Chapter 12

Efficient Capital Markets

At the intuitive level, the efficient market hypothesis (EMH) states that it is impossible to "beat the market": The return you earn by investing in financial assets is proportional to the risk you assume. If certain assets had high returns compared to their risks, investors - who are constantly searching for and analyzing information about companies, commodities, economic indicators, etc. - would rush to buy these assets, pushing up prices until returns were "proportional" to the risk. In this way information about future returns of financial assets are quickly incorporated into prices which in other words fully and instantaneously reflect all available relevant information.

A first attempt to make this intuitive statement more rigorous is to elaborate a little bit on the definition of "relevant information" and to interpret 'reflecting information' as an inability to earn excessive returns. This leads us to the famous *degrees of efficiency:*

1. *Weak-form efficiency.* No investor can earn excess returns developing trading rules based on historical price or return information. In other words, the information in past prices or returns is not useful or relevant in achieving excess returns.

2. *Semistrong-form efficiency.* No investor can earn excess returns from trading rules based on any publicly available information. Examples of publicly available information are annual reports of companies, investment advisory data such as "Heard on the Street" in the Wall Street Journal, or ticker tape information.

3. *Strong-form efficiency.* No investor can earn excess returns using any information, whether publicly available or not.

Of course, we could define efficiency with respect to any information set I_t: No trader can earn excess returns between time t and $t + 1$ given the information I_t.

Once we start thinking about this definition, however, we note that it is

still well short of being rigorous. What do we mean by "risk" and by return being "proportional to risk"? And if indeed prices (in some sense) reflect all available information why are banks and other financial institutions paying analysts to find information, and traders to 'look for arbitrage' in the market? In this note we will analyze the problem of defining market efficiency as follows:

First, we consider the problem of defining excess return. The only 'model free' definition of this concept would be an arbitrage opportunity. Other definitions are invariably linked to a particular model for security returns. If we are unable to find arbitrage opportunities can we then ever reject the hypothesis of 'market efficiency'?

Second, we consider some statistical properties of prices which are often thought of as being consequences of EMH. We try to give some rigorous definitions of these properties and (briefly) discuss if they are in any way necessarily linked to market efficiency.

Third, we look at anomalies and a very interesting challenge to market efficiency which attempts to show that stock prices are 'too volatile' in the sense that they vary much more than accounted for by changes in economic fundamentals.

Finally we mention some attempts to make the definition of market efficiency completely rigorous in a general equilibrium context.

12.1   Excess returns

First, we try to make "excess returns" a little more precise: As a first attempt we could say that for any security P it must be the case that

$$E\left[R_{t+1} \mid I_t\right] = 1 + r_t \tag{12.1}$$

where r_t is the riskless rate of return at time t and

$$R_{t+1} = \frac{P_{t+1}}{P_t}.$$

But having heard of CAPM and risk aversion we immediately object to the use of r_t: Assets with high risk (as measured e.g. through their β's) should (and do!) earn more than the riskless rate. So we restate (12.1) as

$$E\left[R_{t+1} | I_t\right] = 1 + r_t^P \tag{12.2}$$

where r_t^P is some return which is suitable for asset P. Note that we may rewrite (12.2) as

$$\frac{1}{1 + r_t^P} E\left[P_{t+1} \mid I_t\right] = P_t,$$

which states that the properly discounted price is a martingale. If markets are arbitrage free, such a discount factor will exist, so the requirement (12.2) is really nothing more than a statement of no arbitrage. In practice the problem is to find a good model for r_t^P - indeed we spent a lot of time looking at CAPM which tried to do just that. But if we have a model for r_t^P and we see violations of (12.2), it would seem more natural to reject our model for r_t^P than to reject market efficiency. Assume for instance that we use CAPM to model r_t^P. The literature often refers to a rejection of (12.2) as a rejection of "a joint test of CAPM and market efficiency". Given the very strong assumptions needed to derive CAPM, this is a somewhat strange statement (although it is of course logically OK). I would argue that unless there is a truly compelling and extremely realistic model (in terms of assumptions describing markets accurately) for r_t^P, we should consider rejections of (12.2) as rejections of our choice of asset pricing model and not worry about the market efficiency implications. Some proponents of efficient markets would say that it is precisely the belief in efficient markets which causes us to reject our model and look further for a rational explanation of stock price returns, and this is of course a valid point. But can we ever reach a situation in which all conceivable models for rational behavior have been rejected and we have no choice but to reject the EMH?

The closest to this situation would be a case where the model for r_t^P follows from a no arbitrage condition. If, for example, we considered a portfolio of a written call, a put (both European with same exercise price and date) and a stock (the underlying security for both options), we know from put-call-parity that its return should equal the return on a zero-coupon bond. Violations of this relation which are large enough that traders may take advantage of it should clearly not persist in well functioning security markets. In other words, no arbitrage is a necessary condition for markets to be efficient. And only this provides a test of (12.2) which is not linked to a particular choice of asset pricing model. Another way of stating this is that we can use (12.2) to test for violations of relative pricing relations which we derive from an assumption of no arbitrage. Violations would seem to indicate an inefficiency of markets.

But to use (12.2) as a definition of efficiency is complicated when we are looking at "absolute" quantities like stock prices. We will never be able to reject efficiency from this alone - critics will always (justly) be able to argue that our choice of asset pricing model is wrong. A famous paper by DeBondt and Thaler[1] [1986] illustrates the problem well. They did the following:

- Record the returns of a very large number of stocks over a period of

---

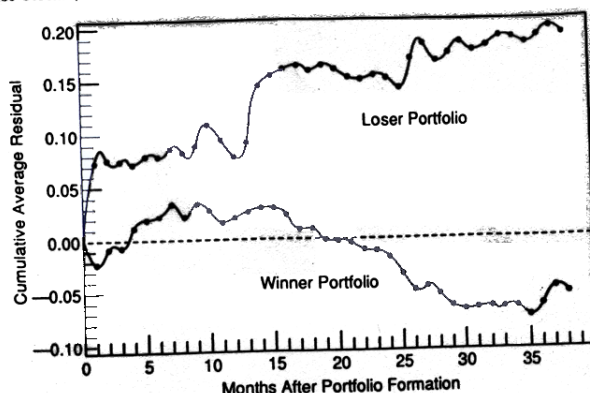[1]See Journal of Finance, July 1986, pp. 793-807.

Figure 12.1: The cumulative excess return (compared to the market) for the "winner" and "loser" portfolios.

(say) 36 months.

• At the end of the period, form one portfolio of "winners" i.e. pick (say) the 35 best performing stocks of the 36 month period.

• Form another portfolio of "losers" - i.e. pick (say) the 35 worst performing (but still alive!) stocks of the 36 month period.

• Compare the returns to that of the market.

The results are illustrated in Figure 12.1, which shows cumulative excess returns compared to the market for the two portfolios. We see that the loser portfolio has a (quite large) positive return over the market, whereas "the winners" are performing badly. Now we could be tempted to see this as a clear violation of "weak-form efficiency" because, intuitively, we are beating the market using only information of past prices. But are we defining "excess return" correctly then? We have seen in the exercises (in a discrete binomial version) that the beta of an option on a stock may be related to the beta of the stock. The equivalent of this in a Black-Scholes model where we model stocks as options on an underlying firm value gives us that

$$\beta_S = N(d_1)\frac{V}{S}\beta_V.$$

A "loser" stock will typically have experienced an increase in leverage. This in turn implies an increase in β_S and therefore the increased return on a "loser" follows directly from option pricing theory and CAPM. In other words, a more careful selection of asset pricing model at least gave us a possible explanation for the results of DeBondt and Thaler. Subsequent research (see for example Chopra, Lakonishok and Ritter[2] [1992]) seems to indicate that correcting for leverage effects does not destroy the result. But the controversy is far from over and the key issue is whether more advanced asset pricing models can be shown to explain the larger returns on losing stocks. Note, however, that the controversy illustrates the sense in which (12.2) is a "hard-to-test" definition of efficiency. We will have to reject it for all conceivable (and reasonable) asset pricing models before an inefficiency is established. This seems to be a tough task, especially considering results of Summers[3] [1986] which indicate that many test traditionally employed do not have very large power (i.e. ability to reject a false hypothesis).

12.2 Martingales, random walks and independent increments

As a motivation for this section consider the footnote in Brealey and Myers:

"When economists speak of stock prices as following a random walk, they are being a little imprecise. A statistician reserves the term *random walk* to describe a series that has a constant expected change each period and a constant degree of variability. But market efficiency does not imply that expected risks and expected returns cannot shift over time."

Clearly, we need to be more precise even about the statistician's definition if we are to discuss the notions precisely. The purpose of this section is to give precise content to such concepts as random walk, independent increments and martingales.

Definition 43 *The stochastic process $X = (X_1, X_2, \ldots)$ is a random walk if it has the form*

$$X_t = \sum_{i=1}^{t} \epsilon_i$$

where $\epsilon_1, \epsilon_2, \ldots$ is a sequence of independent, identically distributed random variables.

---

[2]Journal of Financial Economics, 1992, pp. 235-268

[3]Journal of Finance, July 1986, pp.591-601

We will normally assume that ϵ_i has finite variance (and hence finite expectation as well). If $E\epsilon_i = 0$ we say that the random walk is symmetric.

A weaker concept is that of a process with independent increments:

Definition 44 *The stochastic process* $X = (X_1, X_2, \ldots)$ *is a process with independent increments if for all* t, $\Delta X_t \equiv X_t - X_{t-1}$ *is independent of* $(X_1, X_2, \ldots, X_{t-1})$.

Note that a random walk is a process with independent increments, but a process with independent increments may have the distribution of the increment change with time. Weaker yet is the notion of orthogonal increments where we only require increments to be uncorrelated:

Definition 45 *The stochastic process* $X = (X_1, X_2, \ldots)$ *is a process with orthogonal increments if for all* t, $EX_t^2 < \infty$ *and for all* t, u

$$\mathrm{Cov}(\Delta X_t, \Delta X_u) = 0.$$

Finally a somewhat different (but as we have seen equally important) concept

Definition 46 *The stochastic process* $X = (X_1, X_2, \ldots)$ *is a martingale with respect to the filtration* $(\mathcal{F}_t)$ *if for all* t, $E|X_t| < \infty$ *and*

$$E(X_t | \mathcal{F}_{t-1}) = X_{t-1}.$$

If no filtration is specified it is understood that $\mathcal{F}_{t-1} = \sigma(X_1, X_2, \ldots, X_{t-1})$.

A symmetric random walk is a martingale and so is a process with mean zero independent increments. A martingale with finite second moments has orthogonal increments, but it need not have independent increments: Think for example of the following ARCH process:

$$X_t = \epsilon_t \sigma_t$$

where $\epsilon_1, \epsilon_2, \ldots$ is a sequence of i.i.d. $N(0, 1)$ random variables and

$$\sigma_t^2 = \alpha_0 + \alpha_1 x_{t-1}^2 \qquad \alpha_0, \alpha_1 > 0$$

where x_{t-1} is the observed value of X_{t-1}. The abbreviation ARCH stands for Autoregressive Conditional Heteroskedasticity which means (briefly stated) that the conditional variance of X_t given the process up to time $t-1$ depends on the process up to time $t - 1$.

It is clear from observation of market data, that stock prices have a drift and they are therefore not martingales. However it could be that by an appropriate choice of discount factor, as in (12.2), we get the martingale property. In fact, we know that this will be possible in an arbitrage free market (and for all types of financial securities), but then we are back to the situation discussed above in which the specification of the pricing relation becomes essential.

The independent increment property is often heard as a necessary condition for efficient markets, but research over the last decade has pointed to ARCH effects in stock prices, exchange rates and other financial securities.[4]

Another serious problem with all the properties listed above is that it is perfectly possible to construct sensible general equilibrium models (which in particular are arbitrage free) in which prices and returns are serially correlated (for example because aggregate consumption is serially correlated). Therefore, it seems unnatural to claim that random walks, independent increments and martingale properties of prices are in any way logical consequences of a definition of efficient markets. The most compelling consequence of an efficiency assumption of no arbitrage is that there exists an equivalent measure under which securities are martingales but this does not in general say much about how processes behave in the real world. Indeed, even in the case of futures contracts which would be very natural candidates to being martingales, we have seen that in models where agents are risk-averse, the equivalent martingale measure will in general be different from the empirically measure and hence the martingale property of futures prices is by no means a necessary condition for efficiency.

Also, note that bonds for example have predetermined 'final' value and therefore it is not at all clear that any of the above mentioned properties of price processes are even relevant for this class of securities.

12.3   Anomalies

Another way to challenge the EMH is to look for strange patterns in security prices which seem impossible to explain with any pricing model. Some such anomalies are *week-end effects* and *year-end effects:*

Evidence found in French (1980)[5] seems to suggest that returns on Mondays were significantly negative, compared to returns on other trading days. Although the effect is difficult to exploit due to transactions costs, it is still unclear why the effect exists. Possible explanations have tried to look at

---

[4]For an collection of papers see Engle: *ARCH*.Oxford University Press, 1995.

[5]Journal of Financial Economics, 1980, pp. 55-69.

whether there is a tendency for firms to release bad news on Friday afternoons after trading closes. But if this were the case traders should learn this an adjust prices accordingly earlier. The week-end effect is hard to explain - after all it does seem to suggest that traders who are going to trade anyway should try to sell on Fridays and buy on Mondays.

Another effect documented in several studies is the year-end effect which shows that stocks have a tendency to fall in December and rise in January. An obvious reason for such behavior could be tax considerations but it is still unclear if this explanation is sufficient or whether there is actually an effect which trading rules can take advantage of.

A famous anomaly is the *closed-end mutual fund discount.* A closed-end fund is a mutual fund which holds publicly traded assets, but which only issues a fixed number of shares and where shareholders can only sell their shares in the market. Since the assets are so easy to determine, so should the share price be. However, throughout a long period of time mutual funds seemed to sell at a discount: The values of the funds' assets seemed larger than the market value of shares. Small discounts could be explained by tax-considerations and illiquidity of markets, but no theory could explain a pattern (shown in Figure 12.2) like the one observed for the Tricontinental Corporation[6] during 1960-1986.

It seems that finally in 1986 prices adjusted, but that prices should "instantaneously" have reflected fundamental information seems implausible.

So a newer version of EMH stated in Malkiel [1990][7] could be that "pricing irregularities may well exist and even persist for periods of time, but the financial laws of gravity will eventually take hold and true value will come out." Certainly a weakening of our original version of efficiency.

12.4   Excess volatility.

One of the most interesting challenges to EMH argues that there is too much volatility in stock markets - more than can be explained by *any* sensible asset pricing theory. If indeed markets reflect all relevant information it should be the case that price movements and arrival of new information were somehow in harmony. Trading alone ought not to generate volatility. Several pieces of evidence suggest that trading índeed creates volatility:

1. October 19, 1987. The Dow Jones Industrial Average fell by more than 22% - a much bigger drop than any previous one-day movement. Yet extensive surveys among investment managers seem to suggest that no important
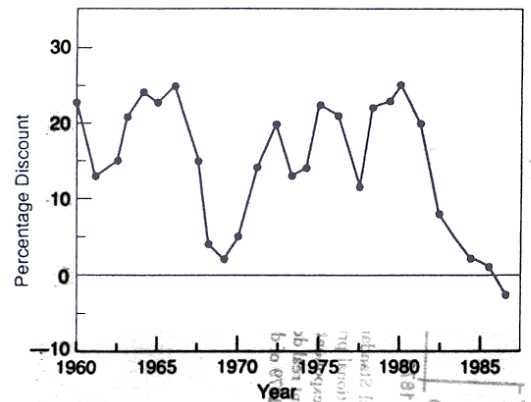
---

[6]See Lee, Shleifer and Thaler, Journal of Finance 1991, pp.75-109.

[7]*A Random Walk Down Wall Street.*

Figure 1 Percentage discount or premium of Tricontinental Corporation at the end of each year during 1960–1986.

The percentage discount is computed as 100 × (NAV − SP); where NAV is the per share net asset value and SP is the share price of the fund. The mean (median) of the percentage discount or premium is 14.43 (15.0). The maximum (minimum) value is 25.0 (−2.5) and the standard deviation is 8.56.

Figure 12.2: Percentage discount (= 100* (net asset value per share - share price of the fund)) at which Tricontinental Corporation was traded 1960-86.

news related to stock prices arrived that day.

2. In the second half of 1986 the stock exchange in New York closed for a series of Wednesdays to catch up on paperwork. Volatility between Tuesday's close and Thursday's opening of trade was significantly smaller than when the exchange was open on Wednesdays. As Thaler [1993][8] puts it, traders react to each others as well as to news.

3. A famous study on "Orange juice and weather" by Roll [1984][9] suggests that surprises in weather forecasts for the Florida area, where 98% of oranges used for orange juice are traded, are too small to explain variations in the futures price (whose main concern is the likelihood of a freeze)

But perhaps the most controversial attack on EMH is the one put forward by Shiller, and Leroy and Porter. Their line of analysis is that what determines stock prices must ultimately be some combination of dividends and earnings of the firm. Since stock price is an expected, discounted value of future quantities, the price should fluctuate less than the quantities themselves.

To understand their line of reasoning, consider the following example: Assume that the quantity x_t (which could be dividends) determines the stock

---

[8]R. Thaler (ed.) *Advances in Behavioral Finance,* Russell Sage Foundation, NY 1993
[9]American Economic Review, December 1984, pp. 861-880.

price through a present value relation of the form:

$$Y_t = \sum_{j=0}^{\infty} \beta^j X_t^e(j),$$

where β is some discount factor, $\beta \in]0,1[$, and

$$X_t^e(j) = E[X_{t+j}|I_t],$$

where I_t can - for simplicity - be thought of as the information generated by the process X up to time t. If we assume that X_t follows an autoregressive process of order one, i.e.

$$X_t - c = \phi(X_{t-1} - c) + \varepsilon_t$$

where $c > 0, |\phi| < 1$ and ε_t's are independent normally distributed random variables, then we can calculate the *coefficient of dispersion* of X_t,

$$CD(X_t) = \frac{\sigma(X_t)}{E(X_t)}$$

and the similar quantity for Y_t. We are interested in $\frac{CD(Y_t)}{CD(X_t)}$. Using a fact from time series analysis, which says that X_t be represented in the form

$$X_t - c = \sum_{j=0}^{\infty} \varphi^j (\varepsilon_{t-j})$$

we find

$$EX_t = c \quad VX_t = \frac{1}{1 - \varphi^2}$$

and

$$EY_t = \frac{c}{1 - \beta} \quad VY_t = \frac{VX_t}{(1 - \beta\varphi)^2}$$

and hence

$$\frac{CD(Y_t)}{CD(X_t)} = \frac{1 - \beta}{1 - \beta\varphi} < 1$$

by our choice of parameters. In other words, the coefficient of dispersion for actual prices is less than that of X_t (dividends, earnings). What Shiller, and Leroy and Porter show is that this result holds for a large class of processes for X *and* that the relationship is severely violated for observed data.

Figure 1

Note: Real Standard and Poor's Composite Stock Price Index (solid line p) and *ex post* rational price (dotted line p^*), 1871–1979, both detrended by dividing a long-run exponential growth factor. The variable p^* is the present value of actual subsequent real detrended dividends, subject to an assumption about the present value in 1979 of dividends thereafter. Data are from Data Set 1, Appendix.
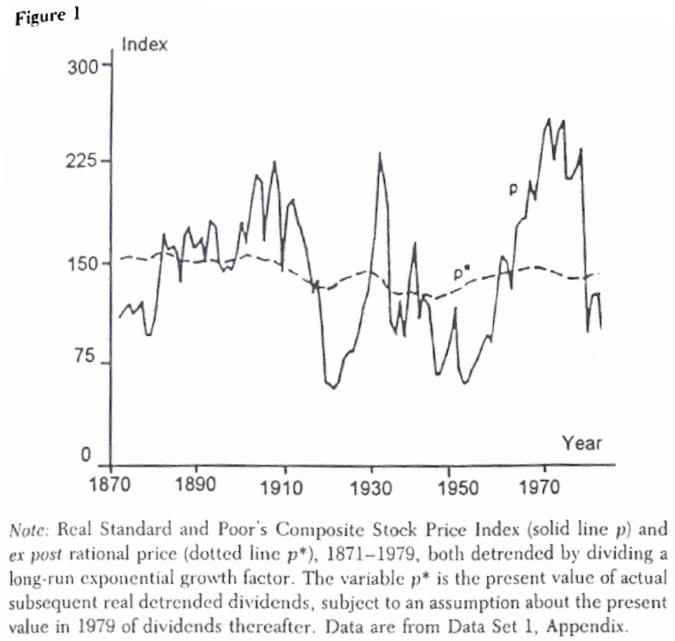
Figure 12.3: Variability of a detrended S&P index (full curve) compared to the variability of dividends (dotted line).

A useful bound also derived in these authors' work is the variance bound on the perfect foresight price

$$P_t^* = \sum_{k=0}^{\infty} X_{t+k} \prod_{j=0}^{k} \gamma_{t+j}$$

(where γ_{t+j} is the discount factor between time j and $j+1$ as recorded at time t) which "knows" all the paid dividends, and the actual price

$$P_t = E_t P_t^*$$

Shiller argues that the bound $V(P_t) \leq V(P_t^*)$ is grossly violated in practice, as illustrated in Figure 12.3.

Marsh and Merton [1986] criticize the work of Shiller and argue that with non-stationary dividend policies, the variance bound relations cannot hold and should in fact be reversed. Future versions of this chapter will discuss this controversy!

12.5   Informationally efficient markets are impossible

The title of this section is the title of a paper by Grossman and Stiglitz[10] (1980) Their first paragraph in that work sums up an essential problem with the assumption of informational efficiency:

"If competitive equilibrium is defined as a situation in which prices are such that all arbitrage profits are eliminated, is it possible that a competitive economy can always be in equilibrium? Clearly not, for then those who arbitrage make no (private) return from their (privately) costly activity. Hence the assumption that all markets, including that for information, are always in equilibrium and always perfectly arbitraged are inconsistent when arbitrage is costly." Grossman and Stiglitz go on to build a model in which gathering (costly) information is part of an equilibrium among informed and uninformed traders.Without going into details of their model, we note here that in equilibrium, prices reflect some but not all of the information. There is a fraction of traders who are informed and spend money to gather information, and a fraction of traders who are uninformed but try to learn as much as possible by observing prices (which they know reflect partly the knowledge that the informed traders have). If some of the informed traders give up information gathering, prices will reflect less information and there will be an incentive for non-informed traders to become informed because the costs of information gathering will more than compensated for by gains from trade. If some uninformed traders in the equilibrium situation decide to become informed prices will reflect more information and there will in fact be an incentive for some of the informed to give up information gathering and 'free-ride' on the information gathered by others which is reflected in prices. Clearly, a situation in which all information is reflected in prices is impossible.

This paper has a number of interesting results that we will not go into here. Suffice it to say that it contains one of the most interesting attempts to give a formal definition of efficiency which does not have the inherent logical problems of the original definitions.

There are numerous other more rigorous attempts to define efficiency. A typical line of approach is to model a situation with asymmetric information and ask whether the equilibrium price (and possibly the portfolio holdings of individuals) would be altered if all of the information held by traders was given to all traders simultaneously. If prices and allocations would not change in this situation, then it would be fair to say that prices reflected all

---

[10]American Economic Review, vol 70, pp. 393-408.

information.

These models will be discussed in a future version of the notes.