CUBS MSc Finance 1999-2000

ASSET MANAGEMENT: LECTURE 5

Dr. D.N. Tambakis

MARKET EFFICIENCY

Textbook: EG17, BKM12, BM 13

Readings: Thaler, R. "The January Effect", Journal of Economic Perspectives 1(1), 1987a, 197-201.

Thaler, R. "Seasonal Movements in Security Prices II: Weekend, Holiday, Turn of the Month, and Intraday Effects", Journal of Economic Perspectives 1(2), 1987b, 169-177.

Haugen, R., E. Ortiz and E. Arjona, "Market Efficiency: Mexico versus the U.S.", Journal of Portfolio Management , Fall 1985, 28-32.

5.1 Alternative Forms of the Efficient Market Hypothesis

The efficient market hypothesis (**EMH**), formally associated with the work of Fama (1970) and its restatement (Fama (1991)), states that the market prices of assets reflect all available information about the assets. The appropriate definition of 'all available' varies and gives rise to alternative testable implications of the EMH. The information set over which markets are said to be efficient can be one of three possibilities, each giving rise to a broader (i.e. stronger) version of the EMH which encompasses more information relevant to asset prices/returns:

5.1.1 *The Weak EMH*

This states that all information contained in historical prices and firm characteristics (such as size, book value etc.) is incorporated in the actual (current) price. All **historical information** is thus reflected in the observed market price. Notice that no claim is made about the inclusion of any other type of information, nor about the speed with which information is incorporated in asset prices.

5.1.2 *The Semi-Strong EMH*

This includes the history of past prices as well as all **publicly available information** about assets' returns, i.e. all disclosures, announcements and reports which are available to all market participants. All such information is reflected in the current price.

5.1.3 *The Strong EMH*

In addition to the the above, the strong form of the EMH also includes all **privately available** information on the assets, i.e. information proprietary to particular analysts and managers. The most common such information is private forecasts of asset returns.

Since each information set is a proper subset of the next one, the strong EMH clearly implies the semi-strong EMH, and in turn the semi-strong form implies the weak form, but not vice versa. A key testable implication of all three forms of the EMH is that investors, trading on the respective information set, should **be unable to realize average excess returns above the normal rate**. For the weak form applied to the stock market, the information set includes the past history of stock prices as well as companies' general characteristics and seasonal (timing) effects. The latter clearly should not have a persistent impact if markets are efficient: empirical 'anomalies' such the January, weekend and holiday effects fall under this category. For the semi-strong form, the relevant information includes all types of companies' announcements, giving rise to the large literature on *event studies*. Finally, strong-form market efficiency also comprises private information available to some group of investors, a quantity which is much more difficult to measure.

Indeed, if all available information is already reflected in the current price, then no particular bit of information can be used to predict future price movements. This general principle applies to the two extreme types of investment analysis, **fundamental** and **technical analysis**, as well as to all intermediate types. Fundamental analysts focus on analyzing the evolution of assets' fundamentals in order to predict their value. In the context of the stock market, this includes reports of earnings, P/E ratios, dividends, etc. In the context of FX markets, this includes all macroeconomic developments and statistics which are considered to have an influence on nominal and real exchange rate dynamics. Most of this information is in the public domain and so belongs to the information set of semi-strong EMH. In contrast, technical analysts analyze the past history of prices (included in the weak form of the EMH), aiming to discover empirical regularities (trends, patterns etc.) which will allow them to predict future price movements.

Therefore, although research of both types is essential in order for the information to get incorporated into prices--e.g., a positive reevaluation by a consensus of market analysts of a firm's future value after a good earnings forecast leads to a rise in its share price--the results of this research should not yield consistently higher returns over the market average. It should be clear that empirical support for the returns' unpredictability proposition depends crucially (*a*) on the particular definition of 'market average', and (*b*) on the *speed* with which new information is incorporated in market prices. None of these two aspects are explicitly addressed by the EMH, but they have a direct impact on the measurements of excess returns.

Regarding the former aspect, the typical modeling assumption is that the normal returns on an asset are constant over time, although recently there has been increasing interest in equilibrium models with time-varying normal asset returns. Abnormal (i.e. excess) asset returns are then computed as the difference between the asset's observed return and its normal return, and forecasts of excess returns are constructed based on the chosen information set. If excess returns are unforecastable, e.g. if they follow a random walk, then the EMH is not rejected. However, a potentially serious problem arises in interpreting a rejection. Since any test of EMH must assume an equilibrium model that defines normal asset returns, if market efficiency is rejected it could be *either* because the market is truly inefficient, *or* because an incorrect equilibrium model has been assumed. This *joint hypothesis* problem implies that we can never be sure of rejecting the EMH as such.

A common **market anomaly** which seems to contradict the EMH is the observed higher excess returns of small firms over those of large firms. In the context of the CAPM, the market portfolio is common to large and small stocks, so any problem in estimating their betas will be equally affecting both sizes of stocks. Consequently there will be no bias in excess returns. However, suppose that the asset pricing technique used is a multi-factor model. Then, it is possible that the risk associated with some factors entering the small stock relationship is systematically underestimated (for example, the risk associated with the probability of survival, which is generally greater for small firms than for large firm). This would imply an underestimate of the expected return for small stocks, and as a result they would appear to yield higher excess returns than large stocks. This example is just one manifestation of the general problem with testing the EMH described above, namely not being able to properly define and measure equilibrium, or **normal returns**.

Regarding aspect (*b*)--the speed of information transmission from news to prices--it may be argued that *perfect* market efficiency at every point in time is an unrealistic benchmark because of various frictions in the short-run. Even in theory, it has been shown that excess returns will exist and may be persistent if there are costs of gathering and processing information. These returns are necessary in order to compensate investors for the expenses associated with the collecting and processing of information. They should not, therefore, be regarded as abnormal when information-related costs are properly accounted for.

Clearly, in large and liquid markets the relatively small magnitude of information costs is unlikely to justify substantial excess returns. It is, however, difficult to say how small excess returns should be, even if information costs could be precisely measured. In contrast, in smaller and/or less liquid markets we could expect higher ‘normal’ excess returns. This is the subject of the reading by Haugen et al. (1985) who compare the NYSE with the Mexican stock exchange. Overall, the notion of *relative*, rather than absolute, market efficiency may be more useful for practical purposes. We now turn to a more detailed discussion of the informational content of price changes and its implications for market efficiency.

5.2 The Arrival of New Information and Price Changes

If current market prices reflect all available (privately and publicly known) information, then observed price changes should only reflect unpredictable information. Therefore the time series of actual returns should be a **random walk**, or, more generally, the process generating asset returns should be a *martingale*, i.e. a ‘fair game’. The random walk is a more restrictive process, as its values all have to be independent (i.e. having zero covariance and correlation) draws from identical probability distributions. These testable hypotheses have led to a whole literature on market predictability based on tests to assess whether stock market time series follow random walk processes. If they do, then the EMH holds and attempts to beat the market are futile: the best forecast of future values is today’s asset price. In other words, the deviations of actual prices/returns from expected prices/returns are unpredictable.

In practice, the different information sets available to different classes of investors imply that the results of such empirical tests have been mixed, with weak support for the random walk over the short-run, but strong evidence for the EMH over longer horizons. In terms of portfolio management techniques, such an approach founded on EMH has led to **passive fund management** strategies which concentrate on replicating the performance of a representative market index such as the FTSE-100. The opposite approach, sometimes referred to as **active fund management**, invests on the premise that some assets are mispriced, i.e. that their returns are forecastable.

If acquiring and processing information is costly, it may be that observed price changes actually reflect a certain part, but not all, of the new information. It will thus take time for *all* of the new information to get incorporated in the price: the information transfer will not be instantaneous. As investors gradually acquire more of the new information, the price gradually rises or falls to its new level. Thus, over relatively short intervals of time, it is possible to observe systematic trends and patterns in asset prices and returns--i.e. *predictable* sequences of price changes--without violating the EMH and the principle of investor rationality. In other words, it is still the case that equilibrium prices reflect all available information (the EMH), but observed prices need not instantaneously be equilibrium prices if information aquisition and processing are costly.

EXAMPLE

A consensus of stock market analysts currently discovers a trend in a share price which, if followed, would lead to a 25% rise over the next 2 months, say from £16 to £20. If the consensus is wide, then as soon as the trend is discovered it disappears because the stock price jumps instantaneously from, say, £17.50 to £20. However, if information about the trend is incorporated gradually into the price--because there is only a minority of analysts who have immediate access to the information--then the trend may persist, although probably not for the full 2-month period. The information cost argument therefore suggests that excess returns are better associated with the strong form of the EMH, since private information is by definition not available to all market participants at the same time.

5.3 Testing the Predictability of Asset Returns

As argued above, all tests of returns predictability are subject to the criticisms concerning variations in the definition of normal returns and the simultaneous availability of all relevant information. However, the number of possible ‘special cases’ is so large that analysts are bound to find some systematic patterns and trends in the time series of prices/returns at some frequency, whether or not these can be used to earn excess returns. To the extent that markets are efficient, trends should not be recurring, or they should be reverting.

In the case of empirical regularities observed at the NYSE, such as higher than average returns on Friday and lower returns on Mondays, the fact that they seem to be persistent and well-understood indicates that transaction costs may not make trading strategies based on them profitable. The same may be said of the January effect, involving monthly returns in January which are higher than average annual returns. This effect is more evident for small stocks and has been found in several stock markets. Since it is clearly not associated with a particular risk factor affecting small stocks, it cannot be captured by a standard asset pricing model such as the CAPM. Despite alternative explanations to do with investors' tax incentives or firms' microstructure, such an effect presents an average investor with a profitable trading strategy (buy in December/sell in January).

Analysts often employ simple statistical tests to test for the predictability of asset returns from their own past history (weak-form EMH) at various horizons. These are based on evaluating the autocorrelation coefficients between past and present returns in OLS regressions of the general form:

$$r_t = b_0 + \sum_{i=1}^{L} b_{t-i} r_{t-i} + e_t$$

The number of lags (e.g. days, weeks, months) is given by L. If L=1, then the value of the regression coefficient b_{t-L} is the first-order autocorrelation coefficient of the asset's returns. As we increase the value of the lag, we are testing for predictability at progressively longer horizons. Correspondingly, tests are classified as searching for short and long-term predictability. However, for any given time series, we are likely to compute the whole spectrum of autocorrelations, e.g. from 1 to 30 (or less) for daily data, 1 to 52 for weekly data and 1 to 12 for monthly data. Note that the *square* of the order-L autocorrelation coefficient represents the amount of current variation in returns explained by the variation in returns L lags (periods) ago. For example, a first-order autocorrelation coefficient of 0.3 on a monthly data set of stock returns implies that 0.09 (9%) of monthly variation in returns is explained by last month's variation.

If the weak form of the EMH holds, then we should find all orders of the autocorrelation coefficients to be statistically insignificant from zero. Notice also that the coefficient b_0 corresponds to the asset's expected return. In regression (1) this has been assumed to be fixed--it is therefore unrelated to past returns. More generally it may be time-varying, thus capturing the positive relationship between assets' actual returns and their averages over time. There are two classes of outcomes other than an insignificant autocorrelation structure in our general regression. Either autocorrelations are significantly nonzero at small lags (i.e. in the short-run) and insignificant at longer lags, o*r* they are insignificant in the short-run and significant in the long-run. In the former case it may be said that the assets' returns have **short memory**, in the latter that they have **long memory**.

In general, both cases are evidence against the EMH, although short memory does not necessarily violate market efficiency if information acquisition and processing is costly. We would, in fact, expect to find a stronger nonzero autocorrelation structure in stock markets which are less liquid and/or have higher transaction costs, reflecting the fact that these markets are likely to be less efficient. This prediction has found empirical support in a number of studies. In either long or short memory, we also expect to have a positive relationship between negative, or successively alternating, autocorrelation signs (from positive to negative and vice versa) and market efficiency. Such autocorrelation patterns suggest that excess returns are **mean-reverting** to the expected market return over time, whereas positive autocorrelations indicate that the trends in asset returns are persistent.

If successive autocorrelations are positive, so that short term trends in returns are persistent, then it is possible to create a profitable trading strategy by using what is known as a *filter rule*. The filter takes advantage of the fact that when returns move upwards from time t to t+1, they are likely to move upwards from t+1 to t+2 as well. The filter thus specifies a particular size of positive (negative) deviation from the current price which triggers a buy (sell) decision, respectively, followed by the reverse position to sell (buy) in order to realize the excess return. These sizes are often implemented in the form of automatic (program) trading. However, to the extent that markets are efficient, the deviations cannot grow too large, so empirically the optimal filter size has been found to be of the order of 1-2% at most. We should also keep in mind that the alternative to a filter rule is the time-honored buy and hold strategy, which clearly involves less transaction costs. Therefore, a successful filter rule has to outperform the buy and hold strategy by an amount increasing in the transaction costs.

Another empirical finding is that the positive autocorrelation structure of portfolios is more significant than that of individual stocks. Moreover, the positive autocorrelation is stronger for small stock portfolios than for large stock portfolios. In principle, this regularity is not difficult to explain using our standard arguments of diversification. A portfolio of many stocks reduces unsystematic (unique) risk to zero and, in the limit, is left with only systematic (market) risk. In Lecture 2 we have seen how, given a sufficiently negative correlation between the individual assets' returns, it is possible for the total risk (variance) of a 2-asset portfolio to be less than the risk of either of its assets. That example generalizes to many assets, and explains the relatively higher predictability of portfolios' returns compared to individual stocks'. However, it should be noted that some of the explained portfolio variability could be due to the fact that usually not all of a portfolio's assets are traded *all* the time. The assets that are traded less frequently are, therefore, likely to be affected by the same new piece of information at different times. Since these influences will be in the same direction, such assets are contributing to the positive autocorrelation of the portfolio itself.

5.4 Strong-Form Efficiency, Market Rationality and Crashes

In contrast to the relatively uncontroversial literature using event studies of returns before and after firms' announcements to test for the semi-strong form of the EMH, tests of strong market efficiency have prompted considerable debate. Trading on inside information is prohibited in all financial centers and insiders have to declare their transactions to the appropriate stock exchange regulating body. However, the small amount of evidence which does not suffer from selection bias suggests that insiders--such as fund and firm managers who hold more than a certain percentage of the outstanding stock--do not earn excess returns over the market's expected return. In empirical studies, this finding shows as *either* (i) a small and insignificant correlation coefficient between actual and forecasted returns, *or* (ii) small and insignificant autocorrelation coefficients between past and present forecasting ability. The latter is usually measured as the percentage of the change in actual returns explained by the change in forecast returns.

Another interesting finding is that excess returns can be earned based on the *aggregate* recommendations of a group of forecasters, rather than on an individual forecasting firm's (stock broker's) recommendation. Thus, although individual information may prove to be incorrect, averaging as much information as possible provides a better forecast of the market's direction. This suggests that information has a value itself, and brings us to the issue of **market rationality**. Markets are said to be perfectly rational, as opposed to informationally efficient, if changes in asset prices reflect changes in the present value of future cash flows. If they do not, i.e. if it can be shown that market volatility is greater or smaller than the volatility of the underlying fundamentals affecting future cash flows, then markets are not perfectly rational.

A body of empirical research has centered on (i) volatility tests comparing actual and theoretical (driven only by future dividends) prices (Shiller 1981, 1984, 1986) and (ii) demostrations of investors' overreactions to particular events and winning/losing streaks (De Bondt and Thaler (1985, 1987), Thaler (1987a,b)). Overall, this evidence supports the idea that markets overreact to irrelevant events, and that the resulting actual volatility of prices and returns far exceeds that of their estimated theoretical levels. Therefore, markets are not perfectly rational. Note that this conclusion does not necessarily contradict the EMH in either its semi-strong or strong forms. However, it does contradict the weak form: there is more than the past history of returns that drives their current levels. In particular, information embodied in market expectations--whether public or private--appears to be a very important determinant of asset price dynamics.

Market irrationality, defined as above, and the role of expectations were evident both in the crash of October 19, 1987 (Black Monday) and in the recent highly volatile trading sessions in most major exchanges. Clearly, such stock price volatility cannot be explained by underlying changes in fundamentals affecting future cash flows. Rather, it may be due to a drastic downward revision of investors' expectations caused by an increase in the degree of market uncertainty (e.g. about possible overvaluation, the effects of the Asian currency crises etc). To the extent that such revisions are affected by the perceived average expectation of all investors, they can be easily reversed, and as a result they contribute to increasing market volatility. Investors' expectations are thus sometimes *self-fulfilling*.

In closing the discussion on market efficiency it should be stressed again that any evidence of predictability of asset returns (or excess returns) that we may find in a particular study may be due to an improper definition of 'normal' returns. In that sense the EMH as such can never be rejected. A particular illustration of this problem is the voluminous literature on the **equity premium puzzle**. Faced with empirical evidence that equity offers substantially higher expected returns over time than fixed income assets, researchers have to account for the difference in terms of the higher risk attached to the underlying risk factors. However, as measurements of the risk factors are often not sufficient to explain the magnitude of that difference, the usual way out has been to assume that the expected return on stocks is time varying, i.e. that the coefficient b_0 in the autocorrelation regression (1) is not fixed. Any predictability in asset returns that is then found can be attributed to the time-varying expected return rather than to informational inefficiency.