Market Efficiency: Definitions and Tests

Why market efficiency matters ..

- 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 is 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.
- Market '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
  - **Saves time** for the analyst
  - **Increases the odds significantly** of finding under and over valued stocks.

What is an efficient market?

- 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 -
- 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 random.

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

Implications of Randomness

- 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 under valued than stocks with high PE ratios.
- 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.

Definitions of Market Efficiency

- 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 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.

Information and Market Efficiency

- *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

- 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 should be random** (50/50). 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.

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 deviations do have to be random.
  - (b) **no investor will 'beat' the market in any time period**. To the contrary, approximately half 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 number are going to beat the market consistently over long periods, not because of their investment strategies but because they are lucky.
- 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

Efficient Markets and Profit-seeking Investors: The Internal Contradiction

- There is an **internal contradiction** in claiming that there is no possibility of beating the market in an efficient market and then requiring profit-maximizing investors to constantly seek out ways of beating the market and thus making it efficient.
- If markets were, in fact, efficient, **investors would stop looking for inefficiencies**, which would lead to markets becoming inefficient again.
- It makes sense to think about an efficient market as a **self-correcting mechanism**, where inefficiencies appear at regular intervals

Market Efficiency and Trading Ease

- **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 an asset, either because open markets do not exist or there are significant barriers to trading, inefficiencies in pricing can continue for long periods.

Market Efficiency and the Cost of Information/Transacting

- **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.
- An Example:
  - **Investing in 'loser' stocks**, i.e., stocks that have done very badly in some prior time period should yields excess returns. 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 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 move.

The Payoff to Establishing Information Advantages

- **Corollary 1: Investors who can estabish 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.
- **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 othe 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.

Market Efficiency and Imitators

- **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.

Testing Market Efficiency

- Tests of market efficiency look at the whether specific investment strategies earn excess returns. Some tests also account for transactions costs and execution feasibility. 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**.
- 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.

1. 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.

Event Study: The First Step

- **Step 1: Identify the event**

  (1) The event to be studied is clearly identified, and the date on which the event was announced pinpointed.

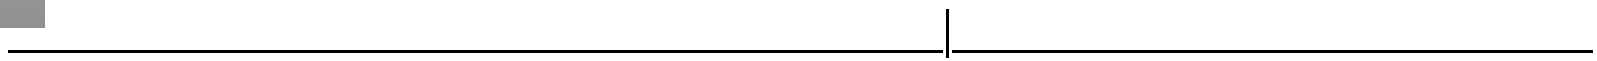

Event Study: Collecting the Returns

- 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

  - by how quickly information is reflected in

  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'.

Return window: -n to +n

Event Study: Controlling for the Market

- 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 day t = Return on day t - Beta * Return on market on day t

ER_{-jn} ..... ER_{j0} .....ER_{+jn}

Return window: -n to +n

Where ER = Excess return on stock j in period t

Event Study: Looking at the Average

- The excess returns, by day, are averaged across all firms in the sample and a standard error is computed.

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

where,

N = Number of events in the event study

Event Study: Estimating Statistical Significance

- 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 significant, the event affects returns; the sign of the excess return determines whether the effect is positive or negative.

An Example: The 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**.

The Steps in Testing the Effect of Listing

- *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.
- *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)

The Steps (Continued)

- *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 -

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

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.

The Results of the Study

Day	*Average Return*	*Cumulative Excess*	*T Statistic*	*Day*	*Average Return*	*Cumulative Excess*	*T Statistic*
-10	0.17%	0.17%	1.30	1	0.17%	0.63%	1.37
-9	0.48%	0.65%	1.66	2	0.14%	0.77%	1.44
-8	-0.24%	0.41%	1.43	3	0.04%	0.81%	1.44
-7	0.28%	0.69%	1.62	4	0.18%	0.99%	1.54
-6	0.04%	0.73%	1.62	5	0.56%	1.55%	1.88
-5	-0.46%	0.27%	1.24	6	0.22%	1.77%	1.99
-4	-0.26%	0.01%	1.02	7	0.05%	1.82%	2.00
-3	-0.11%	-0.10%	0.93	8	-0.13%	1.69%	1.89
-2	0.26%	0.16%	1.09	9	0.09%	1.78%	1.92
-1	0.29%	0.45%	1.28	10	0.02%	1.80%	1.91

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 any one time period, this is repeated for a number of periods.

Steps in Doing a Portfolio Study

(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.

(2) The **data on the variable is collected** for every firm in the defined universe at the start of the testing period, and firms are classified into portfolios based upon the magnitude of the variable.

(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 of each portfolio is 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 period.

Steps in Doing a Portfolio Study (Contd)

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

(6) There are a **number of statistical tests** available to check whether the average excess returns are, in fact, different across the portfolios.

(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.

An Example: Low PE Ratio Stocks as Investments

- 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.

Steps in Testing Low PE ratio Hypotheses

- *Step 1:* Using data on **PE 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.

Testing Low PE Strategy (Continued)

- Step 5: The **raw returns on each portfolio** were computed using data from 1988 to 1992. The following table summarizes the returns each year from 1988 to 1992 for each portfolio.

P/E Class	*1988*	*1989*	*1990*	*1991*	*1992*	*1988-1992*
Lowest	20.65%	30.66%	-1.35%	37.25%	8.22%	19.16%
2	18.56%	33.75%	-3.26%	31.66%	8.51%	18.01%
3	17.01%	28.33%	-3.65%	30.74%	7.70%	15.96%
4	15.56%	30.55%	-4.10%	28.56%	7.10%	15.40%
Highest	15.07%	30.86%	-4.89%	26.51%	6.33%	14.60%
Average	16.81%	31.49%	-3.45%	30.57%	7.58%	16.55%

Low PE Strategy: Excess Returns

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

3. Regressions

- One of the limitations of portfolio studies is that they become increasing unwieldy, as the number of variables that you use in your strategy increases.
- The other problem with portfolio studies is that you group firms into classes and ignore differences across firms within each class. Thus, the stocks in the lowest PE ratio class may have PE ratios that range from the 4 to 12.
- If you believe that these differences may affect the expected returns on your strategy, you could get a better measure of the relationship by running a multiple regression. Your dependent variable would be the returns on stocks and the independent variables would include the variables that form your strategy.

Steps in Regression
Step 1: Identify your dependent variable

- This is the variable that you are trying to explain. In most investment schemes, it will be a measure of the return you would make on the investment but you have to make at least two judgments.
  - The first is whether you plan to use total returns or excess returns; with the latter, you would adjust the returns for risk and market performance, using one the measures discussed earlier in the chapter.
  - The second decision you have to make is on the return interval you will be using – monthly, quarterly, annual or five-year, for instance. This choice will be determined both by your investment strategy – long-term strategies require long-term returns – and the ease with which you can get data on your independent variables for the intervals. For instance, if you use accounting variables such as earnings or book value as independent variables, you will be able to get updates only once every quarter for these variables.

Step 2: Decide on how you will measure the variables that will underlie you strategy

- In this step, you will have to go beyond qualitative variables and come up with quantitative measures. For instance, if you one of the characteristics that you are looking for high growth, you have to choose your measure of growth - growth in earnings or revenues, growth in the past or expected future growth.
- The information on these variables has to be collected at the beginning of your testing period rather than the end. In other words, if you are looking at returns in 2000, you have to measure growth at the beginning of 2000.

Step 3: Check for the nature of the relationship

Panel A: Linear Relationship

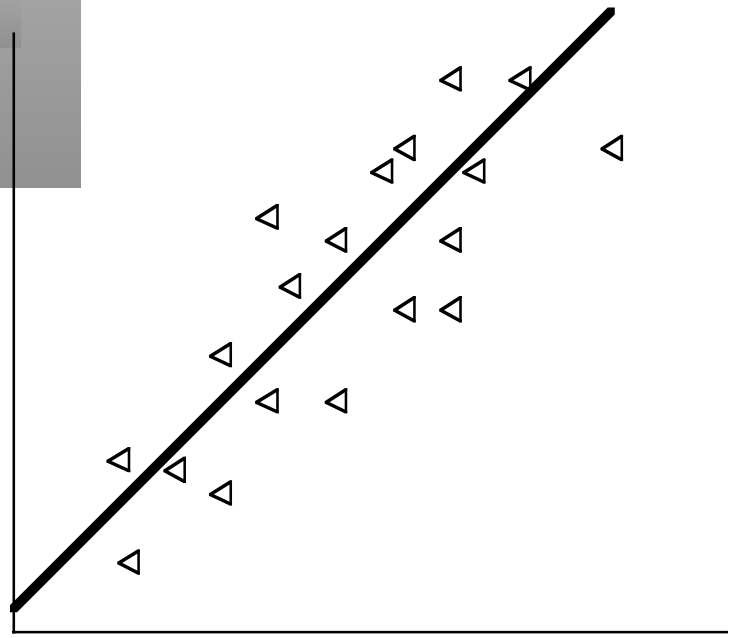

Panel B: Non-linear Relationship

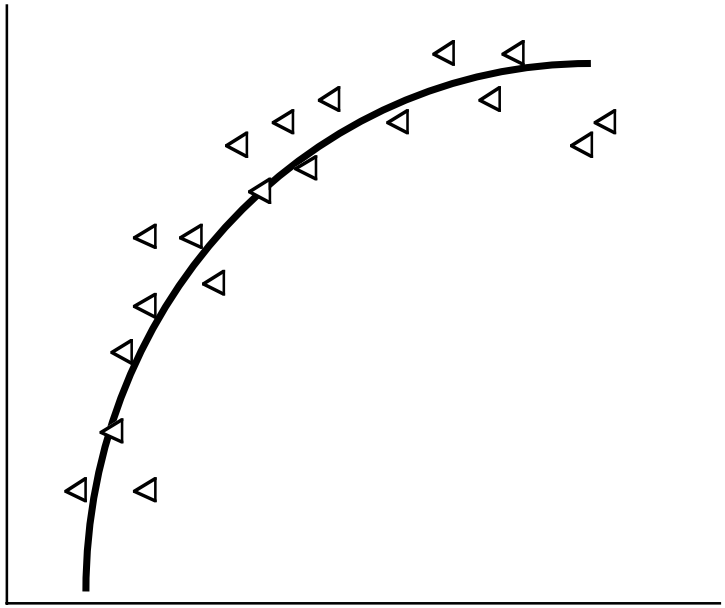

Step 4: Run the regression

- The regression can either be run
  - Across firms or markets at a point in time: this is called a cross sectional regression.
  - For a market across a number of years: this is called a time series regression.
- Once you run the regression, you have to pass it through the tests for statistical significance. In other words, even if all of the coefficients have the right signs, you have to check to ensure that they are significantly different from zero.
  - In most regressions, statistical significance is estimated with a t statistic for each coefficient. This t statistic is computed by dividing the coefficient by the standard error of the coefficient.
  - You can also compute an F statistic to measure whether the regression collectively yield statistically significant results.

The Cardinal Sins in Testing Strategies

- *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.

Data Scheming

- *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.*

Bias in Sample Selection

- *Choosing a biased sample*, The sample on which the test is run may be selected in a "biased" way. Since there are thousands of stocks that could be considered part of this universe, researchers often choose to use a smaller sample. 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.

Failure to Control for Market Performance

- *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.

Missing the Effects of Risk

- *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.

Some Lesser Sins

- *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.

Cost of Trading

- *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.

All in the Execution

- *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.