to

* Corresponding author. Tel.: +44 (0)2920 879695.
E-mail addresses: DoyleJR@cf.ac.uk (J.R. Doyle), C.Chen@mdx.ac.uk (C.H. Chen).
[1] Tel.: +44 (0)20 87104910.

[2] Mulvey (1994) identified the area of security anomalies as a fruitful topic at the intersection of OR and finance. Similarly, Board et al. (2003, p. 12) wrote that: "Finance problems are an excellent application area for OR researchers. OR techniques are used to . . . identify market imperfections..." OST is exactly such a technique.

0377-2217/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2012.11.057

other kinds of problems at the limits of signal/noise discrimination. The interdisciplinary approach adopted and the results of this paper have interesting implications to academic researchers from different disciplines and also to investors. In the short term it may provide speculative investors a short-term boost in returns; but as we discuss later, longer-term it should also provide passive investors who hold the market portfolio the enduring benefit of trading in a slightly more efficient market.

2. Testing EMH as if a RNG

2.1. EMH and standard tests in finance

EMH posits that the stock price fully and fairly reflects all publicly available information, past and present, about that stock, and cannot be predicted in advance, and therefore should have the appearance of a memory-less random walk (Fama, 1965). Many tests have been used to test EMH (Ziemba, 1994). One broad class of tests examines the predictability of sources external to the stock, for instance, the presence of good weather at the exchange (Hirscheifer and Shumway, 2003); national success or failure in prominent sports event (Edmans et al., 2007; Nicolau, 2012); the pronounceability of stocks' ticker symbols (Alter and Oppenheimer, 2006); or calendar effects (Schwert, 2003). The January effect (Rozeff and Kinney, 1976) and the Monday or weekend effect (Cross, 1973; Hui, 2005) are the most scrutinized in past research on calendar effects; also sometimes analyzed together (Holden et al., 2005). The classic findings were that returns on Monday were lower than on Friday (or the rest of the week); and that returns in January were higher than for the other months of the year, although more recent research has modified these findings, such as by acknowledging their changing nature (Doyle and Chen, 2009). While day and month may be known exactly, forecasting weather and sporting outcomes, and measuring ease-of-pronunciation are all measured subject to uncertainty. But what these tests all have in common is to posit that share prices are partially contingent on market–irrelevant information, which can be known in advance, albeit sometimes only approximately. Consequently, if a particular test is significant it contradicts the EMH.

A second broad class of tests uses past stock prices themselves to infer the predictability of the market, including autoregression analysis, runs tests, Variance Ratio tests, Dickey Fuller test, the Box-Pierce test, rescaled range (*R*/*S*) test, and the KPSS test, etc., see Andrews (1991), Campbell et al. (1997), Giraitis et al. (2003), Kim and Shamsuddin (2008), and Tabak and Lima (2009). Some tests predict future returns explicitly, such as autoregression; while others do so implicitly, such as the variance ratio test (Lo and MacKinlay, 1988, 1989), or the runs test (Fama, 1965). Some tests (e.g. Variance Ratio test, total runs tests) examine short-term effects while others (e.g. *R*/*S*, rescaled variance test) test for long-term memories (Fama and French, 1988; Lo, 1991; Campbell et al., 1997). The common principle among these tests is to detect the clustering of similar returns, such as positive/negative autocorrelation (Hurst, 1951; Lo, 1991).

2.2. Efficient markets as RNGs

The design of research into RNGs presents an interesting parallel but contrasting development with that of research into EMH. It is clear that RNGs may be flawed by over- or under-production of any pattern, not just those patterns detected by runs tests. Therefore, it is usual in RNG research to test for anomalies in the distribution of *all* patterns. Because of its extensive practical experience with general pattern detection, RNG research is a useful starting point for similar research into EMH. In this spirit, this paper treats markets *as if* they were RNGs, bringing to bear the same kind of perspectives and methods that have been found useful in RNG research to test for randomness, which is also a central feature of an efficient market.

There are two kinds of RNG (Hayes, 2001). The majority are actually pseudo-RNGs. They are algorithms based on mathematical number theory, which are deterministic and will run through exactly the same sequence if started with the same seed. The second kind are hardware RNGs that aim to capitalize on the unpredictable nature of actual events in the world, such as radioactive decay, thermal electronic noise, or (in past generations) the roll of a dice.

One property of weak-form EMH is that the direction of returns should be unpredictable. If this property were fulfilled, an efficient market (EM) would be a candidate for being a hardware RNG.[3] If we code only the direction of market returns, and thereby effectively remove heteroskedasticity from the sequence, then an efficient market should generate a random stream of 0s and 1s representing negative and positive returns, respectively. Whether a market is efficient can then be tested in much the same way that a hardware RNG might be tested.

The conception of EMs as (hypothetical) RNGs helps us in three ways. First, by stripping the market returns process down to a point that EM becomes a plausible RNG, we gain transparency. Often the opaque part of testing EMH with untransformed returns is dealing with issues such as heteroskedasticity, since there is always some doubt about how best to achieve this. We cite, for instance, the adjustments made for heteroskedasticity in Lo and MacKinlay's (1988, 1989) variance ratio statistic; or the need for GARCH or stochastic volatility components in regression analyses in order to model away heteroskedasticity. Each generates its own uncertainty of interpretation: if the test rejects or does not reject the null hypothesis of an EM, is it due to the essential properties of an EM, or the way that heteroskedasticity has been treated, or even an interaction of the two? But by first transforming returns from interval scale to binary data, the method of removing heteroskedasticity is isolated from subsequent analyses, and the two steps can be examined separately. This advantage would be bought at the price of a decrease in statistical efficiency, were it not a "stylized fact" that idealized Gaussian distributions are seldom seen in this kind of data, as they typically have positive kurtosis – i.e. highly stretched tails – which can itself blight analyses that use classical statistics on interval data.

Second, we can capitalize on the experience of multiple testing in RNG research to guide us towards effective tests that may have been neglected in the finance literature. Researchers assess RNGs using a battery of tests, often packaging them into publicly available software suites. Noteworthy examples are: TestU01 (L'Ecuyer and Simard, 2007); Diehard (Marsaglia, 1996); and NIST Statistical Test Suite for Random and Pseudorandom Number Generators (Rukhin et al., 2008). One test that appears in all three RNG-testing suites is Good's (1953) Overlapping Serial Test (OST). In their study of 57 RNGs, Xu and Tsang (2007) show that the overlapping serial test is at least a powerful as the Gorilla test, which Marsaglia and Tsang (2002) suggest it is one of the three most "difficult-to-pass" tests for a RNG. In this paper we focus on OST as an indicator of market inefficiency. The OST simultaneously analyses all patterns

[3] *Note*: the efficient market as a RNG is a thought experiment, not a practical suggestion. Several properties of EM would disqualify an EM as a RNG. A RNG aims to emit numbers that are uniformly distributed in a fixed range, whereas a typical returns distribution is both dense towards the centre, and has no fixed upper or lower bounds. Moreover, the scale of the returns distribution varies over time with some degree of predictability (known as volatility clustering or heteroskedasticity). Furthermore, an EM is even permitted a slight preponderance of positive (or negative) returns to account for long-term upwards (or downwards) "drift" of its own. However, these should not be problems for this paper when returns are coded as only 0 or 1, and when the general trend of market is taken into account.