

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

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/ © 2017 Elsevier B.V. All rights reserved.

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

fluctuation analysis which derived from studies of correlation of biological molecular chains in deoxyribonucleic acid (DNA). The DME is argued to sense changes in the efficiency of the Shanghai stock exchange despite it being seemingly unrelated to the mechanisms of stock price data series and statistical arbitrage trading.

The usefulness of these kinds of randomness proxy tests for the EMH rely in part upon the methods for detecting randomness and in part upon the random walk hypothesis (RWH) that efficient markets should look like a random walk. Despite evidence against the random walk model itself (Lo and MacKinlay, 1988; Griffin et al., 2010) the mathematical appeal of these tests has ensured their continued use in studies of market efficiency. They remain most useful as binary tests, finding randomness (or not) as a proxy for efficiency (or not). It would be difficult to extend them to a meaningful numerical scale because a degree of randomness does not align intuitively with a degree of efficiency.

2.2. Simulations of arbitrage profitability

A separate group of tests emerged from the premise that an efficient market leaves no opportunities for arbitrage. These tests work by looking for price relations that could be exploited and simulating whether any particular trading algorithm could have been employed profitably on the price series. If an algorithm can be found, it is evidence of an inefficiency. An example is Kang et al. (2002) finding profitable contrarian and momentum strategies in the China Stock market, and showing further that the primary driver is overreaction to information. On the other hand, if no profitable algorithm is found, the result is inconclusive because we may not have tested the optimal strategy. It becomes a joint hypothesis problem with the universe of trading strategies that we choose to explore.

Sometimes it can be sufficient to document the price relation without specifying any particular algorithm to exploit it. The relation may be weak-form predictability of a single security's price, or jointly between two or more seemingly unrelated securities. As an example, Narayan et al. (2010) find the gold and oil futures markets to be jointly inefficient because of a relationship between them. They leave it to the reader to decide how to trade it.

Demonstrating that a market is efficient is harder because it means showing that no trading algorithm could have been profitable. Approaches include Conrad et al. (1997) showing that a significant portion of profits comes from the bid-ask bounce, and the profits disappear at trivial levels of transaction costs.

Intuitively, there is an inverse relation between market efficiency and trading efficiency. When a market is efficient, traders' attempts to profit from mispricing become inefficient, and vice versa. Smith (2001) calls this the difference between market and model efficiency, and develops several quantitative measures of market efficiency based on the degree of market inefficiency implied by successful trading models. Smith use the Sharpe ratio as the measure of model efficiency, representing the highest possible return for a given risk of a long-only portfolio. While Smith's model may match the needs of a quantitative trader looking to profit from weak-form inefficiencies, it suffers two major drawbacks for analyzing market efficiency. First, the Sharpe ratio provides a comparative analysis rather than an absolute scale of efficiency. Second, it is unlikely a successful quantitative trader will publish his or her most profitable algorithms, which means the true level of inefficiency being exploited is likely to remain undocumented.

A more complex approach is to build a matrix of relative inefficiencies within the universe of algorithms being explored. Stephens et al. (2007) consider the excess trading returns for traders who manage portfolios actively compared with buy-and-hold strategies, and they develop an inefficiency matrix between pairs of trading strategies. This matrix attempts to build a complete description of the relative inefficiencies that exist in an entire market, in the same way that a large covariance matrix represents the relations between stocks in the portfolio theory of Markowitz (1952). While a matrix of this kind can in theory provide a complete description, it suffers the same challenges as a large covariance matrix that in practice it becomes too cumbersome and unwieldy to calculate. Even if the matrix can be calculated, it does not provide a simple single measure for the overall degree of market efficiency.

While the tests such as random walk-ness and simulated trading can be computed without an information model and may provide some insight into market characteristics, their nature makes them more suited to finding examples of inefficiencies than conclusive evidence of efficient markets. There is also the risk that the simulations cannot be implemented in practice (Ball et al., 1995). At best they remain best binary tests, providing evidence for or against particular inefficiencies. We would like to move towards numerical measures of how efficient a market may be.

3. Requirements for an information model-free measure

We seek a measure of market efficiency that is a numerical scale rather than a yes-or-no test, and can be calculated from the price series without an information value model.

The progression from a binary test to a numerical measure has been an evolutionary step in many scientific fields of naturally occurring phenomenon. The physics of magnetism was documented as early as the 4th and 5th centuries BC with observations of certain materials attracting iron while others did not (a binary test). Fourteen centuries later the evidence of magnets being used in navigation led to a conclusion that the earth must be magnetic too (another binary question). It was not until the 1800's when an accidental discovery of an interaction between electricity and magnetism enabled researchers to start making precise numerical measurements of the strength of the effect. These numbers in turn enabled relative comparisons between materials and permitted Maxwell to formulate his famous set of equations that still explain electromagnetism today. In summary, the scientific knowledge evolved in three phases: first came the binary observations (research questions of "is the effect present under these circumstances?"), then the numerical measurements (questions of "how strong is it under these circumstances?"), and finally the theory (questions of "why do these circumstances cause it?"). If we map our 115-year discussion of market efficiency into this framework, it would seem

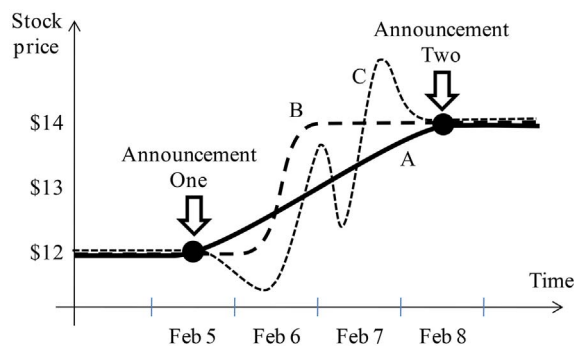


Fig. 1. An illustration of the problem in deciding market efficiency from news announcements. Suppose we observe a stock that issues public announcements on Feb 5 and Feb 8, and its price moves from \$12 to \$14 during the period by one of three hypothetical paths labelled A, B, and C. Path A is a gradual change between the announcements. Path B is an abrupt change at some time between the announcements. Path C involves several apparent overshoots and price corrections along the way. Our task is to explain each possible price movement in the context of the announcements in terms of the weak, semi-strong, and strong forms of the efficient market hypothesis so as to decide whether the market is efficient. Path A could be explained either as a slow response to the first announcement (semi-strong form inefficiency) or an early predictor of the second announcement (strong form efficiency) and we cannot distinguish between these two possibilities. Path B leads to a similar dilemma. We can however say confidently that path C appears to be less efficient than either paths A or B because it has added noise. Whenever the price oscillates along its trajectory in a way that can be characterized or modelled, we can imagine that it can be exploited by a trading strategy. In summary while we cannot make reliable conclusions about the market efficiency along paths A and B, we can say that path C seems less efficient than either A or B.

that we are still in the first phase. We are still asking questions of “is the market efficient?” and repeating this question for different information sets.

To move into the second phase we need to ask research questions of “how efficient is the market?” If we succeed at developing accurate means of measuring efficiency, we then may be able to move into the third phase to investigate why some markets are more efficient than others, and why a market may be more efficient at certain times compared with other times.

3.1. Price movements from announcements and noise

Markets prices move for many reasons. Sometimes there may be a major announcement with large and identifiable price impact. More frequently there are lower-impact announcements which are harder to distinguish from noise. Even more frequent are general news articles and advertisements that motivate people to trade, which may create a lot of the noise. We observe the resulting price series in general without knowing the events that prompted the trades.

Inferring price response from information can be an ambiguous process even when particular information events are identifiable. Consider the hypothetical stock in Fig. 1 which has the two announcement events indicated. Between those events, the stock price is observed to follow one of the three hypothetical paths illustrated. If path A is observed, it could be explained in the context of the EMH as either a slow “semi-strong-form” response to the first announcement or an early “strong form” prediction of the second announcement. These are equally plausible explanations, yet with different conclusions, because the market appears inefficient with respect to one information set and efficient to another. Observing price series B gives a similar dilemma, as the timing and speed of the price response does not necessarily help distinguish between the two explanations. On the other hand, price series C does appear to be less efficient than A or B because it overshoots and oscillates between the two announcements and in the absence of new information. We can be comfortable asserting this without needing to know anything about the information releases and their respective influences. The message in Fig. 1 is that we can infer a greater level of inefficiency when we observe the market pricing overshooting and correcting.

The overshoots and corrections in Fig. 1 may be considered a type of noise trading, but the presence of noise in a market does not necessarily provide a profitable trading opportunity. Black (1986) finds that while noise causes markets to be somewhat inefficient, it also often prevents us from taking advantage of inefficiencies. Shleifer and Vishny (1997) extend the discussion on noise trading and the potential limits of arbitrage. One of the biggest impediments to exploiting small movements is the presence of transaction costs (Conrad et al., 1997) and these must be considered in building any measure of market efficiency.

3.2. Characterizing the mechanisms of market efficiency and inefficiency

We can consider the mechanisms behind market efficiency and inefficiency by categorizing the reasons why people trade. One possible motivation is fundamental or “value investing” where an investor looks for securities that have not yet fully priced the information announced. Another arises from charting and statistical arbitrage, the idea being to trade based on quantitative analysis of the price history. A quantitative trader has a shorter time horizon than a fundamental investor and can seek to profit from short-term price variations, because frequent changes in fundamental valuation should not occur in efficient markets (Lehmann, 1990). A third potential reason may be personal, such as selling shares for tax optimization or to pay a private debt. These personal reasons will be harder for an observer to identify, and while they may seem perfectly rational to the individual investor, they may seem irrational in the context of fundamental and quantitative analyses. This third case is therefore often called “noise”. Noise can also occur from

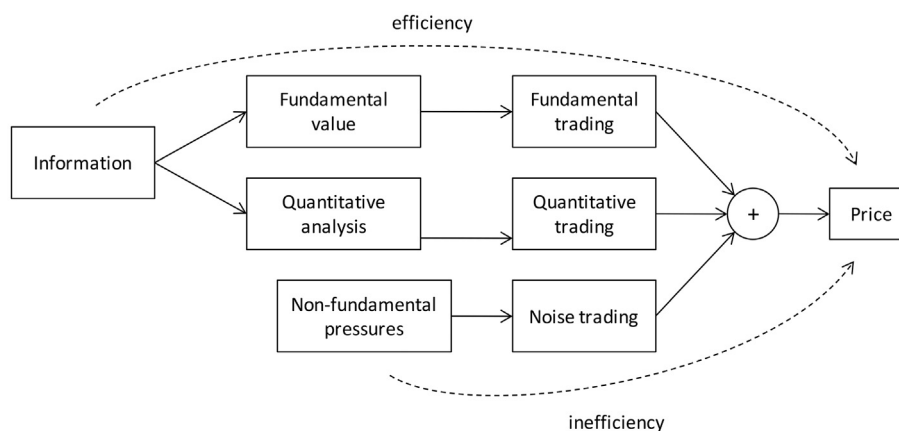


Fig. 2. Characterization of efficient and inefficient pricing mechanisms. In this characterization a security price is determined by a combination of fundamental trading, quantitative trading, and noise trading. When price movements are based on fundamental information being incorporated quickly, the market is said to be efficient. When the market digests information slowly or when prices are dominated by noise trading, the price movements are not related to fundamentals and the market is said to be inefficient. Quantitative trading lies somewhere between the other two, at times contributing to market efficiency by exploiting arbitrage opportunities and at other times creating inefficiencies such as continuing a momentum price trend beyond a reasonable fundamental valuation.

many documented behavioral reasons (Barberis and Thaler, 2003) and also the dissemination of non-information (such as the “phone call from Debbie” scam described by Small and Smith, 2007).

These pathways to security pricing are sketched in Fig. 2. In this characterization, the fundamental trading leads towards a market aggregated decision of a share's intrinsic value. Quantitative trading arises from numerical analyses of the price history and the released company data. Noise trading arises when trades are executed for reasons that are neither fundamental nor quantitative, even if they have a rational basis for the trader personally, and may move the price away from its intrinsic value.

The insight from Fig. 2 is there are two main pathways leading to market price: one which supports efficiency (arising from incorporation of relevant information) and one that promotes inefficiency (by moving the price for other reasons and away from fundamental information). Quantitative trading sits somewhere between the two extremes and could be categorized either way depending on its approach. Consider for example a trader who exploits the momentum effect demonstrated by Jegadeesh and Titman (1993). By trading in the direction of the price trend, the trader will initially help reduce any mispricing by speeding up the convergence process (which contributes to efficiency). By continuing to follow the momentum, the trader is likely to extend the trend too far and cause the price to overshoot its intrinsic value (which creates an inefficiency). Quantitative trading can therefore at times assist market efficiency and at other times hinder it.

Tests of weak-form efficiency put quantitative traders in the same group as the fundamental traders because the past price history becomes part of the information set to be exploited. Sullivan (2010) argues that fundamental and quantitative traders both make decisions grounded in solid theory, and they both foster market efficiency. On the other hand, tests of semi-strong and strong-form efficiency put quantitative trading in the same category as noise trading because it is unrelated to the fundamental information released. The argument that noise traders cannot survive for long (because they are trading against one another in a negative sum game plus transaction costs) can then be applied to the quantitative traders too. In practice however there is plenty of anecdotal evidence of particular quantitative traders being very successful over long periods of time, so the noise trading argument does not explain the whole picture.

The observed price series is the cumulative result of the actions by all traders whether fundamental, quantitative, or noise. Our task is to measure how far the price series leans towards the efficient or inefficient pathway. The approach comes from the quantitative traders in the middle. If a strategy such as momentum trading causes efficiency when supporting and speeding up a trend but inefficiency when overshooting or extending a trend, it would seem that one possible way to study efficiency is to measure the frequency and size of these overshoots or oscillations.

3.3. Desired characteristics for a measure of market efficiency

We want to be able to measure efficiency in any market or more generally any price series. The approach should work in commodity and agricultural markets, futures, foreign exchange rates, bonds, and options, as well as in the stock market. It should work on long and short portfolios as well as individual securities, and it should work in mathematically derived series such as the ratio of the prices between Royal Dutch and Shell as constructed by Rosenthal and Young (1990).

Before looking at a particular measure of market efficiency, we should consider the desired characteristics that any such a measure should demonstrate when applied to a variety of price series. The most basic requirements include:

- Simplicity of computation from the readily-available price data;
- Unambiguity of computation for each price sequence and period; and

Applicability to any price series without needing to know fundamental value.

To be fit for purpose we also require:

An intuitive relationship with information efficiency (it must be a valid proxy);

An absolute scale (for example from zero to 100%); and

A reversible scale (for example 0% efficiency is 100% inefficiency and vice versa).

To be consistent with the practicalities of trading we also require:

Symmetry of calculation (the same efficiency should be computed if a security is constructed as the short sold value of another security or price series);

Scale invariance of calculation (the efficiency should remain the same if a security were constructed to have a multiple or fractional value of another price series);

Parameterization by transaction costs (because we can anticipate efficiency to vary with the limits of arbitrage); and

No conflict with commercial trading secrecy (it must not expose any particular inefficiencies that could be traded, so there no impediment to publication).

We should further test any particular measure in market situations where we can anticipate particular properties. These could include:

When comparing the efficiencies of markets with different levels of transaction costs, we should expect the markets with the higher transaction costs to be more efficient on average because there is less profitable opportunity for arbitrageurs to exploit mispricing; When testing the effects of diversification, we can expect the efficiencies of highly diversified indices to be greater than those of less diversified indices because there should be greater opportunity on average for company-specific arbitrage; and

When testing different time periods, we can expect inefficiency to accumulate over time so that the trading over a longer time period should appear less efficient than over a shorter time period.

With these requirements in mind, we now move toward developing and testing particular numerical measures for market efficiency.

4. A model-free numerical measure of efficiency

The measure of market efficiency developed in this section exploits the relation between market efficiency and the opportunities to arbitrage profitably. In any time period, we compare the optimal performances of passive trading and active trading, computing how much unlevered profit could have been made from each strategy. Intuitively, if the market is efficient, the active traders should not be able to outperform the passive traders. It follows that the extent to which active traders can beat the passive traders provides a measure of inefficiency.

4.1. Definitions

Given a price series, a time period, and a round-trip transaction cost, we can define two types of optimal profit that could have been made by trading that price series:

the *maximum profit from active trading* (MPAT) is the maximum amount that a trader could have earned from the price movements, after transaction costs, by trading a position actively throughout that period, alternately long and short; and the *maximum profit from passive trading* (MPPT) is the maximum amount that a trader could have earned from the price movements, after transaction costs, by holding one position throughout the period, either long or short.

The mathematical definitions are provided in [Appendix A](#).

In both cases the trader is assumed to have perfect senses of timing and direction. The maximum profit from active trading arises when each directional leg is timed perfectly, and the maximum profit from passive trading means the market direction is picked correctly at the beginning of the period. In practice it is difficult for any trader to make correct decisions consistently in real time, but there are no such impediments for researchers to look back and see what optimal decision could have been made in hindsight.

These two profit figures represent the optimal unleveraged active and passive earnings, after transaction costs, during the period. Each is defined to be either positive or zero. The maximum profit from active trading must be at least as great as that from passive trading because the passive trade is just one of the set of possible active trade sequences available. If the price movement throughout the period does not exceed the level of costs, the position can either be left open on an ongoing basis without incurring costs, in which case the maximum profits from active and passive trading can both be treated as zero. This special case should be treated as market efficiency because no abnormal profit could have been made.

We can view the maximum profit from active trading as having two components: one from optimal passive trading over the same period, and one from the excess earned by active trading. The excess profit from active trading will therefore be either positive or zero. When it is zero, the market looks efficient. When it is positive, the market looks inefficient to some degree. The extent of the excess profit available provides a measure of inefficiency.

From the MPAT and MPPT we can define measures of efficiency and inefficiency:

the *measure of efficiency* (MOE) is the ratio of the maximum profit from passive trading to the maximum profit from active trading,




Fig. 3. Sequential computation of maximum profit from active trading. Maximum profit occurs when positions are held alternatively long and short throughout the period, with the changes occurring at local maxima and minima that differ in price by more than the transaction costs. The maximum profit on each leg is the absolute price movement minus the transaction costs.

with a special limiting case of one (100%) if both profits are zero; and the *measure of inefficiency* (MOI) is the ratio of the excess profit from active trading to the maximum profit from active trading, with a special limiting case of zero (0%) if both profits are zero.

Mathematically, $MOE = \frac{MPPT}{MPAT}$ when $MPAT > 0$, and the two measures MOE and MOI sum to 100%. The complete mathematical definitions are presented in [Appendix A](#).

These measures of efficiency and inefficiency have a simple intuition. The measure of efficiency MOE is the proportion of the possible active profit that is achievable through passive trading. The measure of inefficiency MOI is the excess proportion of the possible active profit that is attributable to active trading. The sum of the measures is always one (or 100%) so the market appears fully efficient when $MOE = 1$ (or $MOI = 0$) and fully inefficient when $MOE = 0$ (or $MOI = 1$).

The MPAT, MPPT, MOE, and MOI have many intuitive properties such as $MPPT \geq 0$, $MPAT \geq 0$, $MPAT \geq MPPT$, $0 \leq MOE \leq 1$, $0 \leq MOI \leq 1$, and $MOE + MOI = 1$. These are proven in [Appendix B](#).

4.2. Computational implementation

The definition of MPAT requires the maximum profit to be computed from any possible combination of trades, and this number of combinations grows rapidly with the number of time samples in the price series. Although all possible combinations could in theory be tested in parallel, this would be cumbersome to implement, and in the alternate it will become computationally inefficient in any sequential computing language. Fortunately, the definition of MPAT can be re-specified in a way that can be computed efficiently on a sequential machine.

The sequential algorithm for MPAT is based on the mechanics of trading a price series in practice. [Fig. 3](#) shows a sketch of a hypothetical stock price series. Our task is to measure the MPAT during this period. Intuitively, the optimal trading will occur when each leg maximizes the profit by buying at a local minimum price and selling at a local maximum, or a short selling at a local maximum and repurchasing at a local minimum, when the price movement exceeds the transaction costs. If the price movement does not exceed the transaction costs, then no trade is considered. The optimal trading sequence will involve alternate long and short positions, inverting at local maxima and minima that differ by more than the transaction costs. When a position must be held throughout the period, the first and last legs may be suboptimal without knowledge of the preceding and subsequent points of inflection, but the calculation of MPAT is nonetheless unique within this period.

The sequence for computing the MPAT in [Fig. 3](#) is as follows. The first leg starts at the initial price at time t_0 and proceeds until the price moves by at least the transaction costs. As this first move beyond the transaction costs occurs in a positive price direction, the optimal trading direction would be long. The leg continues until the local extreme at t_2 when the price starts to reverse and will reverse by more than the transaction costs. The profit for the first leg is the absolute difference between the starting price and the value at the local maximum at t_2 , less the transaction costs. The second leg starts at t_2 and proceeds in the opposite direction until the local extreme at t_3 when the price starts to reverse and will do so by more than the transaction costs. We can repeat the procedure for as many legs as needed. The finishing point will typically close the final trade for a profit, or may be left open. There will be a realized profit of zero if the price never moved beyond the transaction costs at all.

An example implementation of this computation is presented in [Appendix C](#) using the R statistical programming language.

While this specification is complete, we can also envisage minor variations to this approach, for example by terminating the first and last legs at the local extremes (t_1 and t_4 in [Fig. 3](#)) rather than the beginning and end points (t_0 and t_5). The more general definition is employed here to have a position open throughout the period, and this keeps the MPAT constructed in a similar manner to the MPPT.

5. Verification of desired characteristics

Having defined the measure of efficiency MOE and shown how it can be implemented computationally, we can test now whether it meets the characteristic properties as required.

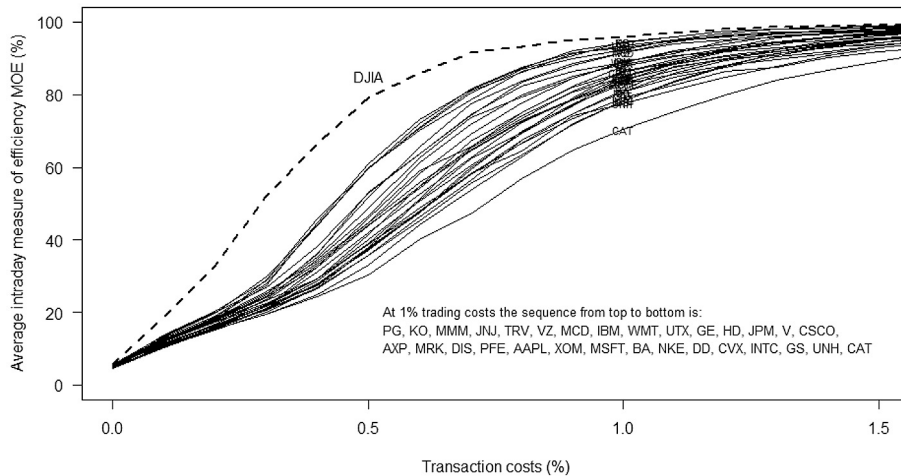


Fig. 4. Average intraday measures of efficiency (MOEs) of the Dow 30 stocks and the Dow Jones Industrial Average over the two years from 20 March 2015 to 19 March 2017. The measures of efficiency (MOEs) for the DJIA and its 30 constituent stocks are seen to increase with transaction costs, as predicted. These dates are chosen so there is no change in the Dow composition during the period.

5.1. Effect of transaction costs

Fama's (1991) overview of market efficiency works on an assumption of no transaction costs. In practice transaction costs are likely to affect market efficiency. While the imposition of trading frictions might seem intuitively to be an inefficiency, higher costs actually improve market efficiency because they impede noise traders. Higher transaction costs make people think more carefully before trading, and make it harder for an active trader to beat a plain buy-and-hold strategy. It follows that prices are more likely to end up near their fundamental values, or within a band of transaction costs from there. While higher costs will widen that band and provide greater opportunity for noise within it, such noise cannot necessarily be traded for profit (Black, 1986; Shleifer and Vishny, 1997).

Fig. 4 shows the average intraday measures of efficiency of the 30 securities in the Dow Jones Industrial Average as a function of transaction costs over the two years from 20 March 2015 to 19 March 2017. These dates have been chosen because there was no change to the composition of the index during this period (the previous change occurring on 19 March 2015). As predicted, we see an increase in the measure of efficiency (MOE) as the transaction costs are increased, and we see this for each security as well as for the index.

The effect of transaction costs has an analogy with the concept of a yardstick in fractals. The maximum profit from active trading MPAT is calculated after the price moves beyond the size of the transaction costs, just as the length of a coastline can be measured by moving beyond the length of the yardstick many times. In the same way that a smaller yardstick enables finer features of the coastline to be measured, for a greater cumulative length, a smaller level of transaction costs enables more trades to be executed profitably, for a greater cumulative profit. Fractals also have a historical relation to stock price analysis because it was Mandelbrot (1982) who popularized them and he also supervised Fama's (1965a, 1965b) PhD thesis on the behavior of stock market prices.

5.2. Effect of diversification

A second property visible in Fig. 4 involves comparing the measure of efficiency of the Dow Jones Industrial Average (DJIA) with the measures of efficiency for the individual stocks. The diversified index shows a higher measure of efficiency than its individual constituent stocks, as expected because the noise and mispricing opportunities within a single stock should typically be greater than in a diversified index.

We can expect the measure of efficiency to become even higher for increasingly diversified portfolios. We can test this by comparing the efficiencies of stock indices of different sizes. Fig. 5 shows the intraday measures of efficiency (MOEs) of the Dow Jones Industrial Average (30 stocks), the Nasdaq Composite (100 stocks), and the Standard and Poors 500 (500 stocks) from 2004 through to March 2017 calculated intraday on one-minute data with no transaction costs. The intraday MOEs are smoothed by a one-year moving average to remove much of the daily variation which would otherwise dominate the graph. We can see the smoothed measure of efficiency of the S & P500 is greater than that of both the Nasdaq 100 and Dow 30 throughout the period, as predicted. The measure of efficiency of the Nasdaq 100 tends to lie above that of the Dow 30, but crosses at times, perhaps because these two indices share no stocks in common. While there is some noise here, overall this is consistent with our prediction that larger diversified portfolios are more efficient.




Fig. 5. Average intraday measures of efficiency (MOEs) under zero transaction costs of three market indices showing the effect of diversification. The intraday measures of efficiency are shown for the Dow Jones Industrial Average (30 stocks), the Nasdaq Composite (100 stocks) and the S & P 500 (500 stocks), for no transaction costs and with each being smoothed by a one-year moving average. The highly diversified S & P 500 index appears consistently more efficient than both the Nasdaq 100 and the Dow 30. The efficiency of the Nasdaq with its 100 stocks tends to lie above the Dow with its 30 stocks, despite the compositions of the Nasdaq and Dow being substantially different (no overlapping stocks). The seemingly low levels of efficiency are a consequence of specifying zero transaction costs which enables a fair comparison between the indices but counts even the smallest movements as inefficiency.

5.3. Effect of time interval

The measured efficiency of a market is likely to be greater over short time intervals because there is less time for inefficiencies to occur. If we calculate the MOE of a diversified stock index over different sizes of time interval, we should expect the efficiency to reduce as the interval grows. Fig. 6 shows the MOE of the well-diversified S & P 500 index measured over intervals of one day, one week, and one month, across the 17 years from 2000 through 2016. The calculations are based on a fixed level of transaction costs at 0.25 index units which is the minimum bid-ask spread for S & P 500 futures quotes. For clarity the results are smoothed by a three-month moving average. We see that the measure of efficiency on daily periods is greater than on weekly periods, which in turn is greater than on monthly periods. This is consistent with our prediction.

5.4. Relation between the measure of efficiency and information efficiency

We would like to test whether the measure of efficiency MOE exhibits characteristics similar to information efficiency. Tests of this type are inevitably hard to design without creating a joint hypothesis problem with an information value model. The approach here is to look for behavior that comes from the price series alone. As an example, a larger presence or build-up of inefficiencies is likely to lead to a subsequent correction in prices. The intuition is that an inefficient response to information will facilitate an




Fig. 6. Average measure of efficiency of the S & P 500 index calculated over different time intervals. Transaction costs are set at 0.25 index units, which is the minimum bid-ask spread for S & P futures quotes, and lead to efficiencies slightly higher than in Fig. 5 where transaction costs were zero. The measures are smoothed with a three-month moving average. The inferred efficiency appears higher on the monthly calculation than the weekly, which is in turn higher than the daily calculation. These results are consistent with the prediction that shorter periods should appear more efficient. Longer periods offer more inefficiencies.




Fig. 7. Daily return on the S & P 500 versus the measure of efficiency MOE from the previous day. The data comprises > 17 years of index prices from January 2000 to March 2017, with transaction costs set at one index point. The chance of a large price movement, either positive or negative, is more likely following a day of low measured efficiency (high inefficiency). This is consistent with the idea that the presence of informational inefficiency is followed by arbitrage and price correction.

arbitrage opportunity and therefore a more significant price movement. This can be tested without needing to model the information that caused the inefficiency.

To test this, we can correlate the measure of efficiency from a previous period (such as the previous day) with the size of the return in the current period (such as the present day) to see whether inefficiency predicts price corrections. We expect to see larger corrections periods of higher inefficiency (lower efficiency).

Fig. 7 shows a scatter plot of the daily return on the S & P 500 index versus the measure of efficiency MOE of the previous day, for > 17 years of data from January 2000 to March 2017, and with transactions costs set at one index point. After days when the market is operating at higher efficiencies, the returns on the following day are relatively small. After days when the market is operating at lower efficiencies, there is a very much wider variation of subsequent returns. This shows an inverse relation between the MOE and the return over the ensuing day, which is as predicted.

The fact that the previous day's measure of efficiency (computed solely from the price history) can predict something about the following day's return, is itself evidence of a weak-form inefficiency. In an efficient market it should not be possible to predict anything about the future price movements.

It can also be argued that Fig. 7 shows a far greater number of samples at the inefficient end of the spectrum, so perhaps the inverse relation claimed is appearing as an artifact from having fewer data samples at the efficient end. It can also be argued that without an information value model we cannot be certain the larger returns observed after inefficient days are corrections related to those inefficiencies (for example some of these inefficiencies may be caused by market nervousness ahead of major announcements). We would need an accurate information value model to be certain. Nonetheless we can be reassured that there is nothing in Fig. 7 that undermines or contradicts our expectations.

5.5. Comparison with volatility

Having seen that the measure of efficiency MOE increases with diversification and decreases when computed over longer periods, we can expect it to have an inverse relation with volatility. The standard deviation of returns, or volatility, reduces with diversification (unless the securities are perfectly correlated) and is proportional to the square root of the time interval (for the ideal case when returns follow a random walk).

Fig. 8 shows the relationships between the average intraday measure of efficiency MOE and the average intraday volatility of the Dow 30 stocks for two levels of transaction costs. Panel (a) shows the relationship if there are no transaction costs. Panel (b) shows the relationship when the round-trip transaction costs are at 1%. There is seemingly no correlation between MOE and volatility in the absence of transaction costs, and a negative correlation as expected when transaction costs are positive (and not unreasonably large). In the extreme case of huge transaction costs the market can be regarded as fully efficient because there would be hardly any trading and a wide band of transaction costs surrounding fundamental value.

We can see from Fig. 8 that although there is an inverse correlation between volatility and the measure of efficiency, they are not measuring the same phenomenon and do not proxy for each other.

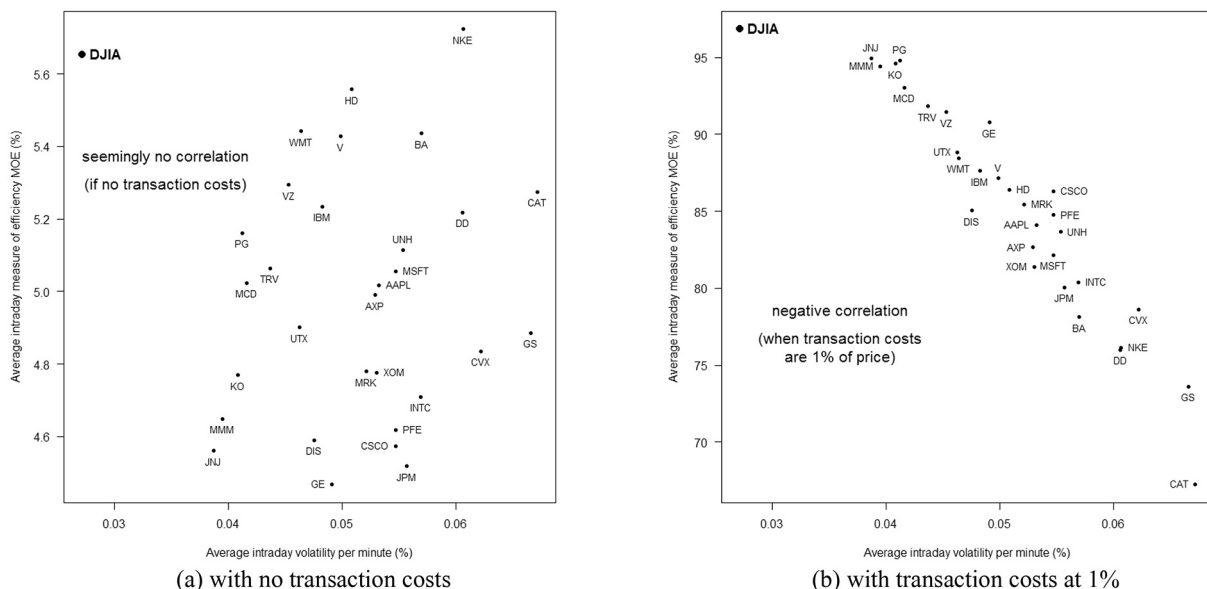


Fig. 8. Scatter plots of the measure of efficiency MOE versus volatility for the Dow Jones Industrial Average (DJIA) and its 30 component stocks for each day in 2016. The measure of efficiency MOE and volatility are calculated intraday from one-minute samples then averaged for each stock over the period. There were no changes to the composition of the Dow during 2016. Panels (a) and (b) show the scatter plots respectively for no transaction costs and for small transaction costs (1%). Panel (a) shows the MOE is uncorrelated with volatility in the absence of transaction costs, while panel (b) shows an inverse relation in the presence of ordinary transaction costs. We can also see in panel (b) the index DJIA has a lower volatility and higher measure of efficiency than each of its component stocks, as expected from diversification.

6. Applications

6.1. Stock market efficiency over time

While the purpose of Fig. 5 is to demonstrate the effect of diversification on the measure of efficiency, there is a further inference that can be drawn. There is an apparent reduction or downwards trend in the measure of efficiency of the stock market over time. Perhaps this is due to lower transaction costs over time and the increasing presence of high-frequency electronic trading, both of

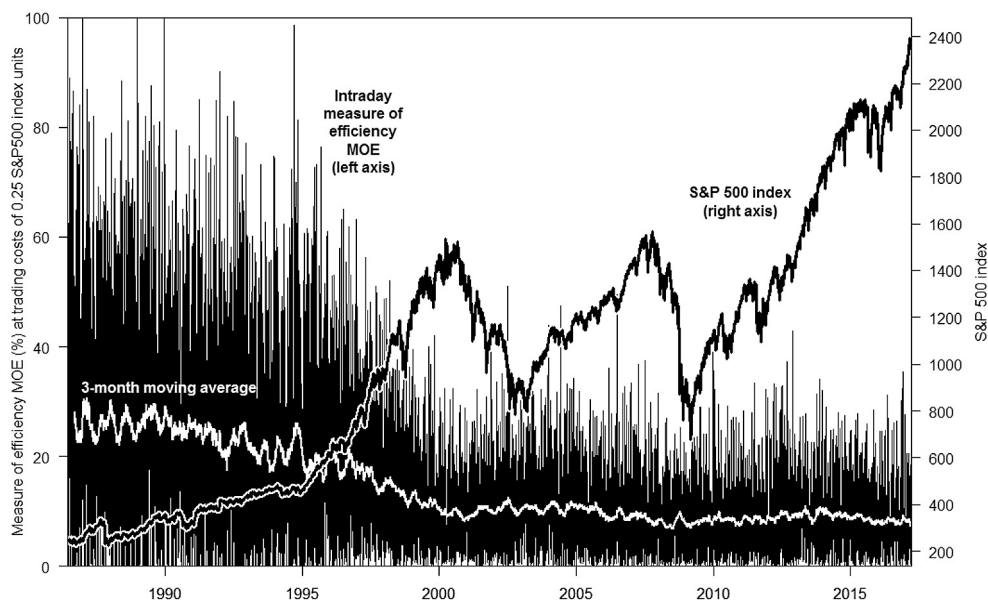


Fig. 9. Average intraday measure of efficiency MOE of the S & P 500 index over a 34-year period. The daily measure of efficiency is shown daily in black as the background noise and then smoothed into the white curve by a three-month moving average. The level of the S & P 500 index is shown in the black curve. It seems that the measure of efficiency has reduced over time, particularly in the decade from 1990 to 2000, and this appears unrelated to the level of the index itself.

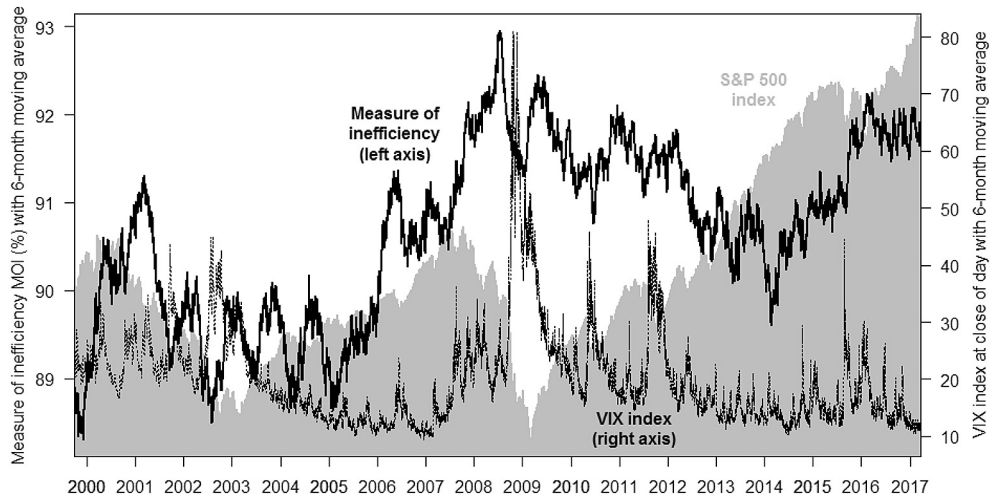


Fig. 10. Comparison of the measure of inefficiency MOI (left axis) with the underlying S & P 500 index (unscaled) and the VIX or fear index (right axis). The measure of inefficiency has been computed with transaction costs at 0.25 index points which is the minimum bid-ask spread for S & P futures quotes, and smoothed with a six-month moving average to mask the daily variation. Despite the delay imposed by the smoothing, the measure rises through 2007 to a peak in mid-2008 which is several months ahead of the crash. By comparison, the spike in the VIX index occurs during the crash in October 2008. The measure of inefficiency therefore appears to be a leading indicator in this case while the fear index has been a reaction.

which would allow for more noise to occur at higher frequencies.

To test this further, Fig. 9 shows the intraday measure of efficiency of the broad S & P 500 index over 34 years from 1983 to March 2017. The daily intraday measure of efficiency is the noisy signal shown in black, and superimposed upon it in white is its three-month moving average, with percentages shown on the left axis. The black curve is the level of the S & P 500 index, shown on the right axis. We see that the measure of efficiency MOE appears to reduce over time and particularly prior to 2001. These moves seem independent of the vicissitudes of the index itself. Overall it looks like the stock market offers more inefficiencies now than it did in the last century.

6.2. Inefficiency of the S & P 500 index before the financial crisis of 2008

If a financial crash is a consequence of a build-up of inefficiencies, it should be possible to observe a higher level of inefficiencies being measured prior to a significant price correction. This is analogous to our previous experiment where we saw a relation on a daily scale between one day's measure of inefficiency and the return the next day. We are investigating a similar concept now but on a larger time scale.

Fig. 10 shows the intraday measure of inefficiency MOI (where $MOI = 1 - MOE$) of the S & P 500 index from 2000 to 2017. We can see that the measure of inefficiency grows and peaks in the lead-up to the crash of October 2008. Even with the smoothing of a six-month moving average, the filtered measure of inefficiency still reaches its peak a few months prior to the crash. The VIX index is plotted for comparison on the right axis and the underlying S & P 500 index level is shown in grey without an axis. While the VIX is interpreted as a fear index (Whaley, 2009) it can be seen reacting to the downturn rather than being a leading indicator. The measure of inefficiency seems more of a leading indicator because it rises in 2007 in the lead up to the crash and then declines as the correction occurs in 2008.

Although the MOI is a leading indicator of the 2008 crash, it is unlikely to be a causal factor. This crash developed from an excess of subprime mortgage lending, and the signs were visible to those who understood where to look. For an analysis of the 2007–2008 crisis see Demanyk and Van Hemert (2011). The rising inefficiency measured in the lead up to the crisis is more likely to be a consequence of the developing situation rather than a cause, for example by reflecting a growing divergence of beliefs. We can conjecture a variety of effects that the MOI is detecting, especially when applied to a diversified index. The synchronicity of stock returns differs significantly between countries as shown by Morck et al. (2000), and this will affect the opportunities for statistical arbitrage on each index. There is scope for future research to measure the MOIs of individual stocks and compare these with the MOI of their index, and to repeat the analysis for indices around the world. This is an example of how the ability to measure efficiency numerically and to answer questions of “how much” in each setting could help us move towards exploring the questions of “why”.

Investigating further to test whether crashes in general are preceded by inefficiency would require trading data from a large number of financial crises, ideally with a magnitude similar to 2008. The necessary high-resolution intraday data required for computing MOE and MOI has become available only in the last 10 to 20 years, which means we can use this technique to study only the most recent crises. By contrast the classic tests of serial correlation can be conducted using older data at daily resolution. Lim et al. (2008) study the Asian crisis of 1997 and compare the behavior of eight Asian stock markets. Kim et al. (2011) study the Dow Jones Industrial Average long term from 1900 through 2009 and conclude there is greater return predictability during economic or political crises. The problem with serial correlation at daily resolution is the endogeneity with the price momentum that occurs naturally

during a crash, an alternate explanation for the greater return predictability. We can also expect that the differential in serial correlation between crisis and non-crisis periods will be exacerbated by the use of overlapping rolling windows. By contrast the MOE and MOI calculations are free from these effects and they are calculated intraday with no rolling windows. We therefore await more crashes to study with high-resolution data.

While we are keen to study more crashes with intraday data, we would hope that such crashes will not recur in our lifetimes. Ironically if they were to become a regular feature of our markets, this would introduce a major inefficiency because the oscillations would then be foreseeable and tradeable by mean reversion.

7. Conclusions

This article presents a measure of efficiency that is a proxy for market efficiency without requiring a model of information value. Intuitively it compares the potential profits from passive and active trading. When a market is efficient in valuing price-sensitive information, it leaves no further room for arbitrageurs to make abnormal profits, and the active traders cannot beat the passive traders. When a market is inefficient, the active trading opportunities should exceed those of passive trading, and the extent of their advantage gives a measure of the inefficiencies.

The measure of efficiency MOE is computed as the ratio of the maximum profits that could have been made by passive trading to those of active trading over the time period of interest and for a specified level of transaction costs. It is a scale rather than a binary test, with a range from 0% (fully inefficient) to 100% (fully efficient). There is a complementary measure of inefficiency MOI with the reverse scale, so the two measures sum to 100%.

When transaction costs are high, there are fewer opportunities to profit, and the MOE is high. When costs are low, mean-reverting strategies can trade profitably on smaller price oscillations and the MOE is lower. The MOE also increases with diversification, decreases with the time interval over which it is measured, and has an inverse relationship with volatility. It is shown to be higher ahead of price corrections, both on a short-term day-to-day scale, and the longer term pattern around the 2008 financial crisis.

The measure of efficiency MOE is a proxy for market efficiency that can be calculated from readily available price information, and it exhibits many of the properties that we require. While debate will inevitably continue as to whether markets are efficient, numerical measures such as this should help facilitate research questions of how efficient markets are, rather than whether or not they are efficient to particular information sets, which will enable relative comparisons between different markets and different time periods.

Acknowledgements

This research has been funded by The University of Western Australia, with security trading data provided by SIRCA. I am grateful for discussions and feedback from Philip Brown, Raymond Da Silva Rosa, Tom Smith, and Robert Faff.

Appendix A. Mathematical definitions of MPAT and MPPT

Consider a price series S_t defined over time period $t_A \leq t \leq t_B$ with round-trip transaction cost $\delta \geq 0$. The maximum profit from passive trading (MPPT) is

$$MPPT = \max(0, |S_{t_B} - S_{t_A}| - \delta) \tag{1}$$

and the maximum profit from active trading (MPAT) is:

$$MPAT = \max\left(0, \sup_{t_A=t_0 \leq t_1 \leq \dots \leq t_N=t_B} \left\{ \sum_{i=1}^N |S_{t_i} - S_{t_{i-1}}| - N\delta \right\}\right) \tag{2}$$

The measure of efficiency is then

$$MOE = \begin{cases} \frac{MPPT}{MPAT} & \text{if } MPAT > 0 \\ 1 & \text{if } MPAT = 0 \end{cases} \tag{3}$$

and the measure of inefficiency is

$$MOI = \begin{cases} \frac{MPAT - MPPT}{MPAT} & \text{if } MPAT > 0 \\ 0 & \text{if } MPAT = 0 \end{cases} \tag{4}$$

Appendix B. Proofs of useful properties

The properties of the maximal profits from passive and active trading are proven as follows:

- (i) $MPPT \geq 0$

$$\begin{aligned} \text{MPPT} &= \max(0, |S_{t_B} - S_{t_A}| - \delta) \\ &\geq \max(0) && \text{because } \max(x, y, z) \geq x \\ &= 0 \end{aligned}$$

(ii) $\text{MPAT} \geq 0$

$$\begin{aligned} \text{MPAT} &= \max\left(0, \sup_{t_A=t_0 \leq t_1 \leq \dots \leq t_N=t_B} \left\{ \sum_{i=1}^N |S_{t_i} - S_{t_{i-1}}| - N\delta \right\}\right) \\ &\geq \max(0) && \text{because } \max(x, y, z) \geq x \\ &= 0 \end{aligned}$$

(iii) $\text{MPAT} \geq \text{MPPT}$

$$\begin{aligned} \text{MPAT} &= \max\left(0, \sup_{t_A=t_0 \leq t_1 \leq \dots \leq t_N=t_B} \left\{ \sum_{i=1}^N |S_{t_i} - S_{t_{i-1}}| - N\delta \right\}\right) \\ &\geq \max\left(0, \sum_{i=1}^N |S_{t_i} - S_{t_{i-1}}| - N\delta\right) \text{ where } N = 1, t_0 = t_A, \text{ and } t_1 = t_B \\ &\quad \text{because } (t_A, t_B) \text{ is one eligible subset of } (t_0, t_1, \dots, t_N) \\ &= \max(0, |S_{t_B} - S_{t_A}| - \delta) \\ &= \text{MPPT} \end{aligned}$$

The properties of the measures of efficiency and inefficiency are proven as follows:

(iv) $0 \leq \text{MOE} \leq 1$

$$\begin{aligned} \text{MOE} &= \begin{cases} \frac{\text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\geq \begin{cases} 0 & \text{by properties (i), (ii), (iii) if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\geq 0 \\ \text{MOE} &= \begin{cases} \frac{\text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\leq \begin{cases} 1 & \text{by properties (i), (ii), (iii) if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\leq 1 \end{aligned}$$

(v) $0 \leq \text{MOI} \leq 1$

$$\begin{aligned} \text{MOI} &= \begin{cases} \frac{\text{MPAT} - \text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\geq \begin{cases} 0 & \text{by properties (i), (ii), (iii) if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\geq 0 \\ \text{MOI} &= \begin{cases} \frac{\text{MPAT} - \text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\leq \begin{cases} 1 & \text{by properties (i), (ii), (iii) if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &\leq 1 \end{aligned}$$

(vi) $\text{MOE} + \text{MOI} = 1$

$$\begin{aligned} \text{MOI} &= \begin{cases} \frac{\text{MPAT} - \text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &= \begin{cases} 1 - \frac{\text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 1 & \text{if } \text{MPAT} = 0 \end{cases} \\ &= 1 - \begin{cases} \frac{\text{MPPT}}{\text{MPAT}} & \text{if } \text{MPAT} > 0 \\ 0 & \text{if } \text{MPAT} = 0 \end{cases} \\ &= 1 - \text{MOE} \end{aligned}$$

Appendix C. Software implementation and R code

This appendix contains example code for computing MPAT, MPPT, MOE, and MOI using the R statistical programming language (<http://cran.r-project.org>). The inputs are a discrete price sequence in a real-valued array $S[i]$ for $i = 1, 2, 3, \dots, N$ and the transaction costs as number C . The code in the R language is:

```
SA <- S[1]
SB <- S[1]
MPAT <- 0
if( N>=2 ) for( i in 2:N ) {
  if( (S[i]-SB) * (SB-SA) > 0 ) {
    SB <- S[i]
  }
  if( abs(S[i]-SB) > C ) {
    if( SB != SA ) {
      MPAT <- MPAT + abs(SB-SA) - C
    }
    SA <- SB
    SB <- S[i]
  }
}
MPAT <- MPAT + max( 0, abs(S[N]-SA)-C )
MPPT <- max( 0, abs(S[N]-S[1])-C )
if( MPAT == 0 ) {
  QE <- 1
}
else {
  QE <- MPPT / MPAT
}
QI <- 1 - QE
```

References

- Bachelier, L., 1900. Théorie de la spéculation. *Ann. Sci. de l'E.N.S.* 3 (17), 21–86.
- Ball, R., Brown, P., 1968. An Empirical Evaluation of Accounting Income Numbers. *J. Account. Res.* 6 (2), 159–178.
- Ball, R., Kothari, S.P., 1989. Nonstationary expected returns: Implications for tests of market efficiency and serial correlation in returns. *J. Financ. Econ.* 25, 51–74.
- Ball, R., Kothari, S.P., Wasley, C.E., 1995. Can we implement research on stock trading rules? *J. Portf. Manag.* 21 (2), 54–63.
- Barberis, N., Thaler, R.H., 2003. A survey of behavioural finance. In: Constantinides, G.M., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance*. Elsevier Science BV, pp. 1052–1121 (Ch 18).
- Black, F., 1986. Noise. *J. Financ.* 41 (3), 529–543.
- Buffett, W.E., 1989. Chairman's letter to the shareholders of Berkshire Hathaway Inc. <http://www.berkshirehathaway.com/letters/1988.html>.
- Conrad, J., Gultekin, M.N., Kaul, G., 1997. Profitability of short-term contrarian strategies: implications for market efficiency. *J. Bus. Econ. Stat.* 15 (3), 379–386.
- Demyanyk, Y., Van Hemert, O., 2011. Understanding the Subprime Mortgage Crisis. *Rev. Financ. Stud.* 24, 1848–1880.
- Fama, E.F., 1965a. The behavior of stock-market prices. *J. Bus.* 38 (1), 34–105.
- Fama, E.F., 1965b. Random walks in stock market prices. *Financ. Anal. J.* 21 (5), 55–59.
- Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. *J. Financ.* 25 (2), 383–417.
- Fama, E.F., 1991. Efficient capital markets: II. *J. Financ.* 46 (5), 1575.
- Fama, E.F., 1998. Market efficiency, long-term returns, and behavioral finance. *J. Financ. Econ.* 49, 283–306.
- Fama, E.F., Fisher, L., Jensen, M.C., Roll, R., 1969. The Adjustment of Stock Prices to New Information. *Int. Econ. Rev.* 10 (1), 1–21.
- Gemmill, G., Thomas, D.C., 2002. Noise trading, costly arbitrage, and asset prices: evidence from closed-end funds. *J. Financ.* 57, 2571–2594.
- Griffin, J.M., Kelly, P.J., Nardari, F., 2010. Do market efficiency measures yield correct inferences? A comparison of developed and emerging markets. *Rev. Financ. Stud.* 23 (8), 3225–3277.
- Ito, M., Sugiyama, S., 2009. Measuring the degree of time varying market inefficiency. *Econ. Lett.* 103, 62–64.
- Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *J. Financ.* 48 (1), 65–91.
- Jensen, M.C., 1978. Some anomalous evidence regarding market efficiency. *J. Financ. Econ.* 6, 95–101.
- Kang, J., Liu, M.H., Ni, S.X., 2002. Contrarian and momentum strategies in the China stock market: 1993–2000. *Pac. Basin Financ. J.* 10, 243–265.
- Kim, J.H., Shamsuddin, A., Lim, K.-P., 2011. Stock return predictability and the adaptive markets hypothesis: evidence from century-long U.S. data. *J. Empir. Financ.* 18, 868–879.
- Lehmann, B.N., 1990. Fads, martingales, and market efficiency. *Q. J. Econ.* 105 (1), 1–28.
- Lim, K.-P., Brooks, R.D., Kim, J.H., 2008. Financial crisis and stock market efficiency: empirical evidence from Asian countries. *Int. Rev. Financ. Anal.* 17, 571–591.
- Lo, A.W., MacKinlay, A.C., 1988. Stock market prices do not follow random walks: evidence from a simple specification test. *Rev. Financ. Stud.* 1 (1), 41–66.
- Mandelbrot, B.B., 1982. *The Fractal Geometry of Nature*. W.H. Freeman and Company.
- Markowitz, H., 1952. Portfolio selection. *J. Financ.* 7 (1), 77–91.
- Morck, R., Yeung, B., Yu, W., 2000. The information content of stock markets: why do emerging markets have synchronous stock price movements? *J. Financ. Econ.*

- 58, 215–260.
- Narayan, P.K., Narayan, S., Zheng, X., 2010. Gold and oil futures markets: are markets efficient? *Appl. Energy* 87 (10), 3299–3303.
- Rosenthal, L., Young, C., 1990. The seemingly anomalous price behavior of Royal Dutch/Shell and Unilever N.V./PLC. *J. Financ. Econ.* 26 (1), 123–141.
- Samuelson, P., 1965. Proof that properly anticipated prices fluctuate randomly. *Ind. Manag. Rev.* 6 (2).
- Scruggs, J.T., 2007. Noise trader risk: evidence from the Siamese twins. *J. Financ. Mark.* 10 (1), 76–105.
- Sewell, M., 2011. History of the efficient market hypothesis. In: Research Note RN/11/04. University College London Department of Computer Science.
- Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. *J. Financ.* 52 (1), 35–55.
- Small, K., Smith, J., 2007. The hot stock tip from Debbie: implications for market efficiency. *J. Behav. Financ.* 8 (4), 191–197.
- Smith, A.D., 2001. In: Investigating some measures of market inefficiency - a technical note. Finance and Investment Conference: Risk, Reward, and Regulation. The Actuarial Profession. (<https://www.actuaries.org.uk/documents/paper-2-investigating-some-measures-market-efficiency> (accessed 24 March 2017)).
- Stephens, C.R., Benink, H.A., Gordillo, J.L., Pardo-Guerra, J.P., 2007. A new measure of market inefficiency. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1009669.
- Sullivan, R.N., 2010. Editor's corner: Competence trumps style. *Financ. Anal. J.* 66 (1), 6–7.
- Urquhart, A., Hudson, R., 2013. Efficient or adaptive markets? Evidence from major stock markets using very long run historic data. *Int. Rev. Financ. Anal.* 28, 130–142.
- Wang, Y., Liu, L., Gu, R., Cao, J., Wang, H., 2010. Analysis of market efficiency for the Shanghai stock market over time. *Phys. A* 389, 1635–1642.
- Whaley, R.E., 2009. Understanding the VIX. *J. Portf. Manag.* 35 (3), 98–105.
- Worthington, A.C., Higgs, H., 2009. Efficiency in the Australian stock market 1875–2006: a note on extreme long-run random walk behavior. *Appl. Econ. Lett.* 16 (3), 301–306.