Pacific-Basin Finance Journal 44 (2017) 97–112

Contents lists available at ScienceDirect

Pacific-Basin Finance Journal

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

Toward a model-free measure of market efficiency

Keith R.L. Godfrey

The University of Western, 35 Stirling Highway, Nedlands 6009, Australia

A B S T R A C T

This article aims to measure market efficiency without an information model. The intuition is that an efficient market leaves no arbitrage opportunities for active traders, so the measure of efficiency (MOE) is the proportion of profits available to passive traders for a given level of transaction costs. It is expressed as a percentage score and defined symmetrically with a measure of inefficiency (MOI). It can be computed sequentially from a price series and a round-trip transaction cost. The measure of efficiency is shown to increase with diversification, reduce in longer time periods, and have an inverse relation with volatility. It is shown to be a leading indicator of price movements on a day-to-day basis and ahead of the financial crisis of 2008.

1. Introduction

Market efficiency is the intuitively appealing concept that describes the speed with which market prices respond to news. An efficient market will price new information rapidly and leave no unexploited opportunities for arbitrage. Despite the apparent simplicity of this concept, research into market efficiency has been debated in the literature for more than one hundred years. Studies reach back to Bachelier's (1900) “Theory of Speculation” with a broad resurgence following Fama's (1965a) work on stock-market price behavior. The fact that this question of efficient markets is still the subject of so much discussion after 115 years is an indicator of how surprisingly complex the topic can be.

While it may seem natural for price discovery to occur faster in a more efficient market, it is rare for the empirical literature to make a relative comparison between markets. The more typical approach is a yes-or-no question of whether or not a market is efficient in respect to a particular information set (Jensen, 1978). The efficient markets hypothesis (EMH) of Fama (1970) formed into three tests for different information sets: the weak-, semi-strong, and strong forms, respectively for historical price information, publicly available information, and private information. In each case the hypothesis proposes that the market trades at prices that fully reflect all the information in the category. A general test of the EMH may comprise a set of any number of yes-or-no tests for particular information sets.

There has been lively theoretical debate in the literature about whether markets are (or even can be) truly efficient. Sewell (2011) gives an overview. Proponents of the EMH such as Fama (1965a, 1970, 1991, 1998) describe mixed empirical support for each of the weak-, semi-strong, and strong subsets. Behavioral economists such as Barberis and Thaler (2003) argue mispricing can and does occur in a predictable way because of human biases and errors of reasoning. Practitioners and investors such as Buffett (1989) argue that they have made money consistently because the market is not always efficient.

The extreme definitions of efficiency and inefficiency can be counterintuitive. Consider for example the case of information leakage around earnings announcements. If an event study finds that stock prices do not move ahead of announcements (implying there is no insider trading) we say the market is inefficient (but legal), while at the other extreme if prices move fully to their new levels before the announcement we say the market is efficient (although the trading is illegal). This example leads to legal trading

E-mail address: keith.godfrey@uwa.edu.au.

http://dx.doi.org/10.1016/j.pacfin.2017.06.002
Received 26 March 2017; Received in revised form 8 May 2017; Accepted 3 June 2017
0927-538X/

being inefficient and illegal trading being efficient. Does it sound right to want an inefficient market?

Most commonly a market is neither fully efficient nor fully inefficient, instead being somewhere in between. There will inevitably be some level of noise trading for rational reasons unrelated to stock information (for example investors needing to sell shares to use the funds for other purposes) and this means we can expect to see weak-form inefficiencies in an otherwise efficient market. There may also be partial price movements ahead of information announcement when a few holders of leaked information exploit it to a limited extent (for example being constrained by the sizes of their trading accounts or their fear of being caught) which means we can expect to see a few strong-form efficiencies in an otherwise inefficient market. Overall it would be useful to have a numerical scale that describes the degree of efficiency or inefficiency, as this can then be used as a comparison between markets and time periods, rather than a binary decision that lands on one extreme or the other.

This paper develops a numerical measure that aims to capture the efficiency of the market without modelling the information driving the prices.

2. Classical tests of information-efficiency

Empirical studies show that market prices respond to information. Ball and Brown (1968) demonstrate a relation with accounting income numbers. Fama et al. (1969) show a relation with stock splits and dividend levels. These relations suggest that the stock market is efficient in the sense that prices adjust to new information.

How rapidly and accurately do stock prices respond to information? Testing the speed that prices respond to information leads to a joint hypothesis problem. On one hand we are measuring how fast the market converges to the expected target price, but on the other hand we are imposing a model of what we think that target price should be. Tests of the EMH become inseparable from the hypothesis that the information valuing model is correct. For this reason, the classical tests of market efficiency have tended to focus on special cases which provide greater certainty about the pricing and timing of information.

One group of tests exploits situations where the expected pricing can be inferred through a relation between the securities. Intrinsic relations between the securities in closed-end funds, depositary receipt shares, and dual-listed or Siamese twin stocks provide unique opportunities for researchers to separate the noise trader risk from the fundamental risk. While noise trading is an essential part of market behavior (Black, 1986) it also introduces inefficiencies. Gemmill and Thomas (2002) analyze 128 closed-end funds and conclude that the noise-trader risk is present although not the cause of the fundamental inequalities, suggesting there are multiple sources of inefficiency. Scruggs (2007) examines the Royal Dutch/Shell and the Unilever NV/Plc pairs and finds that 15% of the weekly return variation is attributable to noise. These analyses were possible in these markets because of the intrinsic relations between the securities.

Another group of tests studies the timing of price movements around isolated high-impact information events such as earnings announcements and central bank decisions. These event studies rely on the premise that major news events dominate the price movements around their times of occurrence, so it becomes possible to measure the speed of response without needing to value the information. The market's convergence to a new price level becomes the proxy for what the “right” price should be. A fast response (or even an anticipatory move in the case of leaked information) is said to be efficient, while a slow response or a period of oscillation before convergence is labelled inefficient.

While related-pricing tests and event studies are useful in studying the markets and time periods when they occur, neither is able to generalize to an arbitrary market at an arbitrary time period. A general test of market efficiency would require an information valuing model, with its associated joint hypothesis.

2.1. Random-walk proxies for market efficiency

One way to avoid the joint hypothesis problem is to look instead for the presence of characteristics that proxy for the desired phenomenon. The most common proxy for market efficiency evolved from random walk theory and the idea that prices in an efficient market behave like a random walk in which returns cannot be predicted. This theory evolved at around the same time as the EMH, with Fama (1965a) linking random walks to market efficiency, and Samuelson's (1965) proof that properly anticipated prices will appear to fluctuate randomly like a random walk.

Tests for random walk-ness are based on statistical properties of the price series such as unit-root stationarity, variance ratios, normal distributions, and lack of serial autocorrelation. These all can be computed without needing an information value model. For a summary of the methods see Ball and Kothari (1989) and Urquhart and Hudson (2013). For further tests of return predictability, Kim et al. (2011) describe an automatic portmanteau test and generalized spectral tests.

Finding characteristics contrary to a random-walk (such as the presence of serial correlation or return predictability) provides conclusive evidence of inefficiencies. Worthington and Higgs (2009) test the Australian stock market and find specific inefficiencies in the price series. Ito and Sugiyama (2009) use the time-varying autocorrelation of stock returns as a proxy for the extent of market inefficiency. Lim et al. (2008) use rolling window tests of serial correlation to study the efficiency of eight markets around the Asian financial crisis of 1997 and find greater inefficiencies in particular markets and periods. On the other hand, finding random-walk-like characteristics such as a lack of serial correlation and a normal distribution of runs does not guarantee that the price series is a random walk, nor does it enable a conclusion that the market is inefficient. This is a limitation for these kinds of tests. They tend to lead to one-sided conclusions by finding evidence of inefficiency.

When tests for randomness become more complex statistically, they risk bearing little or no resemblance to the mechanics of stock market trading. Wang et al. (2010) develop a degree of market efficiency (DME) using a technique of multifractal detrended