Cambridge University Press
0521816033 - Information Efficiency in Financial and Betting Markets
Edited by Leighton Vaughan Williams
Excerpt
More information

# Introduction

A central issue in the analysis of markets is the degree to which they are *efficient*. Although 'efficiency' has a variety of meanings in different contexts, a situation is sometimes termed 'efficient' if it is not possible to increase the well-being (utility) of any one person without reducing the utility of another. This is usually referred to as Pareto efficiency. An implication of Pareto efficiency is *productive efficiency*, a situation which exists when it is not possible to increase the quantity produced of any one good without reducing the quantity produced of another.

In the analysis of betting markets – and, indeed, financial markets more generally – however, the examination of efficiency assumes an informational dimension, the existence of which may well be related to that of Pareto or productive efficiency, but the meaning of which is quite distinct. It is this form of efficiency which is the subject of investigation in this volume. This book traces the development of the idea of informationally efficient markets, and identifies the various precise definitions and variations of the concept extant in the literature on financial markets. The theoretical background is clarified, and empirical tests of information efficiency are reviewed and evaluated.

While most studies of information efficiency are conducted within the framework of conventional financial markets, there are a number of special features of betting markets which warrant particular attention and make them of unique relevance to a study of market efficiency. In particular, these markets not only possess many of the usual attributes of financial markets – notably a large number of investors (or bettors) with potential access to widely available rich information sets – but also the important additional property that each asset (or bet) possesses a well-defined end point at which its value becomes certain. This contrasts with most financial markets, where the value of an asset in the present is dependent both on the present value of future cash flows and also on the uncertain price at which it can be sold at some future point in time.

The defined termination point of betting markets is of particular appeal, therefore, in that it allows researchers employing empirical techniques to avoid many of the difficulties associated with indefinite expected future outcomes. Moreover, by enabling a more productive and clearer learning process, a delineated end point might be expected in particular to promote information efficiency. Evidence of inefficiency in such markets is therefore of special significance. The possibility of insider information and consequent opportunities for insider trading in betting markets is also somewhat analogous to the operation of conventional financial markets, but in some respects easier to measure and assess. For these reasons, the information provided by an examination of betting markets is a convenient and useful perspective from which to consider the evidence and interpretations of consumer and investor behaviour in conventional financial markets, as well as the operation of these markets.

This volume has a two-tiered structure. Part I consists of three chapters. Chapter 1 reviews the academic literature which has investigated the issue of information efficiency in conventional financial markets. The development of the idea of an informationally efficient market is explored, and the various classifications of this issue are identified. Empirical tests of information efficiency in these markets are assessed and evaluated. Chapters 2 and 3 review the academic literature which has investigated the issue of information efficiency in betting markets. The various empirical tests which have been applied in this area are assessed and evaluated. Part II is a collection of hitherto unpublished readings which draws on expertise across the spectrum of research into the issue of information efficiency in betting markets. Each of the contributions is novel and original, but set within the existing framework of literature. As such, this volume will serve as a valuable asset for those who are coming fresh to the subject, as well as for those who are more familiar with the subject matter.

I have greatly enjoyed writing this book, and editing the collection of readings. In great part, this is due to the kindness, support and generosity of family, friends and of colleagues from across the global village of academic research. Special thanks also to all who have contributed to this book. In every case, the contribution offers a new and valuable insight into this fascinating subject.

Welcome to the wonderful world of information efficiency!

# Part I
# The concept of information efficiency

# 1 Information efficiency in financial markets

*Leighton Vaughan Williams*

## 1.1 Introduction

This chapter examines some of the basic issues relating to the theory of information efficiency in financial markets and, in particular, some of the definitions and distinctions which have influenced the academic literature to date. Various empirical tests of information efficiency are then reviewed and assessed.

Section 1.2 outlines the concept of information efficiency and traces the development of the terms, definitions and meanings associated with this idea. Sections 1.3, 1.4 and 1.5 review the methods which have been applied to test for the existence of information efficiency, as variously defined, in financial markets.

## 1.2 The 'efficient markets hypothesis'

In this section, a review is undertaken of the literature which has investigated the concept and existence of information efficiency in financial markets, and in particular the role and relevance of the 'efficient markets hypothesis' in our understanding of the operation of these markets.

### *1.2.1 The efficient markets hypothesis: reviewing the development of an idea*

The concept of information efficiency in a market is contained in the so-called 'efficient markets hypothesis', a standard definition of which can be found in Fama (1991): 'I take the market efficiency hypothesis to be the simple statement that security prices fully reflect all available information' (1991: 1575).

The origin of the ideas central to this hypothesis can be traced back to pioneering work undertaken by Bachelier (1900) into the dynamics of stock price behaviour. His examination of the behaviour of securities prices on La Bourse (the Paris Stock Exchange) led him to conclude that the price changes were identically and independently distributed, so that the next movement in a particular time series could not be predicted from an examination of previous movements. In particular, the stochastic process employed by Bachelier to describe such stock price changes has the characteristic that increments in the process are the result of independent random variables, are normally distributed with a zero mean, and possess a variance increasing in proportion to time elapsed. The implication is that stock prices have no memory and, having no systematic tendencies, cannot be exploited by arbitrage. This proposition that stock price movements observe a normal distribution, and that the price changes follow a 'random walk', laid the basis of much subsequent work into what has come to be known as 'efficient markets theory'.

Kendall (1953), for example, analysed serial correlations in the behaviour of weekly changes in spot prices for wheat, cotton and nineteen indices of UK industrial share prices. His conclusion was that the series appeared 'wandering', 'Almost as if once a week the Demon of Chance drew a random number from a symmetrical population of fixed dispersion and added it to the current price to determine the next week's price' (1953: 13).

A serious challenge to this orthodoxy can be traced to Mandelbrot (1963), whose analysis of the actual distribution of price changes disclosed evidence of high-tail distributions without a finite variance. This work served to cast doubt on the value of the existing standard statistical techniques such as serial correlation analyses to test for dependence, and generated a whole new literature proposing and applying new techniques to test for such dependence.

Another important development in the literature since the late 1950s has been the clarification of hitherto implicit distinctions. In particular, the concepts of a random walk, a 'fair game' and the various 'martingale'[1] specifications are now clearly contrasted. Basically, if prices follow a stochastic process, then this can be identified as a martingale if the best forecast of tomorrow's prices that can be made, based on present information, is today's price. Likewise, the stochastic process is identified as a fair game if the expected gain from forecasting tomorrow's price based on present information is zero, i.e. there is no systematic difference between actual and expected returns. The implication of the above is that if a variable in an investor's information set can be used to predict future returns, then the martingale model is violated, and returns cannot follow a fair game. The stochastic process is identified as a random walk if it satisfies the martingale conditions and also that there is no dependence involving the higher conditional moments of future prices. The random walk specification is, therefore, more restrictive than the martingale. These issues are addressed in more detail in subsection 1.2.2.

The possibility that market inefficiency can exist independently of price dependence, however categorised, can be traced to definitions originally associated with Roberts (1959, 1967), and popularised by Fama (1970), i.e. 'weak form', 'semi-strong form' and 'strong form' efficiency. The idea is that the existence of market efficiency may best be examined in terms of three distinct types of test, each subjecting the efficient markets hypothesis to different levels of strictness.

Fama (1970) discussed the tests in terms of the information subset relevant to changes in security prices. First, weak form tests which are concerned with the information set of historical prices. Second, semi-strong form tests, which are concerned with 'information that is obviously publicly available' (1970: 383). Third, strong form tests 'concerned with whether given investors or groups have monopolistic access to any information relevant for price formation' (1970: 383). The three tests seek to identify which subset of information is relevant in the formation of expectations, and thereby security prices. Weak form information is limited to the price history of the relevant security; semi-strong information is limited to publicly available information; strong form information includes all known relevant information, including private information. These issues are explored in greater detail in subsection 1.2.3.

### *1.2.2 Random walks, fair games and martingales*

The idea that the absence of a random walk by financial variables is sufficient in itself to reject the existence of information efficiency in the relevant financial markets was challenged by Fama (1965). He produced findings that larger than average daily stock price changes in his dataset tended to be followed by larger than average daily price changes. However, the signs of the successor changes appeared random. He concluded that although this represented a contradiction of a random walk by these variables, it did not contradict the existence of information efficiency in the markets exhibiting these characteristics.

This distinction was developed by Fama (1970), where he differentiated between a random walk and a fair game, arguing that a fair game assumption is sufficient for information efficiency, but that a fair game formulation is not sufficient in itself to lead to a random walk. In so doing, he echoed Alexander's (1961) contention that assuming a 'fair game' would take one 'well on the way to picturing the behaviour of speculative prices as a random walk' (1961: 200).

LeRoy (1989) offered a clear presentation of these sorts of distinctions. Specifically, he identified a stochastic process $x_t$ as a martingale[2] with respect to a sequence of information sets $I_t$, if $x_t$ has the property

$$\mathrm{E}(X_{t+1} \text{given } I_t) = X_t$$

Where E ($n$) represents the expected value of $n$.

So, in assuming that $x_t$ is in $I_t$, then if $x_t$ is a martingale, the best forecast of $x_{t+1}$ based on current information $I_t$ would be $x_t$. If the process is a fair game, then the expected gain from forecasting $x_{t+1}$ based on current information $I_t$ is zero.

The implication of the above is that if a variable in an investor's information set can be used to predict future returns the martingale model is violated, and returns cannot follow a fair game. A stochastic process is identified as a random walk if it satisfies the martingale conditions and also that there is no dependence involving the higher conditional moments of $x_{t+1}$. If, for instance, we model security price behaviour in such a way that successive conditional variances of such prices are positively autocorrelated (though not their levels), then this satisfies the martingale conditions, but not the random walk. The existence of risk-neutrality, in which investors are unconcerned about the higher moments of their return distributions, points therefore to a martingale formulation but not a random walk, since investors in such a scenario are not led to bid away serial dependence in these higher conditional moments. The presence of risk aversion, on the other hand, runs contrary to a martingale and a fair game modelling. The reason stems from the fact that risk-averse investors will only hold more risky assets if they are compensated in terms of higher expected returns. As a consequence, knowledge of the riskiness of the current information set implies some knowledge about the level of expected returns. The idea of a submartingale is that expected rates of return (ignoring dividends), conditional on currently available information, are non-negative, *i.e.*

$$\mathrm{E}(p_{\mathrm{t}+1} \text{given } I_{\mathrm{t}}) >= p_{\mathrm{t}}$$

which implies that no trading rule based on the current information set can outperform a strategy of buy-and-hold.

Granger (1992) pointed out that if stock prices were not a martingale, then ignoring transactions costs 'price changes would be consistently forecastable and so a money machine is created and indefinite wealth is created' (1992: 3). Granger took care to differentiate, therefore, between a martingale process and the various interpretations identified with the efficient market hypothesis, expressing his own preference for Jensen's (1978) definition – i.e. that a market is efficient with respect to a given information set if it is impossible to make economic profits[3] by trading on the basis of this information.[4]

Support for Jensen's definition is offered by Fama (1991), in a follow-up to his original 1970 survey of the literature on efficient capital markets. Fama (1991) noted Grossman and Stiglitz's (1980) finding that for security prices to reflect fully all available information then information and trading costs must be zero. Finding this implausible, he preferred Jensen's 'weaker and more sensible version of the efficiency hypothesis [which] says that prices reflect information to the point where the marginal benefits of acting on information (the profits to be made) do not exceed the marginal costs' (1991: 1575).

A related issue is raised by Keane (1993), who highlighted a distinction between rationality and exploitability as aspects of pricing efficiency. For Keane (1993), the market is rational if prices and market movements reflect the best estimates of intrinsic values. It is fair game efficient or non-exploitable if systematic abnormal returns cannot be earned through an analysis of price behaviour. The distinction is made clear in a situation where the market in aggregate is subject to excessive movements that are difficult to identify or are unpredictable in behaviour. In such a situation, irrational market behaviour can co-exist with fair game efficiency or non-exploitability.

The essential issues can, however, be categorised into two parts. First, is there evidence in financial markets of price change dependence as variously defined? Second, can any such evidence be used to secure systematic abnormal returns?

### *1.2.3 Weak, semi-strong and strong form efficiency: classifications of information efficiency*

The weak form of the efficient markets hypothesis holds that current security prices fully and instantaneously reflect all weak form information, and similarly for the semi-strong and strong forms of the hypothesis. In a weak form market it follows that no patterns can be identified which would allow future price movements to be predicted from past price movements, and no trading rule will produce consistent above-average or abnormal returns except by chance. Prices are influenced solely by new economic events and new information. Fama (1991) has proposed extending the categorisation of research in this area to include such variables as dividend yields, interest rates, earnings/price ratios and other term-structure variables. Fama identifies these as tests for return predictability, a more general category which includes weak form tests. In a semi-strong form market, new public information impacts on security prices instantaneously and in an unbiased fashion. Such prices, therefore, most faithfully reflect the available published information. In a strong form market, share prices reflect all information, including that not publicly available.[5]

Dowie (1976) made a basic distinction between the strong form of inefficiency (as hitherto defined) and the other forms of inefficiency (weak and semi-strong). The former tells us about access to and the availability of information, whereas the latter is concerned with how well the market responds to information. Although related, these are quite separate issues. Since strong inefficiency implies the existence of subsets of investors who possess monopolistic access to information (which can be exploited to earn above-average returns), Dowie uses the term ‘equitable’ to describe markets which pass the strong test, and ‘efficient’ to describe those that pass the weak and semi-strong tests.

Keane (1987) also made a clear distinction. Whereas the weak and semi-strong classifications apply to the stock market itself, strong efficiency, he argues, is about a broader concept of capital markets. Specifically, whereas ‘semi-strong efficiency is concerned with how well the market processes the information disclosed to it . . . strong efficiency is concerned primarily with the adequacy of the information disclosure process’ (1987: 6). In this sense, it might be considered misleading to view strong efficiency as a progression from the weak and semi-strong forms, since this confuses the ability of the market to respond to and interpret information with the failure of the market to supply information (what we might call the information production function).

It can be seen that the development of research into information efficiency in recent years has sought to clarify the nature of the distribution of stock price changes, and in this context to develop statistical tests which offer the possibility of testing for dependence between successive price changes. The type and degree of dependence under examination has been clarified, and the concept of information efficiency itself has been broadened and made more explicit.

### *1.2.4 The efficient markets hypothesis: a summary*

An informationally efficient market can in essence be defined as a market which *incorporates all information*. This is a stringent requirement, and so studies of financial markets have also addressed the issue with respect to subsets of the totality of information. The three principal (though not exclusive) levels at which studies of information efficiency have been undertaken are with respect to weak, semi-strong and strong information.

Weak form information is information contained in the set of historical prices. A market is weakly efficient (with respect to information), therefore, if this is fully and (in the strictest form) instantaneously incorporated in present prices. In such a market, present prices reflect all information available in patterns of historical prices, and so future price movements cannot be derived from an examination of past prices.

Semi-strong information is that contained in the set of all public information. A market is semi-strong efficient if this is fully and (in the strictest form) instantaneously incorporated into present prices. In such a market, present prices reflect all available public information, and so future price movements reflect future (and as yet unknown) revelations of publicly available information.

Strong information is that contained in the set of all information, including that privately and monopolistically held. A market is strongly efficient if all information is fully and (in the strictest form) instantaneously incorporated into present prices. In such a market, present prices reflect all information, and so future price movements reflect future (and as yet unknown) revelations of information.

All these definitions of information efficiency require the incorporation of relevant information. In less strict formulations, it is sufficient for efficiency to exist that it is not possible to trade upon this information so as to earn greater than normal profits.

## 1.3 Empirical tests of weak form information efficiency in financial markets

This section reviews some of the empirical tests which have been proposed and applied in the literature to investigate the existence of weak information efficiency in financial markets.

It has already been shown that in a financial market characterised by strict weak form efficiency, no patterns can be identified from the history of price data which would allow one to predict the future pattern of price changes. In a market which is weakly inefficient as so defined, the pattern of incremental prices is well approximated by a random walk specification. Subsections 1.3.1–1.3.4 review the evidence for such a specification: 1.3.1 assesses serial correlation techniques of price dependence, 1.3.2 variance ratio tests, 1.3.3 cointegration approaches and 1.3.4 looks briefly at how rescaled range analysis and chaos theory have been applied to the theory of financial markets.

A less strict form of weak efficiency holds that no information can be gathered from such price data which would allow one to make abnormal returns except by chance. In subsection 1.3.5, a review is undertaken of