
MARKET EFFICIENCY – DEFINITION, TESTS AND EVIDENCE

What is an efficient market? What does it imply for investment and valuation models? Clearly, market efficiency is a concept that is controversial and attracts strong views, pro and con, partly because of differences between individuals about what it really means, and partly because it is a core belief that in large part determines how an investor approaches investing. This chapter provides a simple definition of market efficiency, considers the implications of an efficient market for investors and summarizes some of the basic approaches that are used to test investment schemes, thereby proving or disproving market efficiency. It also provides a summary of the voluminous research on whether markets are efficient.

Market Efficiency and Investment Valuation

The question of whether markets are efficient, and if not, where the inefficiencies lie, is central to investment valuation. If markets are, in fact, efficient, the market price provides the best estimate of value, and the process of valuation becomes one of justifying the market price. If markets are not efficient, the market price may deviate from the true value, and the process of valuation is directed towards obtaining a reasonable estimate of this value. Those who do valuation well, then, will then be able to make 'higher' returns than other investors, because of their capacity to spot under and over valued firms. To make these higher returns, though, markets have to correct their mistakes – i.e. become efficient – over time. Whether these corrections occur over six months or five years can have a profound impact in which valuation approach an investor chooses to use and the time horizon that is needed for it to succeed.

There is also much that can be learnt from studies of market efficiency, which highlight segments where the market seems to be inefficient. These 'inefficiencies' can provide the basis for screening the universe of stocks to come up with a sub-sample that is

more likely to have under valued stocks. Given the size of the universe of stocks, this not only saves time for the analyst, but increases the odds significantly of finding under and over valued stocks. For instance, some efficiency studies suggest that stocks that are 'neglected' by institutional investors are more likely to be undervalued and earn excess returns. A strategy that screens firms for low institutional investment (as a percentage of the outstanding stock) may yield a sub-sample of neglected firms, which can then be valued using valuation models, to arrive at a portfolio of undervalued firms. If the research is correct the odds of finding undervalued firms should increase in this sub-sample.

What is an efficient market?

An efficient market is one where the market price is an unbiased estimate of the true value of the investment. Implicit in this derivation are several key concepts -

(a) Contrary to popular view, market efficiency does not require that the market price be equal to true value at every point in time. All it requires is that errors in the market price be unbiased, i.e., that prices can be greater than or less than true value, as long as these deviations are random¹.

(b) The fact that the deviations from true value are random implies, in a rough sense, that there is an equal chance that stocks are under or over valued at any point in time, and that these deviations are uncorrelated with any observable variable. For instance, in an efficient market, stocks with lower PE ratios should be no more or less likely to be under valued than stocks with high PE ratios.

(c) If the deviations of market price from true value are random, it follows that no group of investors should be able to consistently find under or over valued stocks using any investment strategy.

¹ Randomness implies that there is an equal chance that stocks are under or over valued at any point in time.

Definitions of market efficiency have to be specific not only about the market that is being considered but also the investor group that is covered. It is extremely unlikely that all markets are efficient to all investors, but it is entirely possible that a particular market (for instance, the New York Stock Exchange) is efficient with respect to the average investor. It is also possible that some markets are efficient while others are not, and that a market is efficient with respect to some investors and not to others. This is a direct consequence of differential tax rates and transactions costs, which confer advantages on some investors relative to others.

Definitions of market efficiency are also linked up with assumptions about what information is available to investors and reflected in the price. For instance, a strict definition of market efficiency that assumes that all information, public as well as private, is reflected in market prices would imply that even investors with precise inside information will be unable to beat the market. One of the earliest classifications of levels of market efficiency was provided by Fama (1971), who argued that markets could be efficient at three levels, based upon what information was reflected in prices. Under weak form efficiency, the current price reflects the information contained in all past prices, suggesting that charts and technical analyses that use past prices alone would not be useful in finding under valued stocks. Under semi-strong form efficiency, the current price reflects the information contained not only in past prices but all public information (including financial statements and news reports) and no approach that was predicated on using and massaging this information would be useful in finding under valued stocks. Under strong form efficiency, the current price reflects all information, public as well as private, and no investors will be able to consistently find under valued stocks.

Implications of market efficiency

An immediate and direct implication of an efficient market is that no group of investors should be able to consistently beat the market using a common investment

strategy. An efficient market would also carry very negative implications for many investment strategies and actions that are taken for granted -

(a) In an efficient market, equity research and valuation would be a costly task that provided no benefits. The odds of finding an undervalued stock would always be 50:50, reflecting the randomness of pricing errors. At best, the benefits from information collection and equity research would cover the costs of doing the research.

(b) In an efficient market, a strategy of randomly diversifying across stocks or indexing to the market, carrying little or no information cost and minimal execution costs, would be superior to any other strategy, that created larger information and execution costs. There would be no value added by portfolio managers and investment strategists.

(c) In an efficient market, a strategy of minimizing trading, i.e., creating a portfolio and not trading unless cash was needed, would be superior to a strategy that required frequent trading.

It is therefore no wonder that the concept of market efficiency evokes such strong reactions on the part of portfolio managers and analysts, who view it, quite rightly, as a challenge to their existence.

It is also important that there be clarity about what market efficiency does not imply. An efficient market does not imply that -

(a) stock prices cannot deviate from true value; in fact, there can be large deviations from true value. The only requirement is that the deviations be random.

(b) no investor will 'beat' the market in any time period. To the contrary, approximately half² of all investors, prior to transactions costs, should beat the market in any period.

(c) no group of investors will beat the market in the long term. Given the number of investors in financial markets, the laws of probability would suggest that a fairly large

² Since returns are positively skewed, i.e., large positive returns are more likely than large negative returns (since this is bounded at -100%), less than half of all investors will probably beat the market.

number are going to beat the market consistently over long periods, not because of their investment strategies but because they are lucky. It would not, however, be consistent if a disproportionately large number³ of these investors used the same investment strategy.

In an efficient market, the expected returns from any investment will be consistent with the risk of that investment over the long term, though there may be deviations from these expected returns in the short term.

Necessary conditions for market efficiency

Markets do not become efficient automatically. It is the actions of investors, sensing bargains and putting into effect schemes to beat the market, that make markets efficient. The necessary conditions for a market inefficiency to be eliminated are as follows -

(1) The market inefficiency should provide the basis for a scheme to beat the market and earn excess returns. For this to hold true -

- (a) The asset (or assets) which is the source of the inefficiency has to be traded.
- (b) The transactions costs of executing the scheme have to be smaller than the expected profits from the scheme.

(2) There should be profit maximizing investors who

- (a) recognize the 'potential for excess return'
- (b) can replicate the beat the market scheme that earns the excess return
- (c) have the resources to trade on the stock until the inefficiency disappears

The internal contradiction of claiming that there is no possibility of beating the market in an efficient market and requiring profit-maximizing investors to constantly seek out ways of beating the market and thus making it efficient has been explored by many. If markets were, in fact, efficient, investors would stop looking for inefficiencies, which would lead to

³ One of the enduring pieces of evidence against market efficiency lies in the performance records posted by many of the investors who learnt their lessons from Ben Graham in the fifties. No probability statistics could ever explain the consistency and superiority of their records.

markets becoming inefficient again. It makes sense to think about an efficient market as a self-correcting mechanism, where inefficiencies appear at regular intervals but disappear almost instantaneously as investors find them and trade on them.

Propositions about market efficiency

A reading of the conditions under which markets become efficient leads to general propositions about where investors are most likely to find inefficiencies in financial markets-

Proposition 1: The probability of finding inefficiencies in an asset market decreases as the ease of trading on the asset increases. To the extent that investors have difficulty trading on a stock, either because open markets do not exist or there are significant barriers to trading, inefficiencies in pricing can continue for long periods.

This proposition can be used to shed light on the differences between different asset markets. For instance, it is far easier to trade on stocks than it is on real estate, since markets are much more open, prices are in smaller units (reducing the barriers to entry for new traders) and the asset itself does not vary from transaction to transaction (one share of IBM is identical to another share, whereas one piece of real estate can be very different from another piece, a stone's throw away). Based upon these differences, there should be a greater likelihood of finding inefficiencies (both under and over valuation) in the real estate market.

Proposition 2: The probability of finding an inefficiency in an asset market increases as the transactions and information cost of exploiting the inefficiency increases. The cost of collecting information and trading varies widely across markets and even across investments in the same markets. As these costs increase, it pays less and less to try to exploit these inefficiencies.

Consider, for instance, the perceived wisdom that investing in 'loser' stocks, i.e., stocks that have done very badly in some prior time period should yield excess returns. This may be true in terms of raw returns, but transactions costs are likely to be much higher for these stocks since-

(a) they then to be low priced stocks, leading to higher brokerage commissions and expenses

(b) the bid-ask spread, a transaction cost paid at the time of purchase, becomes a much higher fraction of the total price paid.

(c) trading is often thin on these stocks, and small trades can cause prices to change resulting in a higher 'buy' price and a lower 'sell' price.

Corollary 1: Investors who can establish a cost advantage (either in information collection or transactions costs) will be more able to exploit small inefficiencies than other investors who do not possess this advantage.

There are a number of studies that look at the effect of block trades on prices, and conclude that while they affect prices, that investors will not be exploit these inefficiencies because of the number of times they will have to trade and their transactions costs. These concerns are unlikely to hold for a specialist on the floor of the exchange, who can trade quickly, often and at no or very low costs. It should be pointed out, however, that if the market for specialists is efficient, the value of a seat on the exchange should reflect the present value of potential benefits from being a specialist.

This corollary also suggests that investors who work at establishing a cost advantage, especially in relation to information, may be able to generate excess returns on the basis of these advantages. Thus a John Templeton, who started investing in Japanese and other Asian markets well before other portfolio managers, might have been able to exploit the informational advantages he had over his peers to make excess returns on his portfolio.

Proposition 3: The speed with which an inefficiency is resolved will be directly related to how easily the scheme to exploit the inefficiency can be replicated by other investors. The ease with which a scheme can be replicated itself is inversely related to the time, resources and information needed to execute it. Since very few investors single-handedly possess the resources to eliminate an inefficiency through trading, it is much more likely that an

inefficiency will disappear quickly if the scheme used to exploit the inefficiency is transparent and can be copied by other investors.

To illustrate this point, assume that stocks are consistently found to earn excess returns in the month following a stock split. Since firms announce stock splits publicly, and any investor can buy stocks right after these splits, it would be surprising if this inefficiency persisted over time. This can be contrasted with the excess returns made by some 'arbitrage funds' in index arbitrage, where index futures are bought (sold), and stocks in the index are sold short (bought). This strategy requires that investors be able to obtain information on index and spot prices instantaneously, have the capacity (in terms of margin requirements and resources) to buy and sell index futures and to sell short on stocks, and to have the resources to take and hold very large positions until the arbitrage unwinds. Consequently, inefficiencies in 'index futures pricing' are likely to persist at least for the most efficient arbitrageurs, with the lowest execution costs and the speediest execution times.

Testing market efficiency

Tests of market efficiency look at whether specific investment strategies earn excess returns. Some tests also account for transactions costs and execution feasibility. Since an excess return on an investment is the difference between the actual and expected return on that investment, there is implicit in every test of market efficiency a model for this expected return. In some cases, this expected return adjusts for risk using the capital asset pricing model or the arbitrage pricing model, and in others the expected return is based upon returns on similar or equivalent investments. In every case, a test of market efficiency is a joint test of market efficiency and the efficacy of the model used for expected returns. When there is evidence of excess returns in a test of market efficiency, it can indicate that markets are inefficient or that the model used to compute expected returns is wrong or both. While this may seem to present an insoluble dilemma, if the conclusions of the study are insensitive to different model specifications, it is much more likely that the results are being driven by true market inefficiencies and not just by model misspecifications.

There are a number of different ways of testing for market efficiency, and the approach used will depend in great part on the investment scheme being tested. A scheme based upon trading on information events (stock splits, earnings announcements or acquisition announcements) is likely to be tested using an 'event study' where returns around the event are scrutinized for evidence of excess returns. A scheme based upon trading on a observable characteristic of a firm (price earnings ratios, price book value ratios or dividend yields) is likely to be tested using a 'portfolio' approach, where portfolios of stocks with these characteristics are created and tracked over time to see if, in fact, they make excess returns. The following pages summarize the key steps involved in each of these approaches, and some potential pitfalls to watch out for when conducting or using these tests.

A. Event Study


An event study is designed to examine market reactions to, and excess returns around specific information events. The information events can be market-wide, such as macro-economic announcements, or firm-specific, such as earnings or dividend announcements. The steps in an event study are as follows -

(1) The event to be studied is clearly identified, and the date on which the event was announced pinpointed. The presumption in event studies is that the timing of the event is known with a fair degree of certainty. Since financial markets react to the information about an event, rather than the event itself, most event studies are centered around the announcement date⁴ for the event.

Announcement Date

⁴ In most financial transactions, the announcement date tends to precede the event date by several days and, sometimes, weeks.

(2) Once the event dates are known, returns are collected around these dates for each of the firms in the sample. In doing so, two decisions have to be made. First, the analyst has to decide whether to collect weekly, daily or shorter-interval returns around the event. This will, in part, be decided by how precisely the event date is known (the more precise, the more likely it is that shorter return intervals can be used) and by how quickly information is reflected in prices (the faster the adjustment, the shorter the return interval to use). Second, the analyst has to determine how many periods of returns before and after the announcement date will be considered as part of the 'event window'. That decision also will be determined by the precision of the event date, since more imprecise dates will require longer windows.




where,

$$R_{jt} = \text{Returns on firm } j \text{ for period } t \text{ (} t = -n, \dots, 0, \dots, +n \text{)}$$

(3) The returns, by period, around the announcement date, are adjusted for market performance and risk to arrive at excess returns for each firm in the sample. For instance, if the capital asset pricing model is used to control for risk -

Excess Return on period $t = \text{Return on day } t - (\text{Riskfree rate} + \text{Beta} * \text{Return on market on day } t)$



where,

$$ER_{jt} = \text{Excess Returns on firm } j \text{ for period } t \text{ (} t = -n, \dots, 0, \dots, +n \text{)}$$

(4) The excess returns, by period, are averaged across all firms in the sample and a standard error is computed.

$$\text{Average excess return on day } t = \frac{\sum_{j=1}^{j=N} ER_{jt}}{N}$$

where,

N = Number of events in the event study

(5) The question of whether the excess returns around the announcement are different from zero is answered by estimating the t statistic for each n , by dividing the average excess return by the standard error -

T statistic for excess return on day t = Average Excess Return / Standard Error

If the t statistics are statistically significant⁵, the event affects returns; the sign of the excess return determines whether the effect is positive or negative.

Illustration 8.1: Example of an event study - Effects of Option Listing on Stock prices

Academics and practitioners have long argued about the consequences of option listing for stock price volatility. On the one hand, there are those who argue that options attract speculators and hence increase stock price volatility. On the other hand, there are others who argue that options increase the available choices for investors and increase the flow of information to financial markets, and thus lead to lower stock price volatility and higher stock prices.

One way to test these alternative hypotheses is to do an event study, examining the effects of listing options on the underlying stocks' prices. Conrad(1989) did such a study, following these steps -

Step 1: The date on which the announcement that options would be listed on the Chicago Board of Options on a particular stock was collected.

⁵ The standard levels of significance for a t statistics are -

Level	One-tailed	Two-tailed
1%	2.33	2.55
5%	1.66	1.96

Step 2: The prices of the underlying stock(j) were collected for each of the ten days prior to the option listing announcement date, the day of the announcement, and each of the ten days after.

Step 3: The returns on the stock (R_{jt}) were computed for each of these trading days.

Step 4: The beta for the stock (β_j) was estimated using the returns from a time period outside the event window (using 100 trading days from before the event and 100 trading days after the event).

Step 5: The returns on the market index (R_{mt}) were computed for each of the 21 trading days.

Step 6: The excess returns were computed for each of the 21 trading days -

$$ER_{jt} = R_{jt} - \beta_j R_{mt} \quad \text{..... } t = -10, -9, -8, \dots, +8, +9, +10$$

The excess returns are cumulated for each trading day.

Step 7: The average and standard error of excess returns across all stocks with option listings were computed for each of the 21 trading days. The t statistics are computed using the averages and standard errors for each trading day. Table 6.1 summarizes the average excess returns and t statistics around option listing announcement dates –

Table 6.1: Excess Returns around Option Listing Announcement Dates

<i>Trading Day</i>	<i>Average Excess Return</i>	<i>Cumulative Excess Return</i>	<i>T Statistic</i>
-10	0.17%	0.17%	1.30
-9	0.48%	0.65%	1.66
-8	-0.24%	0.41%	1.43
-7	0.28%	0.69%	1.62
-6	0.04%	0.73%	1.62
-5	-0.46%	0.27%	1.24
-4	-0.26%	0.01%	1.02
-3	-0.11%	-0.10%	0.93
-2	0.26%	0.16%	1.09
-1	0.29%	0.45%	1.28
0	0.01%	0.46%	1.27

1	0.17%	0.63%	1.37
2	0.14%	0.77%	1.44
3	0.04%	0.81%	1.44
4	0.18%	0.99%	1.54
5	0.56%	1.55%	1.88
6	0.22%	1.77%	1.99
7	0.05%	1.82%	2.00
8	-0.13%	1.69%	1.89
9	0.09%	1.78%	1.92
10	0.02%	1.80%	1.91

Based upon these excess returns, there is no evidence of an announcement effect on the announcement day alone, but there is mild⁶ evidence of a positive effect over the entire announcement period.

B. Portfolio Study

In some investment strategies, firms with specific characteristics are viewed as more likely to be undervalued, and therefore have excess returns, than firms without these characteristics. In these cases, the strategies can be tested by creating portfolios of firms possessing these characteristics at the beginning of a time period, and examining returns over the time period. To ensure that these results are not colored by the idiosyncracies of one time period, this analysis is repeated for a number of periods. The steps in doing a portfolio study are as follows -

- (1) The variable on which firms will be classified is defined, using the investment strategy as a guide. This variable has to be observable, though it does not have to be numerical. Examples would include market value of equity, bond ratings, stock price, price earnings ratios and price book value ratios.
- (2) The data on the variable is collected for every firm in the defined universe⁷ at the start of the testing period, and firms are classified into portfolios based upon the magnitude of the

⁶ The t statistics are marginally significant at the 5% level.

variable. Thus, if the price earnings ratio is the screening variable, firms are classified on the basis of PE ratios into portfolios from lowest PE to highest PE classes. The number of classes will depend upon the size of the universe, since there have to be sufficient firms in each portfolio to get some measure of diversification.

(3) The returns are collected for each firm in each portfolio for the testing period, and the returns for each portfolio are computed, generally assuming that the stocks are equally weighted.

(4) The beta (if using a single factor model) or betas (if using a multifactor model) of each portfolio are estimated, either by taking the average of the betas of the individual stocks in the portfolio or by regressing the portfolio's returns against market returns over a prior time period (for instance, the year before the testing period).

(5) The excess returns earned by each portfolio are computed, in conjunction with the standard error of the excess returns.

(6) There are a number of statistical tests available to check whether the average excess returns are, in fact, different across the portfolios. Some of these tests are parametric⁸ (they make certain distributional assumptions about excess returns) and some are non-parametric⁹.

(7) As a final test, the extreme portfolios can be matched against each other to see whether there are statistically significant differences across these portfolios.

⁷ Though there are practical limits on how big the universe can be, care should be taken to make sure that no biases enter at this stage of the process. An obvious one would be to pick only stocks that have done well over the time period for the universe.

⁸ One parametric test is an F test, which tests for equality of means across groups. This test can be conducted assuming either that the groups have the same variance, or that they have different variances.

⁹ An example of a non-parametric test is a rank sum test, which ranks returns across the entire sample and then sums the ranks within each group to check whether the rankings are random or systematic.

Illustration 8.2: Example of a portfolio study - Price Earnings Ratios

Practitioners have claimed that low price-earnings ratio stocks are generally bargains and do much better than the market or stocks with high price earnings ratios. This hypothesis can be tested using a portfolio approach -

Step 1: Using data on price-earnings ratios from the end of 1987, firms on the New York Stock Exchange were classified into five groups, the first group consisting of stocks with the lowest PE ratios and the fifth group consisting of stocks with the highest PE ratios. Firms with negative price-earnings ratios were ignored.

Step 2: The returns on each portfolio were computed using data from 1988 to 1992. Stocks which went bankrupt or were delisted were assigned a return of -100%.

Step 3: The betas for each stock in each portfolio were computed using monthly returns from 1983 to 1987, and the average beta for each portfolio was estimated. The portfolios were assumed to be equally weighted.

Step 4: The returns on the market index was computed from 1988 to 1992.

Step 5: The excess returns on each portfolio were computed using data from 1988 to 1992.

Table 6.2 summarizes the excess returns each year from 1988 to 1992 for each portfolio.

Table 6.2: Excess Returns from 1988 to 1992 for PE Ratio Portfolios

<i>P/E Class</i>	<i>1988</i>	<i>1989</i>	<i>1990</i>	<i>1991</i>	<i>1992</i>	<i>1988-1992</i>
<i>Lowest</i>	3.84%	-0.83%	2.10%	6.68%	0.64%	2.61%
<i>2</i>	1.75%	2.26%	0.19%	1.09%	1.13%	1.56%
<i>3</i>	0.20%	-3.15%	-0.20%	0.17%	0.12%	-0.59%
<i>4</i>	-1.25%	-0.94%	-0.65%	-1.99%	-0.48%	-1.15%
<i>Highest</i>	-1.74%	-0.63%	-1.44%	-4.06%	-1.25%	-1.95%

Step 6: While the ranking of the returns across the portfolio classes seems to confirm our hypothesis that low PE stocks earn a higher return, we have to consider whether the differences across portfolios is statistically significant. There are several tests available, but these are a few:

- An F test can be used to accept or reject the hypothesis that the average returns are the same across all portfolios. A high F score would lead us to conclude that the differences are too large to be random.
- A chi-squared test is a non-parametric test that can be used to test the hypothesis that the means are the same across the five portfolio classes.
- We could isolate just the lowest PE and highest PE stocks and estimate a t statistic that the averages are different across these two portfolios.

The Cardinal Sins in testing Market Efficiency

In the process of testing investment strategies, there are a number of pitfalls that have to be avoided. Some of them are listed below -

1. Using 'anecdotal evidence' to support/reject an investment strategy: Anecdotal evidence is a double edged sword. It can be used to support or reject the same hypothesis. Since stock prices are noisy and all investment schemes (no matter how absurd) will succeed sometimes and fail at other times, there will always be cases where the scheme works or does not work.

2. Testing an investment strategy on the same data and time period from which it was extracted: This is the tool of choice for the unscrupulous investment advisor. An investment scheme is extracted from hundreds through an examination of the data for a particular time period. This investment scheme is then tested on the same time period, with predictable results. (The scheme does miraculously well and makes immense returns.)

An investment scheme should always be tested out on a time period different from the one it is extracted from or on a universe different from the one used to derive the scheme.

3. Choosing a biased universe: The universe is the sample on which the test is run. Since there are thousands of stocks that could be considered part of this universe, researchers often choose to use a smaller universe. When this choice is random, this does limited damage to the results of the study. If the choice is biased, it can provide results which are not true in the larger universe.

4. Failure to control for market performance: A failure to control for overall market performance can lead one to conclude that your investment scheme works just because it makes good returns (Most schemes will make good returns if the overall market does well; the question is did they make better returns than expected) or does not work just because it makes bad returns (Most schemes will do badly if the overall market performs poorly). It is crucial therefore that investment schemes control for market performance during the period of the test.

5. Failure to control for risk: A failure to control for risk leads to a bias towards accepting high-risk investment schemes and rejecting low-risk investment schemes, since the former should make higher returns than the market and the latter lower, without implying any excess returns.

6. Mistaking correlation for causation: Consider the study on PE stocks cited in the earlier section. We concluded that low PE stocks have higher excess returns than high PE stocks. It would be a mistake to conclude that a low price earnings ratio causes excess returns, since the high returns and the low PE ratio themselves might have been caused by the high risk associated with investing in the stock. In other words, high risk is the causative factor that leads to both the observed phenomena – low PE ratios on the one hand and high returns on the other. This insight would make us more cautious about adopting a strategy of buying low PE stocks in the first place.

Some lesser sins that can be a problem

1. Survival Bias: Most researchers start with a existing universe of publicly traded companies and working back through time to test investment strategies. This can create a subtle bias since it automatically eliminates firms that failed during the period, with obvious negative consequences for returns. If the investment scheme is particularly susceptible to picking firms that have high bankruptcy risk, this may lead to an 'overstatement' of returns on the scheme.

For example, assume that the investment scheme recommends investing in stocks which have very negative earnings, using the argument that these stocks are most likely to benefit from a turnaround. Some of the firms in this portfolio will go bankrupt, and a failure to consider these firms will overstate the returns from this strategy.

2. Not allowing for transactions Costs: Some investment schemes are more expensive than others because of transactions costs - execution fees, bid-ask spreads and price impact. A complete test will take these into account before it passes judgment on the strategy. This is easier said than done, because different investors have different transactions costs, and it is unclear which investor's trading cost schedule should be used in the test. Most researchers who ignore transactions costs argue that individual investors can decide for themselves, given their transactions costs, whether the excess returns justify the investment strategy.

3. Not allowing for difficulties in execution: Some strategies look good on paper but are difficult to execute in practice, either because of impediments to trading or because trading creates a price impact. Thus a strategy of investing in very small companies may seem to create excess returns on paper, but these excess returns may not exist in practice because the price impact is significant.

The Evidence on Market Efficiency

This section of the chapter attempts to summarize the evidence from studies of market efficiency. Without claiming to be comprehensive, the evidence is classified into four sections - the study of price changes and their time series properties, the research on the efficiency of market reaction to information announcements, the existence of return anomalies across firms and over time and the analysis of the performance of insiders, analysts and money managers.

Time Series Properties of Price Changes

Investors have used price charts and price patterns as tools for predicting future price movements for as long as there have been financial markets. It is not surprising,

therefore, that the first studies of market efficiency focused on the relationship between price changes over time, to see if in fact such predictions were feasible. Some of this testing was spurred by the random walk theory of price movements, which contended that price changes over time followed a random walk. As the studies of the time series properties of prices have proliferated, the evidence can be classified into two classes - studies that focus on short-term (intraday, daily and weekly price movements) price behavior and research that examines long-term (annual and five-year returns) price movements.

a. Short term Price Movements

The notion that today's price change conveys information about tomorrow's price change is deep rooted in most investors' psyches. There are several ways in which this hypotheses can be tested in financial markets -

a. Serial correlation

The serial correlation measures the correlation between price changes in consecutive time periods, whether hourly, daily or weekly, and is a measure of how much the price change in any period depends upon the price change over the previous time period. A serial correlation of zero would therefore imply that price changes in consecutive time periods are uncorrelated with each other, and can thus be viewed as a rejection of the hypothesis that investors can learn about future price changes from past ones. A serial correlation which is positive, and statistically significant, could be viewed as evidence of price momentum in markets, and would suggest that returns in a period are more likely to be positive (negative) if the prior period's returns were positive (negative). A serial correlation which is negative, and statistically significant, could be evidence of price reversals, and would be consistent with a market where positive returns are more likely to follow negative returns and vice versa.

From the viewpoint of investment strategy, serial correlations can be exploited to earn excess returns. A positive serial correlation would be exploited by a strategy of buying

after periods with positive returns and selling after periods with negative returns. A negative serial correlation would suggest a strategy of buying after periods with negative returns and selling after periods with positive returns. Since these strategies generate transactions costs, the correlations have to be large enough to allow investors to generate profits to cover these costs. It is therefore entirely possible that there be serial correlation in returns, without any opportunity to earn excess returns for most investors.

The earliest studies of serial correlation (Alexander (1964), Cootner (1962) and Fama (1965) all looked at large U.S. stocks and concluded that the serial correlation in stock prices was small. Fama, for instance, found that 8 of the 30 stocks listed in the Dow had negative serial correlations and that most of the serial correlations were less than 0.05. Other studies confirm these findings not only for smaller stocks in the United States, but also for other markets. For instance, Jennergren and Korsvold (1974) report low serial correlations for the Swedish equity market and Cootner (1961) concludes that serial correlations are low in commodity markets as well. While there may be statistical significance associated with some of these correlations, it is unlikely that there is enough correlation to generate excess returns.

The serial correlation in short period returns is affected by market liquidity and the presence of a bid-ask spread. Not all stocks in an index are liquid, and, in some cases, stocks may not trade during a period. When the stock trades in a subsequent period, the resulting price changes can create positive serial correlation. To see why, assume that the market is up strongly on day 1, but that three stocks in the index do not trade on that day. On day 2, if these stocks are traded, they are likely to go up in price to reflect the increase in the market the previous day. The net result is that you should expect to see positive serial correlation in daily or hourly returns in illiquid market indices.

The bid-ask spread creates a bias in the opposite direction, if transactions prices are used to compute returns, since prices have a equal chance of ending up at the bid or the ask

price. The bounce that this induces in prices will result in negative serial correlations in returns. Roll (1984) provides a simple measure of this relationship,


$$\text{Bid-Ask Spread} = -2 (\text{Serial Covariance in returns})$$

where the serial covariance in returns measures the covariance between return changes in consecutive time periods. For very short return intervals, this bias induced in serial correlations might dominate and create the mistaken view that price changes in consecutive time periods are negatively correlated.

b. Filter Rules

In a filter rule, an investor buys an investment if the price rises X% from a previous low and holds the investment until the price drops X% from a previous high. The magnitude of the change (X%) that triggers the trades can vary from filter rule to filter rule, with smaller changes resulting in more transactions per period and higher transactions costs. Figure 6.1 graphs out a typical filter rule.

Figure 6.1: Filter Rule



This strategy is based upon the assumption that price changes are serially correlated and that there is price momentum, i.e., stocks which have gone up strongly in the past are more likely to keep going up than go down. Table 6.4 summarizes results from a study on returns, before and after transactions costs, on a trading strategy based upon filter rules

ranging from 0.5% to 20%. (A 0.5% rule implies that a stock is bought when it rises 0.5% from a previous low and sold when it falls 0.5% from a prior high.)

Table 6.4: Returns on Filter Rule Strategies

Value of X	Return with strategy	Return with Buy and Hold	# Transactions with strategy	Return after transactions costs
0.5%	11.5%	10.4%	12,514	-103.6%
1.0%	5.5%	10.3%	8,660	-74.9%
2.0%	0.2%	10.3%	4,764	-45.2%
3.0%	-1.7%	10.1%	2,994	-30.5%
4.0%	0.1%	10.1%	2,013	-19.5%
5.0%	-1.9%	10.0%	1,484	-16.6%
6.0%	1.3%	9.7%	1,071	-9.4%
7.0%	0.8%	9.6%	828	-7.4%
8.0%	1.7%	9.6%	653	-5.0%
9.0%	1.9%	9.6%	539	-3.6%
10.0%	3.0%	9.6%	435	-1.4%
12.0%	5.3%	9.4%	289	2.3%
14.0%	3.9%	10.3%	224	1.4%
16.0%	4.2%	10.3%	172	2.3%
18.0%	3.6%	10.0%	139	2.0%
20.0%	4.3%	9.8%	110	3.0%

The only filter rule that beats the returns from the buy and hold strategy is the 0.5% rule, but it does so before transactions costs. This strategy creates 12,514 trades during the period which generate enough transactions costs to wipe out the principal invested by the investor. While this test is an dated, it also illustrates a basic strategies that require frequent short term trading. Even though these strategies may earn excess returns prior to transactions costs, adjusting for these costs can wipe out the excess returns.

One popular indicator among investors that is a variant on the filter rule is the relative strength measure, which relates recent prices on stocks or other investments to either average prices over a specified period, say over six months, or to the price at the beginning of the period. Stocks that score high on the relative strength measure are considered good

investments. This investment strategy is also based upon the assumption of price momentum.

c. Runs Tests

A runs test is a non-parametric variation on the serial correlation, and it is based upon a count of the number of runs, i.e., sequences of price increases or decreases, in the price changes. Thus, the following time series of price changes, where U is an increase and D is a decrease would result in the following runs -

UUU DD U DDD UU DD U D UU DD U DD UUU DD UU D UU D

There were 18 runs in this price series of 33 periods. The actual number of runs in the price series is compared against the number that can be expected¹⁰ in a series of this length, assuming that price changes are random. If the actual number of runs is greater than the expected number, there is evidence of negative correlation in price changes. If it is lower, there is evidence of positive correlation. A study of price changes in the Dow 30 stocks, assuming daily, four-day, nine-day and sixteen day return intervals provided the following results -

	DIFFERENCING INTERVAL			
	Daily	Four-day	Nine-day	Sixteen-day
Actual runs	735.1	175.7	74.6	41.6
Expected runs	759.8	175.8	75.3	41.7

Based upon these results, there is evidence of positive correlation in daily returns but no evidence of deviations from normality for longer return intervals.

Again, while the evidence is dated, it serves to illustrate the point that long strings of positive and negative changes are, by themselves, insufficient evidence that markets are not random, since such behavior is consistent with price changes following a random walk. It is

¹⁰ There are statistical tables that summarize the expected number of runs, assuming randomness, in a series of any length.


the recurrence of these strings that can be viewed as evidence against randomness in price behavior.

Long-term Price Movements

While most of the earlier studies of price behavior focused on shorter return intervals, more attention has been paid to price movements over longer periods (one-year to five-year) in recent years. Here, there is an interesting dichotomy in the results. When long term is defined as months rather than years, there seems to be a tendency towards positive serial correlation. Jegadeesh and Titman present evidence of what they call “price momentum” in stock prices over time periods of up to eight months when investors winner and loser stocks. However, when long term is defined in terms of years, there is substantial negative correlation returns, suggesting that markets reverse themselves over very long periods.

Fama and French examined five-year returns on stocks from 1931 to 1986 and present further evidence of this phenomenon. Studies that break down stocks on the basis of market value have found that the serial correlation is more negative in five-year returns than in one-year returns, and is much more negative for smaller stocks rather than larger stocks. Figure 6.2 summarizes one-year and five-years serial correlation by size class for stocks on the New York Stock Exchange.

Figure 6.2: Serial Correlation in Stock Returns



This phenomenon has also been examined in other markets, and the findings have been similar. There is evidence that returns reverse themselves over long time period.

Winner and Loser portfolios

Since there is evidence that prices reverse themselves in the long term for entire markets, it might be worth examining whether such price reversals occur on classes of stock within a market. For instance, are stocks that have gone up the most over the last period more likely to go down over the next period and vice versa? To isolate the effect of such price reversals on the extreme portfolios, DeBontd and Thaler constructed a winner portfolio of 35 stocks, which had gone up the most over the prior year, and a loser portfolio of 35 stocks, which had gone down the most over the prior year, each year from 1933 to 1978, and examined returns on these portfolios for the sixty months following the creation of the portfolio. Figure 6.3 summarizes the excess returns for winner and loser portfolios .

Figure 6.3: Excess Returns for Winner and Loser Portfolios




FIGURE 23.23 Cumulative average residuals for winner and loser portfolios of 35 stocks, one to 60 months after portfolio formation; length of formation period: 5 years.

This analysis suggests that loser portfolio clearly outperform winner portfolios in the sixty months following creation. This evidence is consistent with market overreaction and correction in long return intervals. Jegadeesh and Titman find the same phenomenon occurring, but present interesting evidence that the winner (loser) portfolios continue to go up (down) for up to eight months after they are created and it is in the subsequent periods that the reversals occur.

There are many, academics as well as practitioners, who suggest that these findings may be interesting but that they overstate potential returns on 'loser' portfolios. For instance, there is evidence that loser portfolios are more likely to contain low priced stocks (selling for less than \$5), which generate higher transactions costs and are also more likely to offer heavily skewed returns, i.e., the excess returns come from a few stocks making phenomenal returns rather than from consistent performance. One study of the winner and loser portfolios attributes the bulk of the excess returns of loser portfolios to low-priced stocks and also finds that the results are sensitive to when the portfolios are created. Loser

portfolios created every December earn significantly higher returns than portfolios created every June.

Speculative Bubbles, Crashes and Panics


Historians who have examined the behavior of financial markets over time have challenged the assumption of rationality that underlies much of efficient market theory. They point out to the frequency with which speculative bubbles have formed in financial markets, as investors buy into fads or get-rich-quick schemes, and the crashes with which these bubbles have ended, and suggest that there is nothing to prevent the recurrence of this phenomenon in today's financial markets. There is some evidence in the literature of irrationality on the part of market players.

a. Experimental Studies of Rationality

Some of the most interesting evidence on market efficiency and rationality in recent years has come from experimental studies. While most experimental studies suggest that traders are rational, there are some examples of irrational behavior in some of these studies.

One such study was done at the University of Arizona. In an experimental study, traders were told that a payout would be declared after each trading day, determined randomly from four possibilities - zero, eight, 28 or 60 cents. The average payout was 24 cents. Thus the share's expected value on the first trading day of a fifteen day experiment was \$3.60 (24×15), the second day was \$3.36 The traders were allowed to trade each day. The results of 60 such experiments is summarized in figure 6.4.

Figure 6.4: Trading Price by Trading Day



There is clear evidence here of a 'speculative bubble' forming during periods 3 to 5, where prices exceed expected values by a significant amount. The bubble ultimately bursts, and prices approach expected value by the end of the period. If this is feasible in a simple market, where every investor obtains the same information, it is clearly feasible in real financial markets, where there is much more differential information and much greater uncertainty about expected value.

It should be pointed out that some of the experiments were run with students, and some with Tucson businessmen, with 'real world' experience. The results were similar for both groups. Furthermore, when price curbs of 15 cents were introduced, the booms lasted even longer because traders knew that prices would not fall by more than 15 cents in a period. Thus, the notion that price limits can control speculative bubbles seems misguided.

b. Behavioral Finance

The irrationality sometimes exhibited by investors has given rise to a whole new area of finance called behavioral finance. Using evidence gathered from experimental psychology, researchers have tried to both model how investors react to information and

predict how prices will change as a consequence. They have been far more successful at the first endeavor than the second. For instance, the evidence seems to suggest the following:

- a. Investors do not like to admit their mistakes. Consequently, they tend to hold on to losing stocks far too long, or in some cases, double up their bets (investments) as stocks drop in value.
- b. More information does not always lead to better investment decisions. Investors seem to suffer both from information overload and a tendency to react to the latest piece of information. Both result in investment decisions that lower returns in the long term.

If the evidence on how investors behave is so clear cut, you might ask, why are the predictions that emerge from these models so noisy? The answer, perhaps, is that any model that tries to forecast human foibles and irrationalities is, by its very nature, unlikely to be a stable one. Behavioral finance may emerge ultimately as a trump card in explaining why and how stock prices deviate from true value, but their role in devising investment strategy still remains questionable.

Behavioral Finance and Valuation

In 1999, Robert Shiller made waves in both academia and investment houses with his book titled *Irrational Exuberance*. His thesis is that investors are often not just irrational but irrational in predictable ways- overreacting to some information and buying and selling in herds. His work forms part of a growing body of theory and evidence of behavioral finance, which can be viewed as a congruence of psychology, statistics and finance.

While the evidence presented for investor irrationality is strong, the implications for valuation are less so. You can consider discounted cash flow valuation to be the antithesis of behavioral finance, because it takes the point of view that the value of an asset is the present value of the expected cash flows generated by that asset. With this context, there are two ways in which you can look at the findings in behavioral finance:

- Irrational behavior in finance may explain why prices can deviate from value (as estimated in a discounted cash flow model). Consequently, it provides the foundation for the excess returns earned by rational investors who base decisions on estimated value. Implicit here is the assumption that markets ultimately recognize their irrationality and correct themselves.

- It may also explain why discounted cash flow values can deviate from relative values (estimated using multiples). Since the relative value is estimated by looking at how the market prices similar assets, market irrationalities that exist will be priced into the asset.

Market Reaction to Information Events

Some of the most powerful tests of market efficiency are event studies where market reaction to informational events (such as earnings and takeover announcements) has been scrutinized for evidence of inefficiency. While it is consistent with market efficiency for markets to react to new information, the reaction has to be instantaneous and unbiased. This point is made in Figure 6.5 by contrasting three different market reactions to information announcements -

Figure 6.5: Information and Price Adjustment



Of the three market reactions pictured here, only the first one is consistent with an efficient market. In second market, the information announcement is followed by a gradual increase in prices, allowing investors to make excess returns after the announcement. This is a slow learning market where some investors will make excess returns on the price drift. In the third market, the price reacts instantaneously to the announcement, but corrects itself in the days that follow, suggesting that the initial price change was an over reaction to the information. Here again, an enterprising investor could have sold short after the announcement, and expected to make excess returns as a consequence of the price correction.


a. Earnings Announcements

When firms make earnings announcements, they convey information to financial markets about their current and future prospects. The magnitude of the information, and the size of the market reaction, should depend upon how much the earnings report exceeds or falls short of investor expectations. In an efficient market, there should be an instantaneous reaction to the earnings report, if it contains surprising information, and prices should increase following positive surprises and down following negative surprises.

Since actual earnings are compared to investor expectations, one of the key parts of an earnings event study is the measurement of these expectations. Some of the earlier studies used earnings from the same quarter in the prior year as a measure of expected earnings, i.e., firms which report increases in quarter-to-quarter earnings provide positive surprises and those which report decreases in quarter-to-quarter earnings provide negative surprises. In more recent studies, analyst estimates of earnings have been used as a proxy for expected earnings, and compared to the actual earnings.

Figure 6.6 provides a graph of price reactions to earnings surprises, classified on the basis of magnitude into different classes from 'most negative' earnings reports (Group 1) to 'most positive' earnings reports (Group 10).

Figure 6.6: Price Reaction to Quarterly Earnings Report



The evidence contained in this graph is consistent with the evidence in most earnings announcement studies -

(a) The earnings announcement clearly conveys valuable information to financial markets; there are positive excess returns (cumulative abnormal returns) after positive announcements and negative excess returns around negative announcements.


(b) There is some evidence of a market reaction in the day immediately prior to the earnings announcement which is consistent with the nature of the announcement, i.e., prices tend to go up on the day before positive announcements and down in the day before negative announcements. This can be viewed either as evidence of insider trading or as a consequence of getting the announcement date wrong¹¹.

(c) There is some evidence, albeit weak, of a price drift in the days following an earnings announcement. Thus, a positive report evokes a positive market reaction on the announcement date, and there are mildly positive excess returns in the days following the earnings announcement. Similar conclusions emerge for negative earnings reports.

The management of a firm has some discretion on the timing of earnings reports and there is some evidence that the timing affects expected returns. A study of earnings reports, classified by the day of the week that the earnings are reported, reveals that earnings and dividend reports on Fridays are much more likely to contain negative information than announcements on any other day of the week. This is shown in figure 6.7.


¹¹ The Wall Street Journal or COMPUSTAT are often used as information sources to extract announcement dates for earnings. For some firms, news of the announcement may actually cross the news wire the day before the Wall Street Journal announcement, leading to a misidentification of the report date and the drift in returns the day before the announcement.

Figure 6.7: Earnings and Dividend Reports by Day of the Week



There is also some evidence that earnings announcements that are delayed, relative to the expected announcement date, are much more likely to contain bad news than earnings announcements which are early or on time. This is graphed in Figure 6.8.

FIGURE 6.8: CUMULATED ABNORMAL RETURNS AND EARNINGS DELAY



Earnings announcements that are more than six days late, relative to the expected announcement date, are much more likely to contain bad news and evoke negative market reactions than earnings announcements that are on time or early.

b. Investment and Project Announcements

Firms frequently make announcements of their intentions of investing resources in projects and research and development. There is evidence that financial markets react to these announcements. The question of whether market have a long term or short term perspective can be partially answered by looking at these market reactions. If financial markets are as short term as some of their critics claim, they should react negatively to announcements by the firm that it plans to invest in research and development. The evidence suggests the contrary. Table 6.5 summarizes market reactions to various classes of investment announcements made by the firm.

Table 6.5: Market Reactions to Investment Announcements

<i>Type of Announcement</i>	<i>Abnormal Returns on</i>	
	<i>Announcement Day</i>	<i>Announcement Month</i>
Joint Venture Formations	0.399%	1.412%

R&D Expenditures	0.251%	1.456%
Product Strategies	0.440%	-0.35%
Capital Expenditures	0.290%	1.499%
All Announcements	0.355%	0.984%

This table excludes the largest investments that most make which is acquisitions of other firms. Here, the evidence is not so favorable. In about 55% of all acquisitions, the stock price of the acquiring firm drops on the announcement of the acquisition, reflecting the market's beliefs that firms tend to overpay on acquisitions.

Market Anomalies

Webster's Dictionary defines an anomaly as a "deviation from the common rule" . Studies of market efficiency have uncovered numerous examples of market behavior that are inconsistent with existing models of risk and return and often defy rational explanation. The persistence of some of these patterns of behavior suggests that the problem, in at least some of these anomalies, lies in the models being used for risk and return rather than in the behavior of financial markets. The following section summarizes some of the more widely noticed anomalies in financial markets in the United States and elsewhere.


Anomalies based upon firm characteristics

There are a number of anomalies that have been related to observable firm characteristics, including the market value of equity, price earnings ratios and price book value ratios.

a. The Small Firm Effect

Studies have consistently found that smaller firms (in terms of market value of equity) earn higher returns than larger firms of equivalent risk, where risk is defined in terms of the market beta. Figure 6.9 summarizes returns for stocks in ten market value classes, for the period from 1927 to 1983.


Figure 6.9: Annual Returns by Size Class: 1927-83



The size of the small firm premium, while it has varied across time, has been generally positive. It was highest during the 1970s and lowest during the 1980s. The persistence of this premium has led to several possible explanations.

(a) The transactions costs of investing in small stocks is significantly higher than the transactions costs of investing in larger stocks, and the premiums are estimated prior to these costs. While this is generally true, the differential transactions costs are unlikely to explain the magnitude of the premium across time, and are likely to become even less critical for longer investment horizons. The difficulties of replicating the small firm premiums that are observed in the studies in real time are illustrated in Figure 6.10, which compares the returns on a hypothetical small firm portfolio (CRSP Small Stocks) with the actual returns on a small firm mutual fund (DFA Small Stock Fund), which passively invests in small stocks.

Figure 6.10: Returns on CRSP Small Stocks versus DFA Small Stock Fund



(b) The capital asset pricing model may not be the right model for risk, and betas underestimate the true risk of small stocks. Thus, the small firm premium is really a measure of the failure of beta to capture risk. The additional risk associated with small stocks may come from several sources. First, the estimation risk associated with estimates of beta for small firms is much greater than the estimation risk associated with beta estimates for larger firms. The small firm premium may be a reward for this additional estimation risk. Second, there may be additional risk in investing in small stocks because far less information is available on these stocks. In fact, studies indicate that stocks that are neglected by analysts and institutional investors earn an excess return that parallels the small firm premium.


There is evidence of a small firm premium in markets outside the United States as well. Dimson and Marsh examined stocks in the United Kingdom from 1955 to 1984 and found that the annual returns on small stocks exceeded that on large stocks by 7% annually over the period. Bergstrom, Frashure and Chisholm report a large size effect for French stocks (Small stocks made 32.3% per year between 1975 to 1989, while large stocks made

23.5% a year), and a much smaller size effect in Germany. Chan, Hamao and Lakonishok reports a small firm premium of 5.1% for Japanese stocks between 1971 and 1988.

b. Price Earnings Ratios

Investors have long argued that stocks with low price earnings ratios are more likely to be undervalued and earn excess returns. For instance, Ben Graham, in his investment classic "The Intelligent Investor", uses low price earnings ratios as a screen for finding under valued stocks. Studies that have looked at the relationship between PE ratios and excess returns confirm these priors. Figure 6.11 summarizes annual returns by PE ratio classes for stocks from 1967 to 1988.

Figure 6.11: Annual Returns by PE Ratio Class



F

firms in the lowest PE ratio class earned an average return of 16.26% during the period, while firms in the highest PE ratio class earned an average return of only 6.64%.

The excess returns earned by low PE ratio stocks also persist in other international markets. Table 6.6 summarizes the results of studies looking at this phenomenon in markets outside the United States.

Table 6.6: Excess Returns on Low P/E Ratio Stocks by Country: 1989-1994

<i>Country</i>	<i>Annual Premium earned by lowest P/E Stocks (bottom quintile)</i>
Australia	3.03%
France	6.40%
Germany	1.06%
Hong Kong	6.60%
Italy	14.16%
Japan	7.30%
Switzerland	9.02%
U.K.	2.40%

Annual premium: Premium earned over an index of equally weighted stocks in that market between January 1, 1989 and December 31, 1994. These numbers were obtained from a Merrill Lynch Survey of Proprietary Indices.

The excess returns earned by low price earnings ratio stocks are difficult to justify using a variation of the argument used for small stocks, i.e., that the risk of low PE ratio stocks is understated in the CAPM. Low PE ratio stocks generally are characterized by low growth, large size and stable businesses, all of which should work towards reducing their risk rather than increasing it. The only explanation that can be given for this phenomenon, which is consistent with an efficient market, is that low PE ratio stocks generate large dividend yields, which would have created a larger tax burden in those years where dividends were taxed at higher rates.

c. Price Book Value Ratios

Another statistic that is widely used by investors in investment strategy is price book value ratios. A low price book value ratio has been considered a reliable indicator of undervaluation in firms. In studies that parallel those done on price earnings ratios, the relationship between returns and price book value ratios has been studied. The consistent finding from these studies is that there is a negative relationship between returns and price book value ratios, i.e., low price book value ratio stocks earn higher returns than high price book value ratio stocks.

Rosenberg, Reid and Lanstein (1985) find that the average returns on U.S. stocks are positively related to the ratio of a firm's book value to market value. Between 1973 and 1984, the strategy of picking stocks with high book/price ratios (low price-book values) yielded an excess return of 36 basis points a month. Fama and French (1992), in examining

the cross-section of expected stock returns between 1963 and 1990, establish that the positive relationship between book-to-price ratios and average returns persists in both the univariate and multivariate tests, and is even stronger than the size effect in explaining returns. When they classified firms on the basis of book-to-price ratios into twelve portfolios, firms in the lowest book-to-price (higher P/BV) class earned an average monthly return of 0.30%, while firms in the highest book-to-price (lowest P/BV) class earned an average monthly return of 1.83% for the 1963-90 period.

Chan, Hamao and Lakonishok (1991) find that the book-to-market ratio has a strong role in explaining the cross-section of average returns on Japanese stocks. Capaul, Rowley and Sharpe (1993) extend the analysis of price-book value ratios across other international markets, and conclude that value stocks, i.e., stocks with low price-book value ratios, earned excess returns in every market that they analyzed, between 1981 and 1992. Their annualized estimates of the return differential earned by stocks with low price-book value ratios, over the market index, were as follows:

<i>Country</i>	<i>Added Return to low P/BV portfolio</i>
France	3.26%
Germany	1.39%
Switzerland	1.17%
U.K	1.09%
Japan	3.43%
U.S.	1.06%
Europe	1.30%
Global	1.88%

A caveat is in order. Fama and French point out that low price-book value ratios may operate as a measure of risk, since firms with prices well below book value are more likely to be in trouble and go out of business. Investors therefore have to evaluate for themselves

whether the additional returns made by such firms justifies the additional risk taken on by investing in them.


Temporal Anomalies

There are a number of peculiarities in return differences across calendar time that are not only difficult to rationalize but are also suggestive of inefficiencies. Furthermore, some of these temporal anomalies are related to the small firm effect described in the previous section.

a. The January Effect

Studies of returns in the United States and other major financial markets consistently reveal strong differences in return behavior across the months of the year. Figure 6.12 reports average returns by month of the year from 1926 to 1983.


Figure 6.12: Average Return by Month of the Year



Returns in January are significantly higher than returns in any other month of the year. This phenomenon is called the year-end or January effect, and it can be traced to the first two weeks in January.

The relationship between the January effect and the small firm effect adds to the complexity of this phenomenon. The January effect is much more accentuated for small firms than for larger firms, and roughly half of the small firm premium, described in the prior section, is earned in the first two days of January. Figure 6.13 graphs returns in January by size and risk class for data from 1935 to 1986.

Figure 6.13: Returns in January by Size and Risk Class - 1935-86




A number of explanations have been advanced for the January effect, but few hold up to serious scrutiny. One is that there is tax loss selling by investors at the end of the year on stocks which have 'lost money' to capture the capital gain, driving prices down, presumably below true value, in December, and a buying back of the same stocks¹² in January, resulting

¹² Since wash sales rules would prevent an investor from selling and buying back the same stock within 45 days, there has to be some substitution among the stocks. Thus investor 1 sells stock A and investor 2

in the high returns. The fact that the January effect is accentuated for stocks which have done worse over the prior year is offered as evidence for this explanation. There are several pieces of evidence that contradict it, though. First, there are countries, like Australia, which have a different tax year, but continue to have a January effect. Second, the January effect is no greater, on average, in years following bad years for the stock market, than in other years.

A second rationale is that the January effect is related to institutional trading behavior around the turn of the years. It has been noted, for instance, that ratio of buys to sells for institutions drops significantly below average in the days before the turn of the year and picks to above average in the months that follow. This is illustrated in Figure 6.14.

Figure 6.14: Institutional Buying/Selling around Year-end




sells stock B, but when it comes time to buy back the stock, investor 1 buys stock B and investor 2 buys stock A.

It is argued that the absence of institutional buying pushes down prices in the days before the turn of the year and pushes up prices in the days after.

The universality of the January effect is illustrated in Figure 6.15, which examines returns in January versus the other months of the year in several major financial markets, and finds strong evidence of a January effect in every market.


Figure 6.15: Returns in January vs Other Months - Major Financial Markets



b. The Weekend Effect

The weekend effect is another return phenomenon that has persisted over extraordinary long periods and over a number of international markets. It refers to the differences in returns between Mondays and other days of the week. The significance of the return difference is brought out in Figure 6.16, which graphs returns by days of the week from 1962 to 1978.

Figure 6.16: Average Daily Returns by Day of the Week: 1962-78




The returns on Mondays are significantly negative, whereas the returns on every day of the week are not. There are a number of other findings on the Monday effect that have fleshed it out. First, the Monday effect is really a weekend effect since the bulk of the negative returns is manifested in the Friday close to Monday open returns. The returns from intraday returns on Monday are not the culprits in creating the negative returns. Second, the Monday effect is worse for small stocks than for larger stocks. Third, the Monday effect is no worse following three-day weekends than two-day weekends.

There are some who have argued that the weekend effect is the result of bad news being revealed after the close of trading on Friday and during the weekend. They point to figure 6.16, which reveals that more negative earnings reports are revealed after close of trading on Friday. Even if this were a widespread phenomenon, the return behavior would be inconsistent with a rational market, since rational investors would build in the expectation of the bad news over the weekend into the price before the weekend, leading to an elimination of the weekend effect.

The weekend effect is fairly strong in most major international markets, as shown in Figure 6.17.


Figure 6.17: Weekend Effect in International Markets



The presence of a strong weekend effect in Japan, which allowed Saturday trading for a portion of the period studies here indicates that there might be a more direct reason for negative returns on Mondays than bad information over the weekend.

As a final note, the negative returns on Mondays cannot be just attributed to the absence of trading over the weekend. The returns on days following trading holidays, in general, are characterized by positive, not negative, returns. Figure 6.18 summarizes returns on trading days following major holidays and confirms this pattern.

Figure 6.18




Evidence on Insiders and Investment Professionals

There is a sense that insiders, analysts and portfolio managers must possess an advantage over the average investors in the market and be able to convert this advantage into excess returns. The evidence on the performance of these investors is actually surprisingly mixed.

a. Insider Trading

The SEC defines an insider to be an officer or director of the firm or a major stockholder (holding more than 5% of the outstanding stock in the firm). Insiders are barred from trading in advance of specific information on the company and are required to file with the SEC when they buy or sell stock in the company. If it is assumed, as seems reasonable, that insiders have better information about the company, and consequently better estimates of value, than other investors, the decisions by insiders to buy and sell stock should affect stock prices. Figure 6.19, derived from an early study of insider trading by Jaffe, examines excess returns on two groups of stock, classified on the basis of insider trades. The "buy group" includes stocks where buys exceeded sells by the biggest margin, and the "sell group" includes stocks where sells exceed buys by the biggest margin.

Figure 6.19: Cumulative Returns Following Insider Trading: Buy vs Sell Group



While it seems like the buy group does significantly better than the sell group in this study, advances in information technology have made this information on insider trading available to more and more investors. A more recent study of insider trading examined excess returns around both the date the insiders report to the SEC and the date that information becomes available to investors in the official summary. Figure 6.20 presents the contrast between the two event studies.

Figure 6.20: Abnormal Returns around Reporting Day/ Official Summary Availability Day





Fig. 2. Cumulative daily average prediction errors from 199 days before to 300 days after the first buy insiders' reports are received by the SEC and the availability day of the Official Summary to a portfolio of 769 firms traded by insiders during 1975 to 1981.

Given the opportunity to buy on the date the insider reports to the SEC, investors could have marginal excess returns, but these returns diminish and become statistically insignificant, if investors are forced to wait until the official summary date.

None of these studies examine the question of whether insiders themselves make excess returns. The reporting process, as set up now by the SEC, is biased toward legal and less profitable trades, and away from illegal and more profitable trades. Though direct evidence cannot be offered for this proposition, insiders trading illegally on private information must make excess returns.

b. Analyst Recommendations

Analysts clearly hold a privileged position in the market for information, operating at the nexus of private and public information. Using both types of information, analysts issue buy and sell recommendations to their clients, who trade on its basis.



While both buy and sell recommendations affect stock prices, sell recommendations affect prices much more adversely than buy recommendation affect them positively. Interestingly, Womack (1996) documents that the price effect of buy recommendations tends to be immediate and there is no evidence of price drifts after the announcement, whereas prices


continue to trend down after sell recommendations. Figure 6.21 graphs his findings. Stock prices increase by about 3% on buy recommendations whereas they drop by about 4% on sell recommendations at the time of the recommendations (3 days around reports). In the six months following, prices decline an additional 5% for sell recommendations, while leveling off for buy recommendations.

Though analysts provide a valuable service in collecting private information, or maybe because they do, there is a negative relationship in the cross-section between returns earned by stocks and the number of analysts following the stock. The same kind of relationship exists between another proxy for interest, institutional ownership, and returns.

c. Money Managers


Professional money managers operate as the experts in the field of investments. They are supposed to be better informed, smarter, have lower transactions costs and be better investors overall than smaller investors. The earliest study of mutual funds by Jensen suggested that this supposition might not hold in practice. His findings, summarized in Figure 6.22, as excess returns on mutual funds, were that the average portfolio manager actually underperformed the market between 1955 and 1964.

Figure 6.22: Mutual Fund Performance: 1955-64 - The Jensen Study



These results have been replicated with mild variations in the conclusions. In the studies that are most favorable for professional money managers, they break even against the market after adjusting for transactions costs, and in those that are least favorable, they underperform the market even before adjusting for transactions costs.


The results, when categorized on a number of different basis, do not offer much solace. For instance, Figure 6.23 shows excess returns from 1983 to 1990, and the percentage of money managers beating the market, categorized by investment style.



Money managers in every investment style underperform the market index.

Figure 6.24, from the same study, looks at the payoff to active portfolio management by looking at the added value from trading actively during the course of the year and finds that returns drop from 0.5% to 1.5% a year as a consequence.

**Figure 9.22: The Payoff to Active Money Management:
Equity Funds**



This table measures the difference between actual return on equity funds and return on hypothetical portfolio frozen at beginning of period.

Finally, the study, like others before it, found no evidence of continuity in performance. It classified money managers into quartiles and examined the probabilities of movement from one quartile to another each year from 1983 to 1990. The results are summarized in Table 6.7.

Table 6.7: Probabilities of Transition from One Quartile to Another

<i>Ranking this period</i>	<i>Ranking next period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>1</i>	26%	24%	23%	27%
<i>2</i>	20%	26%	29%	25%
<i>3</i>	22%	28%	26%	24%
<i>4</i>	32%	22%	22%	24%

This table indicates that a money manager who was ranked in the first quartile in a period had a 26% chance of being ranked in the first quartile in the next period and a 27% chance of being ranked in the bottom quartile. There is some evidence of reversal in the portfolio

managers in the lowest quartile, though some of that may be a reflection of the higher risk portfolios that they put together.

While the evidence is depressing for active portfolio management as a whole, there are a few bright spots. Carhart (1992) looked at mutual funds and concluded that there was some persistence at the extremes – a small group of exceptional managers who outperform a passive strategy and another group who consistently underperform, largely because they have high expenses.

Market Inefficiencies and Money Manager Performance

The evidence on markets is contradictory. On the one hand, there seem to be numerous patterns in stock prices – stock prices reverse course in the long term and returns are higher in January – and evidence of market anomalies – small market-cap firms with low price to book and price to earnings ratios seem to handily beat the market. On the other, there seems to be little evidence of money managers being able to exploit these findings to beat the market.

There are a number of possible explanations. The most benign one is that the inefficiencies show up mostly in hypothetical studies and that the transactions costs and execution problems associated with converting these inefficiencies into portfolios overwhelms the excess returns. A second possible explanation is that the studies generally look at the long term – many are over 20 to 50 years. Over shorter periods, there is substantially more uncertainty about whether small stocks will outperform large stocks and whether buying losers will generate excess returns. There are no investment strategies that are sure bets for short periods. Pradhuman (2000) illustrates this phenomenon by noting that small cap stocks have underperformed large cap stocks in roughly one out of every four years in the last 50 years. Bernstein (1998) notes that while value investing (buying low PE and low Price to book value stocks) may earn excess returns over long periods, growth investing has outperformed value investing over many five-year periods during the last three decades. A third explanation is that portfolio managers do not consistently follow any one

strategy and that they jump from one strategy to another, increasing both their expenses and reducing the likelihood that the strategy can generate excess returns in the long term.

Conclusion

The question of whether markets are efficient will always be a provocative one, given the implications that efficient markets have for investment management and research. If an efficient market is defined as one where the market price is an unbiased estimate of the true value, it is quite clear that some markets will always be more efficient than others and that markets will always be more efficient to some investors than to others. The capacity of a market to correct inefficiencies quickly will depend, in part, on the ease of trading, the transactions cost and the vigilance of profit-seeking investors in that market.

While market efficiency can be tested in a number of different ways, the two most widely used tests to test efficiency are 'event studies' which examine market reactions to information events and 'portfolio studies' which evaluate the returns of portfolios created on the basis of observable characteristics. It does make sense to be vigilant, because bias can enter these studies, intentionally or otherwise, in a number of different ways and can lead to unwarranted conclusions, and, worse still, wasteful investment strategies.

There is substantial evidence of irregularities in market behavior, related to systematic factors such as size, price-earnings ratios and price book value ratios, as well as to time - the January and the weekend effects. While these irregularities may be inefficiencies, there is also the sobering evidence that professional money managers, who are in a position to exploit these inefficiencies, have a very difficult time consistently beating financial markets. Read together, the persistence of the irregularities and the inability of money managers to beat the market is testimony to the gap between empirical tests on paper and real world money management in some cases, and the failure of the models of risk and return in others.

1. Which of the following is an implication of market efficiency? (There may be more than one right answer)

- (a) Resources are allocated among firms efficiently (i.e. put to best use)
- (b) No investor will do better than the market in any time period
- (c) No investor will do better than the market consistently
- (d) No investor will do better than the market consistently after adjusting for risk
- (e) No investor will do better than the market consistently after adjusting for risk and transactions costs
- (f) No group of investors will do better than the market consistently after adjusting for risk and transactions costs.

2. Suppose you are following a retailing stock which has a strong seasonal pattern to sales. Would you expect to see a seasonal pattern in the stock price as well?

3. Tests of market efficiency are often referred to as joint tests of two hypotheses - the hypothesis that the market is efficient and an expected returns model. Explain. Is it ever possible to test market efficiency alone? (i.e. without jointly testing an asset pricing model)

4. You are in a violent argument with a chartist. He claims that you are violating the fundamental laws of economics by trying to find intrinsic value. "Price is determined by demand and supply... not by some intrinsic value'. Is finding an intrinsic value inconsistent with demand and supply?

5. You are testing the effect of merger announcements on stock prices. (This is an event study.). Your procedure goes through the following steps.

Step 1: You choose the twenty biggest mergers of the year

Step 2: You isolate the date the merger became effective as the key day around which you will examine the data

Step 3: You look at the returns for the five days after the effective merger date

By looking at these returns (.13%) you conclude that you could not have made money on merger announcements. Are there any flaws that you can detect in this test? How would you correct for them? Can you devise a stronger test?

6. In an efficient market, the market price is defined to be an 'unbiased estimate' of the true value. This implies that

- (a) the market price is always equal to true value.
- (b) the market price has nothing to do with true value
- (c) markets make mistakes about true value, and investors can exploit these mistakes to make money

(d) market prices contain errors, but the errors are random and therefore cannot be exploited by investors.

(e) no one can beat the market.

7. Evaluate whether the following actions are likely to increase stock market efficiency, decrease it or leave it unchanged, and explain why.

a. The government imposes a transaction tax of 1% on all stock transactions.

Increase Efficiency___ Decrease Efficiency___ Leave unchanged___

b. The securities exchange regulators impose a restriction on all short sales to prevent rampant speculation.

Increase Efficiency___ Decrease Efficiency___ Leave unchanged___


c. An options market, trading call and put options, is opened up, with options traded on many of the stocks listed on the exchange.

Increase Efficiency___ Decrease Efficiency___ Leave unchanged___

d. The stock market removes all restrictions on foreign investors acquiring and holding stock in companies.

Increase Efficiency___ Decrease Efficiency___ Leave unchanged___

8. The following is a graph of cumulative abnormal returns around the announcement of asset divestitures by major corporations.



How best would you explain the

(a) market behavior before the announcement?

(b) market reaction to the announcement ?

(c) market reaction after the announcement?

9. What is the phenomenon of the size effect in stock performance? How does it relate to the 'turn-of-the-year' effect? Can you suggest any good reasons why small stocks, after adjusting for beta, still do

better than large stocks? What strategy would you follow to exploit this anomaly? What factors do you have to keep in mind?

10. A study examining market reactions to earnings surprises found that prices tend to drift after earnings surprises. What does this tell you market's capacity to learn from events and new information? What cross-sectional differences would you expect to find in this learning behavior? (i.e. Would you expect to see a greater price drift in some types of firms than in others? Why?) How would you try to exploit this anomaly? What possible costs would you have to keep in mind?

11. One explanation of the turn-of-the-year or January effect has to do with sales and purchases related to the tax year.

(a) Present the tax effect hypothesis

(b) Studies have shown that the January effect occurs internationally, even in countries where the tax year does not start in January. Speculate on a good reason for this.

12. The following are the expected price appreciation and dividend yield components of returns on two portfolios - a 'high dividend yield' portfolio and a 'low dividend yield' portfolio.

<i>Portfolio</i>	<i>Expected Price Appreciation</i>	<i>Expected Dividend Yield</i>
High Yield	9%	5%
Low Yield	12%	1%

You are a taxable investor who faces a tax rate of 40% on dividends. What would your tax rate on capital gains need to be for you to be indifferent between these two portfolios?

13. Answer true or false to the following questions –

a. Low price-earnings stocks, on average, earn returns in excess of expectations, while high price-earnings stocks earn less than expected. This is primarily because lower P/E ratio stocks have lower risk.

TRUE

FALSE

b. The small firm effect, which refers the positive excess returns earned, on average, by small firms, is primarily caused by a few small firms that make very high positive returns.

TRUE

FALSE

c. Investors generally cannot make money on analyst recommendations, because stock prices are not affected by these recommendations.

TRUE

FALSE

14. You are examining the performance of two mutual funds. AD VALUE Fund has been in existence since January 1, 1988 and invests primarily in low Price Earnings Ratio stocks, with high dividend yields. AD GROWTH Fund has also been in existence since January 1, 1988 but it invests primarily in high growth stocks, with high PE ratios and low or no dividends. The performance of these funds over the last five years is summarized below:

Average from 1988-1992		
Price Appreciation	Dividend Yield	Beta

NYSE Composite	13%	3%	1.0
AD VALUE	11%	5%	0.8
AD GROWTH	15%	1%	1.2

The average riskfree rate during the period was 6%. The current riskfree rate is 3%.

- a. How well or badly did these funds perform after adjusting for risk?
- b. Assume that the front-end load on each of these funds is 5% (i.e. if you put \$1000 in each of these funds today, you would only be investing \$950 after the initial commission). Assume also that the excess returns you have calculated in part (a) will continue into the future and that you choose to invest in the fund that outperformed the market. How many years would you have to hold this fund to break even?