fractiles of the posterior distributions of β_i for each of the stocks, as was done for IBM.

One apparently direct solution to this problem is to work with a sample period that covers more than five years of monthly data. We can see from equations (26) and (27) that in sampling from the assumed stationary bivariate normal distribution of $\tilde{R}_{it}$ and $\tilde{R}_{mt}$, the variance of the sampling distribution of b_i decreases as the sample size increases. Thus a larger sample would seem to be the most direct way to reduce uncertainty about β_i. The validity of this approach depends, however, on the assumption that the joint distribution of $\tilde{R}_{it}$ and $\tilde{R}_{mt}$ is stationary through time, and especially on the implication of this assumption that β_i itself is stationary through time. If this is not true, then a larger sample does not necessarily imply a more reliable estimate of the value of β_i at the end of the sampling period.

The evidence of Blume (1968), Gonedes (1973), and L. Fisher (1970) indicates that over long periods, the β_i values of individual stocks do indeed change. The work of Gonedes and Fisher further indicates that with monthly data, the assumption that β_i is constant is a reasonable approximation for periods of up to seven years. With more than seven years of data, the estimates of the β_i of individual securities are likely to be less reliable than if shorter periods are used. With monthly data, the optimal estimation period is apparently five to seven years.

III. Conclusions

It seems that, at least for individual securities, we must learn to live with substantial uncertainty about the values of β_i. For many purposes, the problem is not serious. When we conduct tests requiring estimates of β_i, it is often possible to work with estimates for portfolios rather than individual securities, and it turns out that the β_p's of portfolios can be estimated far more reliably than those of individual securities. This is a matter we shall study in more detail when the need arises.

CHAPTER 5

Efficient Capital Markets

Much of the recent literature in finance is concerned with capital market efficiency. This chapter introduces the theory and discusses tests. The ideas and tests of them reappear in later chapters.

I. An Efficient Capital Market: Introduction

An efficient capital market is a market that is efficient in processing information. The prices of securities observed at any time are based on "correct" evaluation of all information available at that time. In an efficient market, prices "fully reflect" available information.

An efficient capital market is an important component of a capitalist system. In such a system, the ideal is a market where prices are accurate signals for capital allocation. That is, when firms issue securities to finance their activities, they can expect to get "fair" prices, and when investors choose among the securities that represent ownership of firms' activities, they can do so under the assumption that they are paying "fair" prices. In short, if the capital market is to function smoothly in allocating resources, prices of securities must be good indicators of value.

The statement that prices in an efficient market "fully reflect" available information conveys the general idea of what is meant by market efficiency, but the statement is too general to be testable. Since the goal is to test the extent to which the market is efficient, the proposition must be restated in a testable form. This requires a more detailed specification of the process of price formation, one that gives testable content to the term "fully reflect."

The process of price formation described below is far from the most general model that can be used to give testable content to the theory of capital market efficiency. The goals are (a) to present a simple model but one that is nevertheless sufficient to illustrate the problems that arise in testing market efficiency and (b) to describe and give some critical perspective on the types of tests that are commonly done.

II. An Efficient Capital Market: Formal Discussion

Assume that all events of interest take place at discrete points in time, $t-1$, t, $t+1$, etc. Then define

ϕ_{t-1} = the set of information available at time $t-1$, which is relevant for determining security prices at $t-1$.

ϕ^m_{t-1} = the set of information that the market uses to determine security prices at $t-1$. Thus ϕ^m_{t-1} is a subset of ϕ_{t-1}; ϕ^m_{t-1} contains at most the information in ϕ_{t-1}, but it could contain less.

$p_{j,t-1}$ = price of security j at time $t-1$, $j = 1, 2, \ldots, n$, where n is the number of securities in the market.

$f_m(p_{1,t+\tau}, \ldots, p_{n,t+\tau} | \phi^m_{t-1})$ = the joint probability density function for security prices at time $t+\tau (\tau \geqslant 0)$ assessed by the market at time $t-1$ on the basis of the information ϕ^m_{t-1}.

$f(p_{1,t+\tau}, \ldots, p_{n,t+\tau} | \phi_{t-1})$ = the "true" joint probability density function for security prices at time $t+\tau (\tau \geqslant 0)$ that is "implied by" the information ϕ_{t-1}.

To keep the notation manageable, the security prices $p_{1,t+\tau}, \ldots, p_{n,t+\tau}$ that appear as arguments in f and f_m are taken to be the prices of the securities at time $t+\tau$, plus any interest or dividend payments at $t+\tau$. The prices $p_{1,t-1}, \ldots, p_{n,t-1}$, are just actual prices at time $t-1$.

The set of information ϕ_{t-1} available at time $t-1$ includes what might be called the "state of the world" at time $t-1$: e.g., current and past values of any relevant variables, like the earnings of firms, GNP, the "political climate," the tastes of consumers and investors, etc. Since ϕ_{t-1} includes the past history of all relevant variables, ϕ_{t-1} includes ϕ_{t-2}; equivalently, ϕ_{t-2} is a subset of ϕ_{t-1}. In addition to current and past values of relevant variables, ϕ_{t-1} is also assumed to include whatever is knowable about relationships among variables. This includes relationships among current and past values of the same or different variables, and also whatever can be predicted about future states of the world from the current state. In short, ϕ_{t-1}, the information available at $t-1$, includes not only the state of the world at $t-1$, but also whatever is knowable about the process that describes the evolution of the state of the world through time. We assume that one of the things that is knowable about the process is the implication of the current state of the world for the joint probability distributions of security prices at future times. Thus ϕ_{t-1} is assumed to imply the joint density functions $f(p_{1,t+\tau}, \ldots, p_{n,t+\tau} | \phi_{t-1})$, $\tau = 0, 1, 2, \ldots$

The process of price formation at time $t-1$ is then assumed to be as follows. On the basis of the information ϕ^m_{t-1}, the market assesses a joint distribution of security prices for time t, $f_m(p_{1t}, \ldots, p_{nt} | \phi^m_{t-1})$. From this assessment of the distribution of prices at t, the market then determines appropriate current prices, $p_{1,t-1}, \ldots, p_{n,t-1}$, for individual securities. The appropriate current prices are determined by some model of market equilibrium—that is, by a model that determines what equilibrium current prices should be on the basis of characteristics of the joint distribution of prices at t. The term "equilibrium" has its usual economic meaning. A market equilibrium at time $t-1$ is achieved when the market sets prices $p_{1,t-1}, \ldots, p_{n,t-1}$ for individual securities at which the demand for each security by investors is equal to the outstanding supply of the security. In other words, a market equilibrium implies a market-clearing set of prices for individual securities.

When we say that "the market" assesses a joint distribution of security prices for time t and then uses the characteristics of its assessed distribution to determine equilibrium prices for securities at $t-1$, we speak metaphorically. To say that "the market" does something is just a convenient way of summarizing the decisions of individual investors and the way these decisions interact to determine prices. The metaphor allows us to save for the end of the chapter, when the issues can be better appreciated, the discussion of some of the subtle and not too subtle simplifications of the world that are built into the model.

In our model of price formation, the hypothesis that the capital market is efficient is stated as

$$\phi_{t-1}^m = \phi_{t-1}; \tag{1}$$

that is, ϕ_{t-1}^m, the information that the market uses to determine security prices at $t-1$, includes all the information available. Market efficiency also implies that

$$f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m) = f(p_{1t}, \ldots, p_{nt}|\phi_{t-1}); \tag{2}$$

that is, the market understands the implications of the available information for the joint distribution of ~~returns~~. Since ϕ_{t-1}, the set of available information, includes whatever is knowable about the process that describes the evolution of the state of the world through time, equation (1) can be taken to imply (2). Stating the two conditions separately, however, emphasizes that market efficiency means that the market is aware of all available information and uses it correctly.

Having correctly assessed the joint distribution of prices for t, the market then uses some model of equilibrium to set prices at $t-1$. The model says what the current prices of securities, $p_{1,t-1}, \ldots, p_{n,t-1}$, should be in light of the correctly assessed joint distribution of security prices for t. In this sense, both the joint density function $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$ and the current prices $p_{1,t-1}, \ldots, p_{n,t-1}$ that are based on this joint density function "fully reflect" all the information available at $t-1$.

Tests of market efficiency are concerned with whether or not the market does correctly use available information in setting security prices. Most common are tests that try to determine whether prices fully reflect specific subsets of information. For example, one possible source of information about future prices is the history of past prices and returns on securities. A nontrivial segment of the empirical literature on efficient markets is concerned with whether current security prices fully reflect any information in past prices and returns. Other sources of publicly available information are also fertile ground for tests of market efficiency. For example, there are studies of the adjustment of stock prices to the information in a stock split, a merger, an earnings announcement, the announcement of a new issue of securities by a firm, and so forth. In these tests, the goal is to determine whether prices adjust fully and instantaneously to the public announcement of the event of interest. Finally, another sort of test of market efficiency is concerned with whether there are individuals or groups—for example, managers of mutual funds—who are adept at investment selection in the sense that their choices reliably provide higher returns than comparable choices by other investors. If prices always fully reflect available information, this sort of investment adeptness is ruled out. For if such adeptness exists, it implies that some investors either have access to information that is not utilized by the market in setting prices or that they are better able to evaluate available information than the market. In either case, the market is not efficient.

The process of price formation in an efficient market, as described so far, is not sufficient for such tests of market efficiency. All we have said is that an efficient market correctly uses all available information in assessing the joint distribution of future prices, which is the basis of current equilibrium prices. Since we cannot observe $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$, we cannot determine whether (2) holds, and so we cannot determine whether the real-world capital market is efficient. Equations (1) and (2) are formal notation for the statement that prices in an efficient market fully reflect available information, but this is not sufficient to make the statement testable.

What the model lacks is a more detailed specification of the link between $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$ and $p_{1,t-1}, \ldots, p_{n,t-1}$. We must specify in more detail how equilibrium prices at $t-1$ are determined from the characteristics of the market-assessed joint distribution of prices for t. Some model of market equilibrium, however simple, is required. This is the rub in tests of market efficiency. Any test is simultaneously a test of efficiency and of assumptions about the characteristics of market equilibrium. If the test is successful—that is, if the hypothesis that the market is efficient cannot be rejected—then this also implies that the assumptions about market equilibrium are not rejected. If the tests are unsuccessful, we face the problem of deciding whether this reflects a true violation of market efficiency (the simple proposition that prices fully reflect available information) or poor assumptions about the nature of market equilibrium.

It turns out that a few simple models of market equilibrium produce many successful tests of market efficiency or, more precisely, many successful joint tests of market efficiency and of the models of market equilibrium. We now discuss the most popular models and tests of market efficiency derived from them.

III. Four Models of Market Equilibrium

Four basic models of market equilibrium are used in tests of market efficiency. We discuss them in order of complexity.

A. *Expected Returns Are Positive*

The joint distribution $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$ of security prices for time t assessed by the market at time $t-1$ implies a marginal distribution $f_m(p_{jt}|\phi_{t-1}^m)$ for the price at t of any security j. This marginal distribution has mean or expected value $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$.* The first model of market equilibrium simply says that at any time $t-1$ the market sets the price of any security j in such a way that the market's expected return on the security from time $t-1$ to time t is positive.

Formally, the one-period return on security j from time $t-1$ to t is

$$\tilde{R}_{jt} = \frac{\tilde{p}_{jt} - p_{j,t-1}}{p_{j,t-1}}. \tag{3}$$

At time $t-1$ the market assesses a probability distribution on $\tilde{p}_{jt}$ given by the density function $f_m(p_{jt}|\phi^m_{t-1})$. A distribution for the return $\tilde{R}_{jt}$ is not defined, however, until the market sets $p_{j,t-1}$. The model of market equilibrium which we are discussing posits that the market always sets $p_{j,t-1}$ so that the mean of the resulting distribution of $\tilde{R}_{jt}$ is strictly positive. That is, the market always sets $p_{j,t-1}$ so that, given its assessment of the expected price at t, $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$,

$$E_m(\tilde{R}_{jt}|\phi^m_{t-1}) = \frac{E_m(\tilde{p}_{jt}|\phi^m_{t-1}) - p_{j,t-1}}{p_{j,t-1}} > 0. \tag{4}$$

Equivalently, the market sets $p_{j,t-1}$ at a value less than its assessment of the expected future price, $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$.

Suppose now that we join this model of market equilibrium with the proposition that the market is efficient. Market efficiency says that in assessing distributions of future prices, the market uses all available information and uses it correctly:

$$f_m(p_{jt}|\phi^m_{t-1}) = f(p_{jt}|\phi_{t-1}), \tag{5}$$

which implies

$$E_m(\tilde{p}_{jt}|\phi^m_{t-1}) = E(\tilde{p}_{jt}|\phi_{t-1}) \tag{6}$$

$$E_m(\tilde{R}_{jt}|\phi^m_{t-1}) = E(\tilde{R}_{jt}|\phi_{t-1}). \tag{7}$$

In words, market efficiency says that at time $t-1$ the market correctly assesses the distribution of the price of any security for time t, which means that the expected value of the future price assessed by the market is the true expected value, which in turn means that when the market sets the prices of securities at time $t-1$, its assessment of the expected return on any security is the true expected return. If the market sets prices so that equation (4) holds, then the true expected return on any security is always positive:

*Tildes (~) are used to denote random variables. When referring to any specific value of a random variable, the tilde is dropped. Thus, $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$ is the expected value of the random variable $\tilde{p}_{jt}$, but we write $f_m(p_{jt}|\phi^m_{t-1})$ to denote the density function for specific values of the variable.

$$E(\tilde{R}_{jt}|\phi_{t-1}) > 0. \tag{8}$$

This is not to say that a positive return on security j will be observed at t. The return observed at t will be the result of a drawing from $f(p_{jt}|\phi_{t-1})$, and the drawing may yield a negative return. Rather, the hypothesis that the market is efficient (prices correctly reflect available information), when combined with a model of market equilibrium which says that $E_m(\tilde{R}_{jt}|\phi^m_{t-1}) > 0$ (the market sets current prices so that its expected returns on securities are positive), implies that at time $t-1$ the true expected return on any security j, $E(\tilde{R}_{jt}|\phi_{t-1})$, is positive.

If the market is efficient and if this model of market equilibrium is correct, then any investor or market analyst who disagrees with the market and posits a negative expected return on a security is incorrect. Many stock market analysts feel that they can identify times when expected returns on individual securities or on the market, as represented by some portfolio of securities, are negative. These analysts would agree with the proposition that the market always sets prices so that its assessed expected returns $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ are positive. But they would disagree with the proposition that the market is efficient. They feel that in setting prices, the market sometimes neglects relevant information or draws incorrect inferences from it, so that sometimes the true expected returns $E(\tilde{R}_{jt}|\phi_{t-1})$ are negative. They feel that they see more information or are better able to analyze available information than the market.

Such analysts are potentially a fertile source of tests of market efficiency. If they record the times when they assess negative expected returns on securities, then one can simply compute the returns that are later realized. One or a few such observations are not much evidence for or against market efficiency; but as a history of the predictions of an analyst is built up, a reliable average return for periods when he assesses negative expected returns can be obtained. If the average is negative and if the sample of predictions is sufficiently large to make the negative average return a low-probability event if true expected returns are positive, then we can conclude that the analyst is able to identify periods when true expected returns $E(\tilde{R}_{jt}|\phi_{t-1})$ are negative. If we are willing to stick by the model of market equilibrium which says that the market always sets prices so that its expected returns $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ are positive, then the predictions of the analyst establish that the market sometimes either neglects available information in setting prices or analyzes information incorrectly. In either case, the analyst is living evidence for the existence of market inefficiency.

The model summarized by equations (4) to (8) has been used to test the claims of one group of analysts about market inefficiency. This group, collectively known as chartists or "technical" analysts, claims that market prices only react slowly and over fairly long periods to new information. If new information implies a price increase, the increase will be spread across time, as will any decrease in prices that is implied by negative information. This slow adjustment process posited by the chartists is in sharp contrast to the theory of efficient markets. When the market is efficient, prices fully reflect available information, which means that the market adjusts prices fully and instantaneously when new information becomes available.

The chartists further claim that the reaction of the market to new information is so slow that one need not be concerned with the information itself. By studying patterns in the sequence of past prices, they argue, one can learn how the price of the security tends to react to new information. The patterns in the price sequence will be strong enough and will recur frequently enough for a trained eye to predict the future price movement of a security on the basis of its recent past movement and knowledge of the typical patterns in the price behavior of the security. In short, the chartists claim the market is inefficient in the sense that in setting prices, the market does not even take full account of the obvious information in the historical behavior of prices.

Given the expected return model summarized by (4), an empirical confrontation between the claims of the chartist and those of the theory of capital market efficiency is easily devised. The basic proposition of the chartist is that because the market adjusts slowly to new information, price movements tend to persist. When prices have moved up in the recent past, one can expect them to continue to move up, and there is likewise persistence in downward price movements. Consider the following trading rule, suggested by Alexander (1961; 1964) and close in spirit to the various trading rules proposed by chartists. If the price of a security moves up at least y percent, buy and hold the security until its price moves down at least y percent from a subsequent high, at which time simultaneously sell and go short.* The short position is maintained until the price rises at least y percent above a subsequent low, at which time one covers the short position and goes long. Moves less than y percent in either direction are ignored. Such a system is called a y percent filter. Its sequence of successive long and short positions formalizes the proposition of the chartists that upward price movements tend to persist and to be followed by downward movements, which also tend to persist and to be followed by upward movements, and so on.

*In the jargon of the capital market, when one buys a security, this is known as going long. When one owns the security, this is called a long position in the security. The opposite of a long position is a short position. Selling short involves borrowing a security from someone who has a long position in the security, with the borrower promising to return the security to the lender at some future date and to pay to the lender any dividends or interest that are paid on the security while the short position is "open," that is, before the securities are returned. Upon borrowing the security, the borrower or short-seller immediately sells the security in the market. He then repurchases the security in the market when it comes time to return it to the lender, and in this way "closes" or "covers" his short position. If the price of the security falls during the period the short position is open, and if it falls by more than the amount of any dividends or interest paid on the security, then the short-seller profits. Otherwise he loses.

A short sale is equivalent to issuing a security with precisely the characteristics of the security that is sold short. Short-selling is thus a device whereby investors can issue securities that are identical to those issued by firms—assuming, of course, that the investor can deliver on the promises involved in the short sale. These concepts are discussed in Chapter 7.

If the capital market is efficient and if the market sets prices so that its expected returns are positive, then filter rules are nonsense. If the market correctly uses available information and if it sets prices so that expected returns are positive, then the best trading rule for any security is to buy and hold. If the market is efficient, then the buy-and-hold strategy has higher expected returns or profits than any strategies that involve periods when the security is not held or, like the filter rules, involve periods when the security is sold short. In contrast, the chartist would say that because the market does not correctly use available information, there are periods when true expected returns are negative. This implies that there are strategies for trading in a security that have higher expected returns or profits than the buy-and-hold strategy. Most chartists would believe that some of the filters could systematically beat a buy-and-hold strategy.

Tests of filter rules are reported by Alexander (1961; 1964) and by Fama and Blume (1966). To present their results would involve a long discussion of technical details, none of which would be useful in any of our future work. We shall simply discuss conclusions and let the reader check the original sources. Thus, Alexander (1961; 1964) reports extensive tests of filter rules using daily data on price indexes from 1897 to 1959 and filters from 1 to 50 percent. In his final paper on the subject, Alexander concludes (1964, p. 351):

> In fact, at this point I should advise any reader who is interested only in practical results, and who is not a floor trader and so must pay commissions, to turn to other sources on how to beat buy and hold.

Further evidence is provided by Fama and Blume (1966), who compare the profitability of various filters to a buy-and-hold strategy for daily data on the individual stocks of the Dow-Jones Industrial Average. (The data are those discussed in Chapter 1.) Fama and Blume conclude that for the most part their evidence is in favor of buy and hold, and they reject the hypothesis that there is any important information in past prices that the market neglects in setting current prices.

Looking hard, however, one can find evidence in the filter tests of both Alexander and Fama-Blume that is inconsistent with capital market efficiency, if efficiency is interpreted in a strict sense. In particular, the results for very small filters (1 percent in Alexander's tests and 0.5, 1.0, and 1.5 percent in the tests of Fama-Blume) indicate that it is possible to devise trading schemes based on very short-term (preferably intraday, but at most daily) price swings that on average outperform buy and hold. The average profits on individual transactions from such schemes are minuscule, but they generate transactions so frequently that over longer periods and ignoring commissions they outperform buy and hold by a substantial margin. These results are evidence of persistence in very short-term price movements of the type posited by the chartists.

When one takes account of even the minimum trading costs that would be generated by small filters, however, their advantage over a buy-and-hold strategy disappears. For example, even a floor trader—that is, a person who owns a seat on the New York Stock Exchange—must pay clearinghouse fees on his trades that amount to about 0.1 percent per turnaround transaction (sale plus purchase). Fama and Blume show that because small filters produce such frequent trades, these minimum trading costs are sufficient to wipe out the advantage of the small filters over buy and hold. Strictly speaking, then, the filters uncover evidence of market inefficiency, but the departures from efficiency do not seem sufficient for any trader to reject the hypothesis that the market is efficient so far as his own activities are concerned.

Remember that no null hypothesis, such as the hypothesis that the market is efficient, is a literally accurate view of the world. It is not meaningful to interpret the tests of such a hypothesis on a strict true-false basis. Rather, one is concerned with testing whether the model at hand is a reasonable approximation to the world, which can be taken as true, at least until a better approximation comes along. What is a reasonable approximation depends on the use to which the model is to be put. For example, since traders cannot use filters to beat buy and hold, it is reasonable for them to assume that they should behave as if the market were efficient, at least for the purposes of trading on information in past prices.

B. *Expected Returns Are Constant*

The filter tests are the only tests of market efficiency based on the model of market equilibrium which simply assumes that expected returns are positive. Somewhat more common are tests based on a model in which the expected return is assumed to be constant through time. Specifically, at time $t-1$ the market assesses a joint distribution for security prices at time t, $f_m(p_{1t}, \dots, p_{nt}|\phi^m_{t-1})$, which implies a distribution $f_m(p_{jt}|\phi^m_{t-1})$ for the price of security j at t, and this distribution has mean or expected value $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$. Having assessed $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$, the market then sets the price of the security at $t-1$ so that the expected return on the security from $t-1$ to t is equal to some constant, call it $E(\tilde{R}_j)$, which is the same for every period. Formally, at every time $t-1$, the market sets the current price of security j so that, given its assessment of the expected value of the future price $E_m(\tilde{p}_{jt}|\phi^m_{t-1})$,

$$E_m(\tilde{R}_{jt}|\phi^m_{t-1}) = \frac{E_m(\tilde{p}_{jt}|\phi^m_{t-1}) - p_{j,t-1}}{p_{j,t-1}} = E(\tilde{R}_j). \tag{9}$$

The model says that $E(\tilde{R}_j)$ is constant through time, but different securities are allowed to have different expected returns, based perhaps on differences in risk, and some may even have negative expected returns.

If the market is also efficient—that is, if it correctly uses all available information to assess $f_m(p_{jt}, \dots, p_{nt}|\phi^m_{t-1})$—then this assessed distribution is the true distribution $f(p_{jt}, \dots, p_{nt}|\phi_{t-1})$, which implies that equations (5) to (7) hold. Combining (7) with the assumption of a constant expected return, we have

$$E(\tilde{R}_{jt}|\phi_{t-1}) = E_m(\tilde{R}_{jt}|\phi^m_{t-1}) = E(\tilde{R}_j). \tag{10}$$

In words, at any time $t-1$ the market sets the price of security j in such a way that its assessment of the expected return on the security, $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$, is the constant $E(\tilde{R}_j)$. Since an efficient market correctly uses all available information, $E(\tilde{R}_j)$ is also $E(\tilde{R}_{jt}|\phi_{t-1})$, the true expected return on the security.

This particular combination of a model of market equilibrium with market efficiency has a directly testable implication. There is no way to use any information available at time $t-1$ as the basis of a correct assessment of the expected return on security j which is other than $E(\tilde{R}_j)$. If the market is efficient and sets prices so that the expected return on security j is constant through time, then any market analyst who assesses an expected return for security j that is different from $E(\tilde{R}_j)$ is necessarily incorrect. But if the analyst systematically shows an ability to identify periods when the expected return on security j is not equal to $E(\tilde{R}_j)$, and if we insist on the model of market equilibrium which says that the market sets prices so that its expected return on security j is always $E(\tilde{R}_j)$, then the predictions of the analyst are evidence that the market does not correctly use all available information in setting prices. In this case, equation (7) does not hold, and the market is inefficient.

For the statistically sophisticated, equation (10) implies that for all ϕ_{t-1}, $E(\tilde{R}_{jt}|\phi_{t-1})$, the regression function of $\tilde{R}_{jt}$ on ϕ_{t-1} is the constant $E(\tilde{R}_j)$. Thus, if one takes any elements from the set of information available at $t-1$ and then estimates the regression of $\tilde{R}_{jt}$ on these information variables, all the coefficients except for the intercept should be indistinguishable from zero. If some of the variables have nonzero coefficients, (10) must be rejected; that is, the joint hypothesis that the market is efficient and that it sets prices so that equilibrium expected returns are constant through time is rejected.

Tests of market efficiency based on the assumption that equilibrium expected returns are constant have focused primarily on one subset of ϕ_{t-1}, the potential information about current expected returns that appears in time series of past returns. If the market is efficient and equilibrium expected returns are constant through time, the past returns on security j are a source of information about $E(\tilde{R}_j)$, which, after all, is unknown.* If the market is efficient, however, the past returns are not a source of information about the expected value of the deviation of $\tilde{R}_{jt}$ from $E(\tilde{R}_j)$. For any sequence of past returns $R_{j,t-1}, R_{j,t-2}, \ldots$, the conditional expected value

$$E(\tilde{R}_{jt}|R_{j,t-1}, R_{j,t-2}, \ldots) = E(\tilde{R}_j).$$

In words, if the market is efficient, there is no way to use any information available at time $t - 1$ as the basis for a correct assessment of an expected value of $\tilde{R}_{jt}$ which is different from the assumed constant equilibrium expected return $E(\tilde{R}_j)$. Since part of the information available at $t - 1$ is the time series of past returns, there is no way to use the past returns as the basis for a correct assessment of the expected return from $t - 1$ to t which is other than $E(\tilde{R}_j)$.

This proposition is easily tested with a tool introduced in Chapter 4. If the correct assessment of the expected value of $\tilde{R}_{jt}$ is $E(\tilde{R}_j)$, then for any $R_{j,t-\tau}$

$$E(\tilde{R}_{jt}|R_{j,t-\tau}) = E(\tilde{R}_j); \tag{11}$$

that is, there is no way to use the past return $R_{j,t-\tau}$ as the basis of a current assessment of an expected value of $\tilde{R}_{jt}$ which is other than $E(\tilde{R}_j)$. In formal terms, the regression function of $\tilde{R}_{jt}$ on $R_{j,t-\tau}$, $E(\tilde{R}_{jt}|R_{j,t-\tau})$, is the constant $E(\tilde{R}_j)$.

To test this proposition, we introduce an alternative hypothesis which says that the regression function is linear in $R_{j,t-\tau}$:

$$E(\tilde{R}_{jt}|R_{j,t-\tau}) = \delta_\tau + \gamma_\tau R_{j,t-\tau}. \tag{12}$$

From Chapter 4 we recognize γ_τ as the autoregression or autocorrelation coefficient for lag τ, also denoted $\rho(\tilde{R}_{jt}, \tilde{R}_{j,t-\tau})$. Thus market efficiency, in

*If we are willing to assume that the distribution of $\tilde{R}_{jt}$ is constant through time, then frequency distributions of historical returns are information about the distribution of $\tilde{R}_{jt}$. This is the basis of the empirical work in Chapter 1. The assumption that the distribution of $\tilde{R}_{jt}$ is constant through time is, of course, stronger than the assumption that the mean of the distribution is constant.

combination with the assumption that equilibrium expected returns are constant through time, implies that the autocorrelations of the returns on any security j are zero for all values of the lag τ.

In Chapter 4 we looked at sample autocorrelations of monthly returns for common stocks on the NYSE and concluded that the autocorrelations were close to zero. There we used the sample autocorrelations to test the assumption of random sampling that underlies the statistical inferences drawn from market model coefficient estimates. Now that we want to examine sample autocorrelations to test the hypothesis that the market is efficient, it is well to look at more of them.

Table 5.1, taken from Fama (1965), shows sample autocorrelations of daily returns for each of the 30 Dow-Jones Industrials, for time periods that vary slightly from stock to stock but usually run from about the end of 1957 to September 26, 1962. (The data are discussed in Chapter 1.)* For each stock, the table shows sample autocorrelations for lags of from one to ten days. Recall from Chapter 4 that when the true autocorrelation is zero, the sampling distribution of the sample autocorrelation, $r(\tilde{R}_{jt}, \tilde{R}_{j,t-\tau})$, is approximately normal, with approximate mean and standard deviation

$$E[r(\tilde{R}_{jt}, \tilde{R}_{j,t-\tau})] \doteq -1/(T-\tau)$$
$$\sigma[r(\tilde{R}_{jt}, \tilde{R}_{j,t-\tau})] \doteq \sqrt{1/(T-\tau)},$$

where T is the number of returns in the sample.

In Table 5.1 the sample autocorrelations that are at least two standard deviations to the left or to the right of $-1/(T-\tau)$ are indicated by asterisks. The values of sample autocorrelations so marked might be regarded as extreme in the sense that they are low-probability events if the true autocorrelations are zero. Of the 30 sample autocorrelations between successive daily returns ($\tau = 1$), 11 are extreme in this sense and 9 of these 11 are positive. Moreover, 22 of the 30 sample autocorrelations between successive daily returns are positive. Since market efficiency says that the true autocorrelations between successive returns are zero, one might interpret the results as evidence against market efficiency: there seems to be positive autocorrelation between successive daily returns.

There are several reasons why one might conclude that the results in Table 5.1 are not sufficient to overturn the hypothesis of market efficiency. First, the 30 autocorrelations for lag $\tau = 1$ (or for any other specific lag) are not independent. From our study of the market model in Chapter 4 we know that returns on individual securities are all related to the return on the mar-

*These are continuously compounded returns, but recall from Chapter 1 that continuously compounded daily returns are numerically close to simple returns.

TABLE 5.1

Sample Autocorrelations of Daily Return on the Dow-Jones Industrials for Lags $\tau = 1, 2, \ldots, 10$

STOCK	LAG (τ) 1	2	3	4	5	6	7	8	9	10	T
Allied Chemical	.017	-.042	.007	-.001	.027	.004	-.017	-.026	-.017	-.007	1223
Alcoa	.118*	.038	-.014	.022	-.022	.009	.017	.007	-.001	-.033	1190
American Can	-.087*	-.024	.034	-.065*	-.017	-.006	.015	.025	-.047	-.040	1219
AT&T	-.039	-.097*	.000	.026	.005	-.005	.002	.027	-.014	.007	1219
American Tobacco	.111*	-.109*	-.060*	-.065*	.007	-.010	.011	.046	.039	.041	1283
Anaconda	.067*	-.061*	-.047	-.002	.000	-.038	.009	.016	-.014	-.056	1193
Bethlehem Steel	.013	-.065*	.009	.021	-.053	-.098*	-.010	.004	-.002	-.021	1200
Chrysler	.012	-.066*	-.016	-.007	-.015	.009	.037	.056*	-.044	.021	1692
Du Pont	.013	-.033	.060*	.027	-.002	-.047	.020	.011	-.034	.001	1243
Eastman Kodak	.025	.014	-.031	.005	-.022	.012	.007	.006	.008	.002	1238
General Electric	.011	-.038	-.021	.031	-.001	.000	-.008	.014	-.002	.010	1693
General Foods	.061*	-.003	.045	.002	-.015	-.052	-.006	-.014	-.024	-.017	1408
General Motors	-.004	-.056*	-.037	-.008	-.038	-.006	.019	.006	-.016	.009	1446
Goodyear	-.123*	.017	-.044	.043	-.002	-.003	.035	.014	-.015	.007	1162
International Harvester	-.017	-.029	-.031	.037	-.052	-.021	-.001	.003	-.046	-.016	1200
International Nickel	.096*	-.033	-.019	.020	.027	.059*	-.038	-.008	-.016	.034	1243
International Paper	.046	-.011	-.058*	.053*	.049	-.003	-.025	-.019	-.003	-.021	1447
Johns Manville	.006	-.038	-.027	-.023	-.029	-.080*	.040	.018	-.037	.029	1205
Owens Illinois	-.021	-.084*	-.047	.068*	.086*	-.040	.011	-.040	.067*	-.043	1237
Procter and Gamble	.099*	-.009	-.008	.009	-.015	.022	.012	-.012	-.022	-.021	1447
Sears	.097*	.026	.028	.025	.005	-.054	-.006	-.010	-.008	-.009	1236
Standard Oil (Calif.)	.025	-.030	-.051*	-.025	-.047	-.034	-.010	.072*	-.049*	-.035	1693
Standard Oil (N. J.)	.008	-.116*	.016	.014	-.047	-.018	-.022	-.026	-.073*	.081*	1156
Swift and Co.	-.004	-.015	-.010	.012	.057*	.012	-.043	.014	.012	.001	1446
Texaco	.094*	-.049	-.024	-.018	-.017	-.009	.031	.032	-.013	.008	1159
Union Carbide	.107*	-.012	.040	.046	-.036	-.034	.003	-.008	-.054	-.037	1118
United Aircraft	.014	-.033	-.022	-.047	-.067*	-.053	.046	.037	.015	-.019	1200
U.S. Steel	.040	-.074*	.014	.011	-.012	-.021	.041	.037	-.021	-.044	1200
Westinghouse	-.027	-.022	-.036	-.003	.000	-.054*	-.020	.013	-.014	.008	1448
Woolworth	.028	-.016	.015	.014	.007	-.039	.013	.003	-.088*	-.008	1445

*Sample autocorrelation is at least two standard deviations to the left or to the right of its expected value under the hypothesis that the true autocorrelation is zero.

Source: Eugene F. Fama, "The Behavior of Stock Market Prices," *Journal of Business* 38 (January 1965): 72.

ket. For current purposes, this means that the sample autocorrelations of the returns on individual securities all reflect to some extent the sample autocorrelation of the return on the market. Thus, it is not necessarily surprising that for a given lag the sample autocorrelations in Table 5.1 are predominantly positive or negative.

Even if we are willing to conclude that there is evidence in Table 5.1 of positive dependence between successive daily returns, it is reasonable to argue that the evidence is not sufficient to reject the hypothesis that the market is efficient. With 1,200 to 1,700 observations per stock, a sample autocorrelation as small as .05 is for some stocks more than two standard deviations to the right of its expected value under the hypothesis that the true value of the coefficient is zero. Thus, a sample coefficient as small as .05 is extreme in the statistical sense, and so is fairly convincing statistical evidence against the hypothesis that the true value of the coefficient is zero. Suppose, however, that the true value of an autocorrelation is as much as twice .05, or $\rho(\tilde{R}_{jt}, \tilde{R}_{j,t-\tau}) = .10$. The square of the autocorrelation between $\tilde{R}_{jt}$ and $\tilde{R}_{j,t-\tau}$ is the proportion of the variance of $\tilde{R}_{jt}$ that can be attributed to the linear regression function relationship between $\tilde{R}_{jt}$ and $\tilde{R}_{j,t-\tau}$. Thus, the squared autocorrelation can be interpreted as a measure of the information that $\tilde{R}_{j,t-\tau}$ carries for $\tilde{R}_{jt}$; it tells how much we can reduce the variance of $\tilde{R}_{jt}$ if we have exact knowledge about the linear regression function relationship between $\tilde{R}_{jt}$ and $\tilde{R}_{j,t-\tau}$. In these terms, an autocorrelation $\rho(\tilde{R}_{it}, \tilde{R}_{j,t-\tau}) = .10$ says that $\tilde{R}_{j,t-\tau}$ doesn't carry much information about $\tilde{R}_{jt}$, since only 1 percent of the variance of $\tilde{R}_{jt}$ can be attributed to the linear relationship between $\tilde{R}_{jt}$ and $\tilde{R}_{j,t-\tau}$. Thus, even though the true autocorrelation is nonzero, it is close enough to zero for us to conclude that market efficiency is a reasonable description of the world.

The evidence in Table 5.1 is actually good support for the hypothesis that the market is efficient. The sample autocorrelations are close to zero in magnitude and in terms of "proportion of variance explained." Although the true autocorrelations might be nonzero, given the large sample sizes and the small observed autocorrelations it is unlikely that the true autocorrelations are much different from zero, which means that is is unlikely that the deviation of $\tilde{R}_{j,t-\tau}$ from $E(\tilde{R}_j)$ carries much information about the deviation of $\tilde{R}_{jt}$ from $E(\tilde{R}_j)$. Thus, at least with respect to potential information in past daily returns, the hypothesis that the market is efficient seems to be a good approximation to the world.

For each of the 30 Dow-Jones Industrial stocks, Table 5.2 shows sample autocorrelations of monthly returns for lags $\tau = 1, 2, 3$, that is, for returns one, two, and three months apart. The time period is July 1963–June 1968. Although the sample autocorrelations in Table 5.2 are generally close to zero, they are also more variable and thus larger in absolute value than those for the daily returns in Table 5.1. This is to be expected, since the sample size in Table 5.2 is only $T = 60$, whereas in Table 5.1 the samples include from 1,200 to 1,700 daily returns. As a consequence, the standard deviations for the autocorrelations in Table 5.2 are about .13, while those for the autocorrelations in Table 5.1 are generally less than .03. Thus, the results in both tables are consistent with market efficiency, but those for the larger samples in Table 5.1 give a much more precise feeling for how close the true autocorrelations of returns are to zero.

TABLE 5.2
Autocorrelations of Monthly Returns on the Dow-Jones Industrials for July 1963–June 1968

COMPANY	$r(R_{jt}, R_{j,t-1})$	$r(R_{jt}, R_{j,t-2})$	$r(R_{jt}, R_{j,t-3})$
Allied Chemical	.017	−.236	.144
Alcoa	−.306*	.076	.172
American Can	−.061	.003	.162
AT&T	−.117	.096	.173
American Tobacco	−.282*	−.058	.156
Anaconda	−.097	−.170	.156
Bethlehem Steel	−.034	−.044	−.101
Chrysler	.207	−.020	−.093
Du Pont	−.076	−.023	.234
Eastman Kodak	.098	−.175	.088
General Electric	−.028	−.093	−.006
General Foods	−.001	−.023	.070
General Motors	−.091	−.060	.254
Goodyear	−.034	−.294*	−.114
International Harvester	−.050	.236	.140
International Nickel	−.196	−.043	−.058
International Paper	−.010	−.367*	.089
Johns Manville	.080	−.128	−.113
Owens Illinois	.139	−.176	−.288*
Procter and Gamble	−.193	.193	−.077
Sears	−.105	−.020	.253
Standard Oil (Calif.)	−.111	.093	.207
Standard Oil (N. J.)	−.025	−.032	.242
Swift and Co.	.020	.005	−.020
Texaco	.076	−.148	.004
Union Carbide	−.080	.022	.047
United Aircraft	−.143	.136	.159
U.S. Steel	−.113	.023	.067
Westinghouse	.099	−.005	−.094
Woolworth	.078	.062	.098
Averages	−.044	−.016	.065

*Sample autocorrelation is at least two standard deviations to the left or to the right of its expected value under the hypothesis that the true autocorrelation is zero.

The success of the tests of market efficiency based on autocorrelations is somewhat fortuitous. The tests derive from a model of market equilibrium in which the equilibrium expected return on any security is constant through time. If this assumption is incorrect, tests of market efficiency based on autocorrelations could fail even though the market is efficient. For example, suppose the equilibrium expected return on security j, $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$, instead of being constant at the value of $E(\tilde{R}_j)$, tends to wander around $E(\tilde{R}_j)$, which we now interpret as the long-run average value of $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$. Moreover, suppose, as indicated in Figure 5.1, $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ tends to stay above

FIGURE 5.1
Hypothetical Behavior of Returns in an Efficient Market Where Equilibrium Expected Returns Wander Substantially Through Time

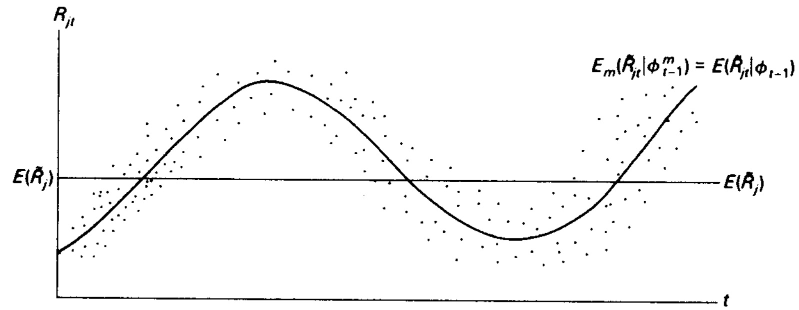

or below $E(\tilde{R}_j)$ for fairly long periods. If the market is efficient, then $E_m(\tilde{R}_{jt}|\phi^m_{t-1}) = E(\tilde{R}_{jt}|\phi_{t-1})$, the equilibrium return expected by the market is the true expected return. With an efficient market, the deviations of $\tilde{R}_{jt}$ from $E(\tilde{R}_{jt}|\phi_{t-1})$ would be more or less as shown in Figure 5.1; the current deviation is unpredictable from the past deviations. In this example, however, the deviation of $\tilde{R}_{jt}$ from $E(\tilde{R}_j)$ is quite predictable from the behavior of the most recent past deviations. Thus, if we used autocorrelations computed from an assumed constant average return to test market efficiency, we would conclude that the market is inefficient, when in fact the high autocorrelations in the returns would be due to the wandering of the equilibrium expected return. This sort of behavior of the equilibrium expected return is in no way ruled out by market efficiency.

The point, of course, is that any test of market efficiency is simultaneously a test of assumptions about market equilibrium. Since tests based on autocorrelations yield evidence consistent with the hypothesis that the market is efficient, the tests can also be interpreted as evidence consistent with the assumption that, at least for common stocks, equilibrium expected returns are constant through time. This does not say, however, that the evidence proves the assumption. Like any statistical evidence, it is at best consistent with the general model in the sense that it does not lead to rejection either of the hypothesis that the market is efficient or of the hypothesis that equilibrium expected returns are constant through time. This just means that, at least as far as the evidence from the autocorrelations is concerned, the hypotheses are reasonable models of the world. Like any models, however, they are just approximations that are useful for organizing our thinking about the phenomena of interest. They do not necessarily rule out other models which might also be reasonable and useful approximations.

For example, the evidence from the autocorrelations is also consistent with a world where the equilibrium expected return is not literally constant but where its variation is trivial relative to other sources of variation in the return on the security. Such a world might be as shown in Figure 5.2. The equilib-

FIGURE 5.2

Hypothetical Behavior of Returns in an Efficient Market Where Equilibrium Expected Returns Wander Through Time, but Only Slightly

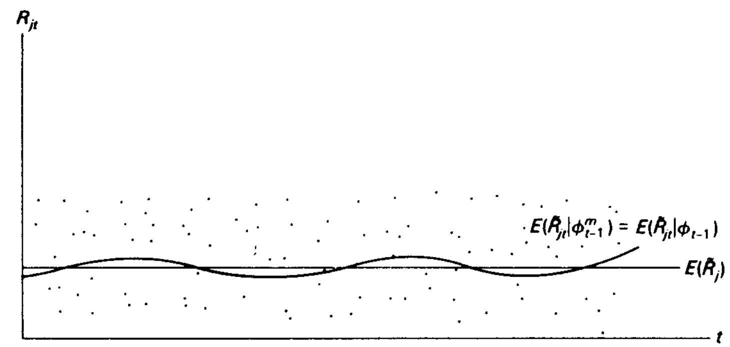

rium expected return $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ wanders through time about its long-run average value $E(\tilde{R}_j)$, but its wanderings are slight compared to those pictured in Figure 5.1. In Figure 5.2, the deviations of $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ from $E(\tilde{R}_j)$ are so small relative to the deviations of $\tilde{R}_{jt}$ from $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ that the wanderings of $E_m(\tilde{R}_{jt}|\phi^m_{t-1})$ would only be a source of slight positive autocorrelations in successive values of $\tilde{R}_{jt}$.

Thus, autocorrelations of $\tilde{R}_{jt}$ that are close to zero are consistent with a world where the market is efficient and equilibrium expected returns are constant through time. But they are also consistent with a world where the market is efficient and where equilibrium expected returns wander over time,

but not sufficiently to have any important effect on the autocorrelations of $\tilde{R}_{jt}$. Since we are primarily concerned with testing market efficiency, the choice between these two models of equilibrium expected returns is not important. All we need to say about equilibrium expected returns is that apparently they do not wander enough or in such a way as to invalidate autocorrelations as a tool for testing the hypothesis that the market is efficient, at least with respect to any information in historical returns.

C. *Returns Conform to the Market Model*

The tests of market efficiency discussed above are concerned with whether prices of securities fully reflect any information in past prices or returns. Historically, this was the first concern. When the results seemed to support the market efficiency hypothesis (see, for example, the various studies reported in Cootner 1964), attention turned to tests in which the concern was the speed of price adjustment to other publicly available information, like announcements of stock splits, earnings reports, new security issues, mergers, and so forth. As the tests of market efficiency moved in the direction of new information subsets, the models of market equilibrium on which the tests were based also became more complex.

THE MARKET MODEL AND MARKET EQUILIBRIUM

One of the models used extensively in more advanced tests of market efficiency is the market model of Chapters 3 and 4. In these chapters the market model is treated as an implication of the assumption that the joint distribution of security returns is multivariate normal. For current purposes, we formulate the model in part as an outgrowth of the process by which market equilibrium is attained.

The return on security j from time $t-1$ to time t is

$$\tilde{R}_{jt} = \frac{\tilde{p}_{jt} - p_{j,t-1}}{p_{j,t-1}} = \frac{\tilde{p}_{jt}}{p_{j,t-1}} - 1.0. \tag{13}$$

If the true distribution of $\tilde{p}_{jt}$, $f(p_{jt}|\phi_{t-1})$, is normal, then for any given price set by the market at time $t-1$, the distribution of $\tilde{R}_{jt}$, $f(R_{jt}|\phi_{t-1})$, will also be normal, since $\tilde{R}_{jt}$ is just a linear transformation of $\tilde{p}_{jt}$. Moreover, if the true joint distribution of the prices of different securities at time t, $f(p_{1t}, \ldots, p_{nt}|\phi_{t-1})$ is multivariate normal, the joint distribution of security returns, $f(R_{1t}, \ldots, R_{nt}|\phi_{t-1})$, is multivariate normal. According to Chapter 3, this implies that the market model holds. Thus,

$$E(\tilde{R}_{jt}|\phi_{t-1}, R_{mt}) = \alpha_j + \beta_j R_{mt} \tag{14}$$

with

$$\beta_j = \frac{\text{cov}(\tilde{R}_{jt}, \tilde{R}_{mt})}{\sigma^2(\tilde{R}_{mt})} \text{ and } \alpha_j = E(\tilde{R}_{jt}|\phi_{t-1}) - \beta_j E(\tilde{R}_{mt}|\phi_{t-1}). \qquad (15)$$

As in earlier chapters, the market portfolio m contains all common stocks on the NYSE, and R_{mt} is just the average of the returns on these stocks from $t-1$ to t. The return on security j at time t will not, of course, be equal to its conditional expected value as given by (14). The returns at t can be described in terms of the market model equation

$$\tilde{R}_{jt} = \alpha_j + \beta_j \tilde{R}_{mt} + \tilde{\epsilon}_{jt}, \qquad (16)$$

where the disturbance $\tilde{\epsilon}_{jt}$ is the deviation of $\tilde{R}_{jt}$ from its conditional expected value, and equation (14) implies

$$E(\tilde{\epsilon}_{jt}|\phi_{t-1}, R_{mt}) = 0.0. \qquad (17)$$

Equations (14) to (17) describe properties of the true bivariate normal joint distribution of $\tilde{R}_{jt}$ and $\tilde{R}_{mt}$, $f(R_{jt}, R_{mt}|\phi_{t-1})$, implied by the assumption that the joint distribution of security prices for time t, $f(p_{1t}, \ldots, p_{nt}|\phi_{t-1})$ is multivariate normal, and given the security prices set by the market at time $t-1$. The market is assumed to set prices at time $t-1$ in the usual way. That is, on the basis of the information ϕ_{t-1}^m, the market assesses a joint distribution on prices at time t, $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$, and then sets equilibrium prices at time $t-1$ on the basis of characteristics of $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$. If $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m)$ is the density function of a multivariate normal distribution, then $f_m(R_{jt}, R_{mt}|\phi_{t-1}^m)$ is the density function of a bivariate normal distribution, and the market's assessments imply market model equations, which, by analogy with (14) to (17), are

$$E_m(\tilde{R}_{jt}|\phi_{t-1}^m, R_{mt}) = \alpha_j^m + \beta_j^m R_{mt} \qquad (18)$$

$$\beta_j^m = \frac{\text{cov}_m(\tilde{R}_{jt}, \tilde{R}_{mt})}{\sigma_m^2(\tilde{R}_{mt})}, \text{ and } \alpha_j^m = E_m(\tilde{R}_{jt}|\phi_{t-1}^m) - \beta_j^m E_m(\tilde{R}_{mt}|\phi_{t-1}^m) \qquad (19)$$

$$\tilde{R}_{jt} = \alpha_j^m + \beta_j^m R_{mt} + \tilde{\epsilon}_{jt}^m \qquad (20)$$

$$E_m(\tilde{\epsilon}_{jt}^m|\phi_{t-1}^m, R_{mt}) = 0.0. \qquad (21)$$

To indicate that equations (18) to (21) describe the market model as seen by the market, subscript and superscript m's are included in the notation for the various parameters. As usual, if the market is efficient, the market's view is the correct view, so that $\phi_{t-1}^m = \phi_{t-1}$ and $f_m(p_{1t}, \ldots, p_{nt}|\phi_{t-1}^m) = f(p_{1t}, \ldots, p_{nt}|\phi_{t-1})$. Then the various parameters in equations (18) to (21) are identical to those in (14) to (17).

With all of the additional interpretation in terms of the process by which market equilibrium is attained, we still have only presented the market model as an implication of multivariate normality. In tests of market efficiency, an interpretation in economic terms is also given. The market return $\tilde{R}_{mt}$ is presumed to reflect information that becomes available at time t that, to a greater or lesser extent, affects the returns on all securities. When security prices are set at time $t-1$, $\tilde{R}_{mt}$ is unknown. It has a true distribution $f(R_{mt}|\phi_{t-1})$ which, in formal terms, is implied by the joint distribution of security prices, $f(p_{1t}, \ldots, p_{nt}|\phi_t)$, and the prices of securities set at $t-1$. But in economic terms, $f(R_{mt}|\phi_{t-1})$ is presumed to capture the uncertainty at time $t-1$ about information that will become available at time t which will affect the returns on all securities. The market model coefficient β_j in (14) to (16) therefore measures the sensitivity of the return on security j to $\tilde{R}_{mt}$ and thus, indirectly, to information about marketwide factors.

$t-1$

While $\tilde{R}_{mt}$ is presumed to reflect new information at time t that affects returns on all securities, the disturbance $\tilde{\epsilon}_{jt}$ in (16) is presumed to reflect information that becomes available at t that is more specifically relevant to the prospects of security j. The disturbance $\tilde{\epsilon}_{jt}$ has a true distribution $f(\epsilon_{jt}|\phi_{t-1}, R_{mt})$ that summarizes the uncertainty about the company-specific information which will become available at time t. The value of $\tilde{\epsilon}_{jt}$ observed at t will be a drawing from this distribution. Tests of market efficiency based on the market model are primarily concerned with the adjustment of prices to company-specific information, like earnings announcements, new issues of securities, stock splits, and so on. Thus, the tests concentrate on the behavior of $\tilde{\epsilon}_{jt}$ or, more precisely, on the behavior of estimates of $\tilde{\epsilon}_{jt}$.

Specifically, in empirical tests of market efficiency based on the market model, it is (implicitly) assumed that during each period the market sets prices so that $f_m(R_{jt}, R_{mt}|\phi_{t-1}^m)$, its perceived bivariate normal joint distribution of $\tilde{R}_{jt}$ and $\tilde{R}_{mt}$, is constant through time. This means that the market sets prices so that α_j^m, β_j^m, and its perceived distribution on $\tilde{\epsilon}_{jt}$ are the same, period after period. Moreover, it is assumed that it is possible for the market to set prices so that the true joint distribution of $\tilde{R}_{jt}$ and $\tilde{R}_{mt}$, $f(R_{jt}, R_{mt}|\phi_{t-1})$, is constant through time, which means that α_j, β_j and the true distribution of $\tilde{\epsilon}_{jt}$ are the same, period after period.

Suppose now that the market is efficient, so that $f_m(R_{jt}, R_{mt}|\phi_{t-1}^m)$ and $f(R_{jt}, R_{mt}|\phi_{t-1})$ coincide. If the joint distribution of security returns is stationary through time, then the market model can be estimated from time series data on $\tilde{R}_{jt}$ and $\tilde{R}_{mt}$, using the least squares procedures of Chapters 3 and 4. The result is the estimated version of (16),

$$\tilde{R}_{jt} = \tilde{a}_j + \tilde{b}_j \tilde{R}_{mt} + \tilde{e}_{jt},$$

where $\tilde{a}_j$, $\tilde{b}_j$ and $\tilde{e}_{jt}$ are unbiased estimators of $\alpha_j = \alpha_j^m$, $\beta_j = \beta_j^m$, and $\tilde{e}_{jt} = \tilde{e}_{jt}^m$ in (16) and (20). Thus, when the market is efficient and the joint distribution of security returns is constant through time,

$$E(\tilde{e}_{jt}\,|\,\phi_{t-1}, R_{mt}) = E_m(\tilde{e}_{jt}^m\,|\,\phi_{t-1}^m, R_{mt}) = 0.$$

In words, with an efficient market and stationary return distributions, the deviation of $\tilde{e}_{jt}$ from zero results solely from new information that becomes available at t; there is no way to use information available at $t - 1$ as the basis of a correct nonzero assessment of the expected value of $\tilde{e}_{jt}$. For example, if new information about the earnings of firm j is available at $t - 1$, this affects the price of the security set at $t - 1$, which in turn determines $\tilde{e}_{j,t-1}$. But in an efficient market, the earnings information available at $t - 1$ is fully utilized in setting the price of the security at $t - 1$. This means that at t, the deviation of $\tilde{e}_{jt}$ from zero cannot be due to the earnings information that was available at $t - 1$. On the other hand, if the market is inefficient, and in particular if there is some lag in the adjustment of prices to new company-specific information, then the residual for period t is to some extent predictable for information available at $t - 1$; that is, ϕ_{t-1} and ϕ_{t-1}^m no longer coincide, so that

$$E(\tilde{e}_{jt}\,|\,\phi_{t-1}, R_{mt}) \neq 0.$$

Rather than continuing this general and excessively formal discussion of how tests of market efficiency can be approached in the context of the market model, we let the details of the approach arise naturally in the course of a discussion of a specific study, the work on stock splits by Fama, Fisher, Jensen, and Roll (1969), henceforth FFJR, which is the first study that uses the market model as the basis of a test of market efficiency.

SPLITS AND THE ADJUSTMENT OF STOCK PRICES TO NEW INFORMATION

Since the only apparent result of a stock split is to multiply the number of shares per shareholder, without changing any shareholder's claims on the firm's assets, splits in themselves are not necessarily sources of new information. The presumption of FFJR is that splits may be associated with more fundamentally important information. The idea is to examine security returns around split dates to determine whether there is any unusual behavior and, if so, to what extent it can be accounted for by relationships between splits and more fundamental variables.

The FFJR sample includes all 940 stock splits (involving 622 different common stocks) on the NYSE during 1927-1959 where the split was at least 5 new shares for 4 old shares, and where the security was listed for at least 12 months before and after the split. Since any information in a split is likely to be company-specific, the search for unusual behavior in the returns on split securities is confined to market model residuals. Thus, the first step is to obtain estimates of the market model coefficients α_j and β_j of (16) for each of the 622 different securities in the sample. To estimate α_j and β_j, FFJR use all of the monthly return data available for security j during the 1926-1960 period. They then compute the market model residuals for each security for the period from 29 months before to 30 months after any split of the security.

FFJR are concerned with generalizations about the types of return behavior typically associated with splits, rather than with the effects of a split on any individual common stock.* To abstract from the eccentricities of specific cases, they rely on the process of averaging. They concentrate attention on the behavior of cross-sectional averages of estimated regression residuals in the months surrounding split dates. The procedure is as follows: For a given split, define month 0 as the month in which the effective date of a split occurs. Thus, month 0 is not the same chronological date for all securities. Some securities split more than once and hence have more than one month 0. Month 1 is then defined as the month immediately following the split month, month -1 is the month preceding, and so forth. Now define the average residual for month s, with s measured relative to the split month, as

$$\bar{e}_s = \frac{\sum_{j=1}^{N_s} e_{js}}{N_s}$$

where e_{js} is the sample market model residual for security j in month s and N_s is the number of splits for which data are available in month s. The principal tests involve examining the behavior of $\bar{e}_s$ for s in the interval $-29 \leqslant s \leqslant 30$, that is, for the 60 months surrounding the split month. Since FFJR are also interested in the cumulative effects of abnormal return behavior in months surrounding the split month, they also study the behavior of the cumulative average residual U_s, defined as

$$U_s = \sum_{k=-29}^{s} \bar{e}_k.$$

The average residual $\bar{e}_s$ can be interpreted as the average deviation, in month s relative to the split month, of the returns of split stocks from their normal relationships with the market. Similarly, the cumulative average residual U_s can be interpreted as the cumulative deviation from month -29 to month s; it shows the cumulative effects of the wanderings of the returns of split stocks from their normal relationships with the market.

Since the hypothesis about the effects of splits on returns developed by FFJR centers on the dividend behavior of split shares, in some of their tests

*Much of the discussion that follows is taken directly from FFJR.

they separately examine splits that are associated with increased dividends and splits that are associated with decreased dividends. In order to abstract from general changes in dividends across the market, "increased" and "decreased" dividends are measured relative to the average dividends paid by all securities on the New York Stock Exchange during the relevant time periods. The dividends are classified as follows: Define the dividend change ratio as total dividends (per equivalent unsplit share) paid in the 12 months after the split, divided by total dividends paid during the 12 months before the split. Dividend "increases" are then defined as cases where the dividend change ratio of the split stock is greater than the ratio for the NYSE as a whole, while dividend "decreases" include cases of relative dividend decline. FFJR then define $\bar{e}_s^+$, $\bar{e}_s^-$ and U_s^+, U_s^- as the average and cumulative average residuals for splits followed by "increased" ($^+$) and "decreased" ($^-$) dividends.

The most important empirical results of the FFJR study are summarized in Table 5.3 and Figures 5.3a–b and 5.4a–d. Table 5.3 presents the average residuals, cumulative average residuals, and the sample size for each of the two dividend classifications ("increased" and "decreased") and for the total of all splits for each of the 60 months surrounding the split. Figures 5.3a–b.

FIGURE 5.3a
Average Residuals—All Splits

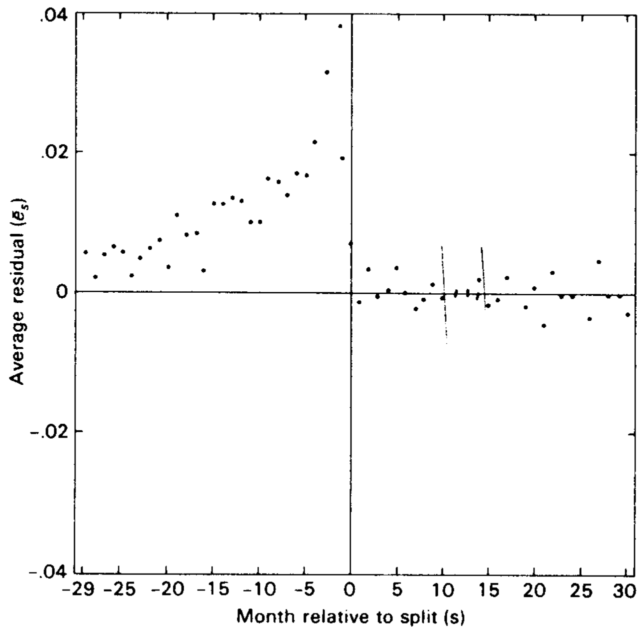

FIGURE 5.3b
Cumulative Average Residuals—All Splits

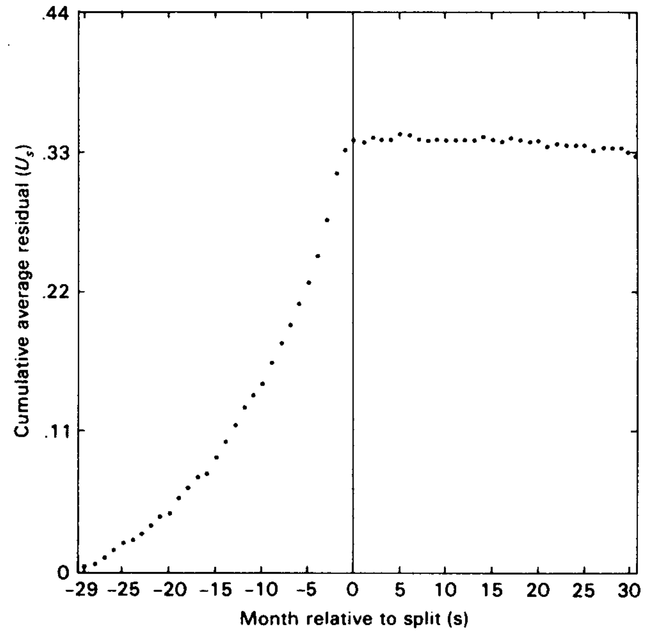

Source: Eugene F. Fama, Lawrence Fisher, Michael Jensen, and Richard Roll, "The Adjustment of Stock Prices to New Information," *International Economic Review* 10 (February 1969): 1–21. Reprinted by permission.

present graphs of the average and cumulative average residuals for the total sample of splits, and Figures 5.4a–d present these graphs for each of the two dividend classifications.

Figures 5.3a, 5.4a and 5.4b show that the average residuals in the 29 months prior to the split are uniformly positive for all splits and for both classes of dividend behavior. This can hardly be attributed entirely to the splitting process. FFJR cite evidence that in only about 10 percent of the splits was the time between the announcement date and the effective date greater than four months. Thus, it seems safe to say that the split cannot account for the behavior of the residuals as far as $2\frac{1}{2}$ years in advance of the split date. Rather, FFJR suggest that there is probably a sharp improvement, relative to the market, in the earnings prospects of a company sometime during the years immediately preceding a split.

Note from Figure 5.3a and Table 5.3 that when all splits are examined together, the largest positive average residuals occur in the three or four months

TABLE 5.3

Analysis of Residuals in Months Surrounding Stock Splits on the NYSE, 1927–1959

	SPLITS FOLLOWED BY DIVIDEND "INCREASES"			SPLITS FOLLOWED BY DIVIDEND "DECREASES"			ALL SPLITS		
(1) MONTH s	(2) AVERAGE $\bar{e}_s^+$	(3) CUMULATIVE U_s^+	(4) SAMPLE SIZE N_s^+	(5) AVERAGE $\bar{e}_s^-$	(6) CUMULATIVE U_s^-	(7) SAMPLE SIZE N_s^-	(8) AVERAGE $\bar{e}_s$	(9) CUMULATIVE U_s	(10) SAMPLE SIZE N_s
−29	0.0062	0.0062	614	0.0033	0.0033	252	0.0054	0.0054	866
−28	0.0013	0.0075	617	0.0030	0.0063	253	0.0018	0.0072	870
−27	0.0068	0.0143	618	0.0007	0.0070	253	0.0050	0.0122	871
−26	0.0054	0.0198	619	0.0085	0.0155	253	0.0063	0.0185	872
−25	0.0042	0.0240	621	0.0089	0.0244	254	0.0056	0.0241	875
−24	0.0020	0.0259	623	0.0026	0.0270	256	0.0021	0.0263	879
−23	0.0055	0.0315	624	0.0028	0.0298	256	0.0047	0.0310	880
−22	0.0073	0.0388	628	0.0028	0.0326	256	0.0060	0.0370	884
−21	0.0049	0.0438	633	0.0131	0.0457	257	0.0073	0.0443	890
−20	0.0044	0.0482	634	0.0005	0.0463	257	0.0033	0.0476	891
−19	0.0110	0.0592	636	0.0102	0.0565	258	0.0108	0.0584	894
−18	0.0076	0.0668	644	0.0089	0.0654	260	0.0080	0.0664	904
−17	0.0072	0.0739	650	0.0111	0.0765	260	0.0083	0.0746	910
−16	0.0035	0.0775	655	0.0009	0.0774	260	0.0028	0.0774	915
−15	0.0135	0.0909	659	0.0101	0.0875	260	0.0125	0.0900	919
−14	0.0135	0.1045	662	0.0100	0.0975	263	0.0125	0.1026	925
−13	0.0148	0.1193	665	0.0099	0.1074	264	0.0134	0.1159	929
−12	0.0138	0.1330	669	0.0107	0.1181	266	0.0129	0.1288	935
−11	0.0098	0.1428	672	0.0103	0.1285	268	0.0099	0.1387	940
−10	0.0103	0.1532	672	0.0082	0.1367	268	0.0097	0.1485	940
− 9	0.0167	0.1698	672	0.0152	0.1520	268	0.0163	0.1647	940
− 8	0.0163	0.1862	672	0.0140	0.1660	268	0.0157	0.1804	940
− 7	0.0159	0.2021	672	0.0083	0.1743	268	0.0138	0.1942	940
− 6	0.0194	0.2215	672	0.0106	0.1849	268	0.0169	0.2111	940
− 5	0.0194	0.2409	672	0.0100	0.1949	268	0.0167	0.2278	940
− 4	0.0260	0.2669	672	0.0104	0.2054	268	0.0216	0.2494	940
− 3	0.0325	0.2993	672	0.0204	0.2258	268	0.0289	0.2783	940
− 2	0.0390	0.3383	672	0.0296	0.2554	268	0.0363	0.3147	940
− 1	0.0199	0.3582	672	0.0176	0.2730	268	0.0192	0.3339	940

	SPLITS FOLLOWED BY DIVIDEND "INCREASES"			SPLITS FOLLOWED BY DIVIDEND "DECREASES"			ALL SPLITS		
(1) MONTH s	(2) AVERAGE $\bar{e}_s^+$	(3) CUMULATIVE U_s^+	(4) SAMPLE SIZE N_s^+	(5) AVERAGE $\bar{e}_s^-$	(6) CUMULATIVE U_s^-	(7) SAMPLE SIZE N_s^-	(8) AVERAGE $\bar{e}_s$	(9) CUMULATIVE U_s	(10) SAMPLE SIZE N_s
0	0.0131	0.3713	672	−0.0090	0.2640	268	0.0068	0.3407	940
1	0.0016	0.3729	672	−0.0088	0.2552	268	−0.0014	0.3393	940
2	0.0052	0.3781	672	−0.0024	0.2528	268	0.0031	0.3424	940
3	0.0024	0.3805	672	−0.0089	0.2439	268	−0.0008	0.3416	940
4	0.0045	0.3851	672	−0.0114	0.2325	268	0.0000	0.3416	940
5	0.0048	0.3898	672	−0.0003	0.2322	268	0.0033	0.3449	940
6	0.0012	0.3911	672	−0.0038	0.2285	268	−0.0002	0.3447	940
7	0.0008	0.3919	672	−0.0106	0.2179	268	−0.0024	0.3423	940
8	−0.0007	0.3912	672	−0.0024	0.2155	268	−0.0012	0.3411	940
9	0.0039	0.3951	672	−0.0065	0.2089	268	0.0009	0.3420	940
10	−0.0001	0.3950	672	−0.0027	0.2062	268	−0.0008	0.3412	940
11	0.0027	0.3977	672	−0.0056	0.2006	268	0.0003	0.3415	940
12	0.0018	0.3996	672	−0.0043	0.1963	268	0.0001	0.3416	940
13	−0.0003	0.3993	666	0.0014	0.1977	264	0.0002	0.3418	930
14	0.0006	0.3999	653	0.0044	0.2021	258	0.0017	0.3435	911
15	−0.0037	0.3962	645	0.0026	0.2047	258	−0.0019	0.3416	903
16	0.0001	0.3963	635	−0.0040	0.2007	257	−0.0011	0.3405	892
17	0.0034	0.3997	633	−0.0011	0.1996	256	0.0021	0.3426	889
18	−0.0015	0.3982	628	0.0025	0.2021	255	−0.0003	0.3423	883
19	−0.0006	0.3976	620	−0.0057	0.1964	251	−0.0021	0.3402	871
20	−0.0002	0.3974	604	0.0027	0.1991	246	0.0006	0.3409	850
21	−0.0037	0.3937	595	−0.0073	0.1918	245	−0.0047	0.3361	840
22	0.0047	0.3984	593	−0.0018	0.1899	244	0.0028	0.3389	837
23	−0.0026	0.3958	593	0.0043	0.1943	242	−0.0006	0.3383	835
24	−0.0022	0.3936	587	0.0031	0.1974	238	−0.0007	0.3376	825
25	0.0012	0.3948	583	−0.0037	0.1936	237	−0.0002	0.3374	820
26	−0.0058	0.3890	582	0.0015	0.1952	236	−0.0037	0.3337	818
27	−0.0003	0.3887	582	0.0082	0.2033	235	0.0021	0.3359	817
28	0.0004	0.3891	580	−0.0023	0.2010	236	−0.0004	0.3355	816
29	0.0012	0.3903	580	−0.0039	0.1971	235	−0.0003	0.3352	815
30	−0.0033	0.3870	579	−0.0025	0.1946	235	−0.0031	0.3321	814

SOURCE: Eugene F. Fama, Lawrence Fisher, Michael Jensen, and Richard Roll, "The Adjustment of Stock Market Prices to New Information," *International Economic Review* 10 (February 1969): 10–11. Reprinted by permission.

FIGURE 5.4a

Average Residuals for Dividend "Increases"

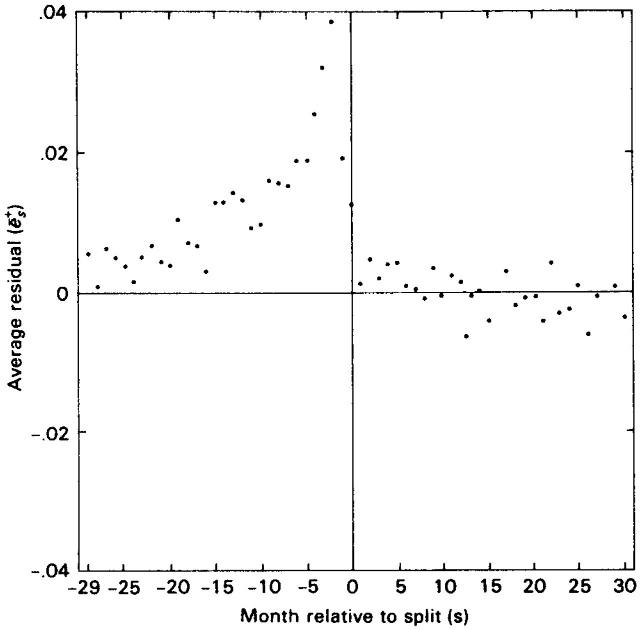

FIGURE 5.4b

Average Residuals for Dividend "Decreases"

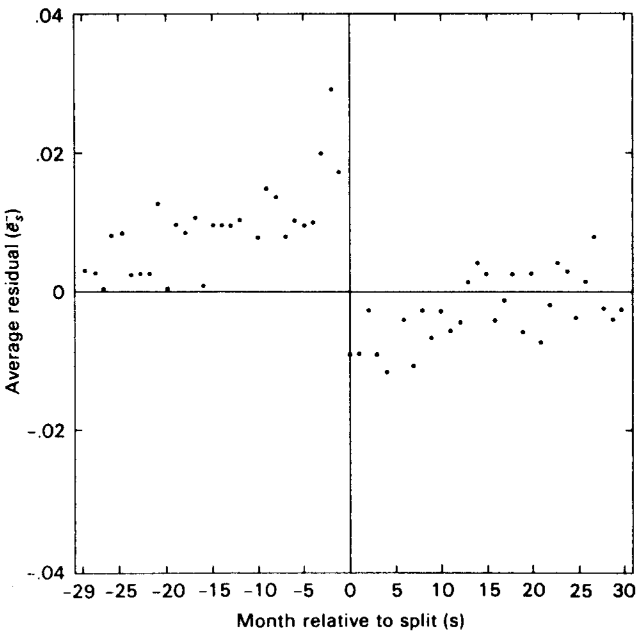

Source: Figures 5.4a–5.4d from Eugene F. Fama, Lawrence Fisher, Michael Jensen, and Richard Roll, "The Adjustment of Stock Prices to New Information," ***International Economic Review*** 10(February 1969): 1-21. Reprinted by permission.

FIGURE 5.4c

Cumulative Average Residuals for Dividend "Increases"

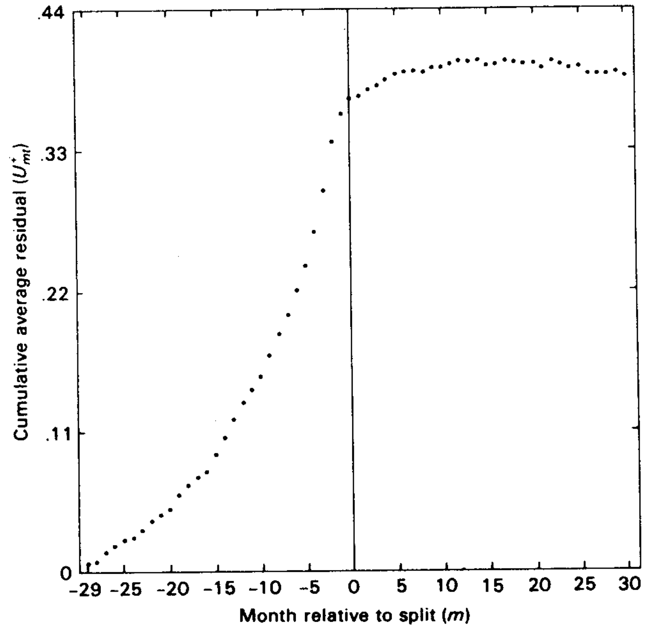

FIGURE 5.4d

Cumulative Average Residuals for Dividend "Decreases"

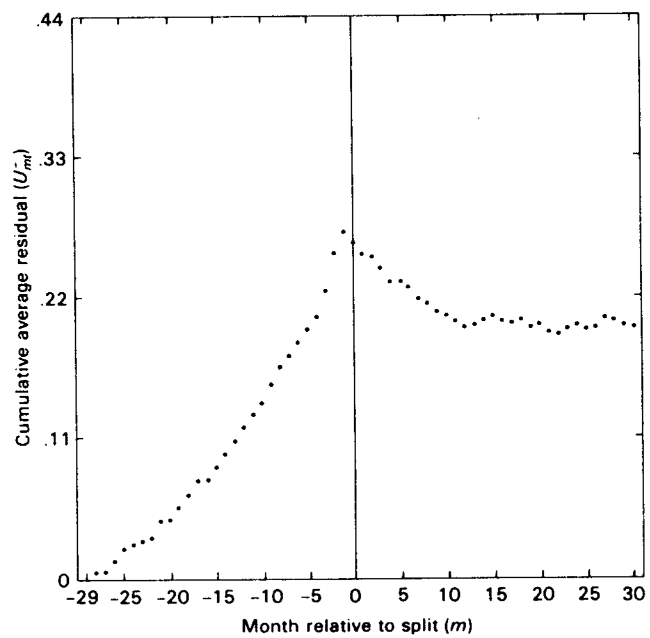

immediately preceding the split, but that after the split the average residuals are randomly distributed about 0. Equivalently, in Figure 5.3b the *cumulative* average residuals rise up to the split month, but there is almost no further systematic movement thereafter. During the first year after the split, the cumulative average residual changes by less than one-tenth of one percentage point, and the total change in the cumulative average residual during the $2\frac{1}{2}$ years following the split is less than one percentage point. This is especially striking because 71.5 percent (672 out of 940) of all splits experience greater percentage dividend increases in the year after the split than the average for all securities on the NYSE.

The explanation offered by FFJR for this behavior of the average residuals is as follows. When a split is announced or anticipated, the market interprets this, and correctly so, as greatly improving the probability that dividends will soon be substantially increased. In many cases the split and the dividend increase are announced at the same time. If, as Lintner (1956) suggests, firms are reluctant to reduce dividends, then a split, which implies an increased expected dividend, is a signal to the market that the company's directors are confident that future earnings will be sufficient to maintain dividend payments at a higher level. If the market agrees with the judgments of the directors, then it is possible that the large price increases in the months immediately preceding a split are due to altered expectations concerning the future earning potential of the firm and thus of its shares, rather than to any intrinsic effects of the split itself.

If the information effects of actual or anticipated dividend increases explain the behavior of common stock returns in the months immediately surrounding a split, then return behavior subsequent to the split should be substantially different in cases where the dividend increase materializes than in cases where it does not. It is apparent from Figures 5.4a–d that the differences are in fact substantial, and FFJR argue that they are in the direction predicted by their hypothesis.

Thus, the fact that the cumulative average residuals for both dividend classes rise sharply in the few months before the split is consistent with the hypothesis that the market recognizes that splits are usually associated with higher dividend payments. In some cases, however, the dividend increase, if it occurs, is declared sometime during the year after the split. Thus, it is not surprising that the average residuals (Figure 5.4a) for stocks in the "increased" dividend class are in general slightly positive in the year after the split, so that the cumulative average residuals for these stocks (Figure 5.4c) drift upward. The fact that this upward drift is only slight can be explained in two ways. First, in many cases the dividend increase associated with a split is declared and the corresponding price adjustments take place before the end of the split

month. Second, according to the FFJR hypothesis, when the split is declared, even if no dividend announcement is made, there is some price adjustment in anticipation of future dividend increases. Thus, only a slight additional adjustment is necessary when the dividend increase actually takes place. By one year after the split, the returns on stocks which have experienced dividend "increases" have resumed their normal relationships to market returns, since from this point onward the average residuals are small and randomly scattered about zero.

FFJR contend that the behavior of the residuals for stock splits associated with "decreased" dividends provides the strongest evidence for their split hypothesis. For stocks in the "decreased" dividend class the average and cumulative average residuals (Figures 5.4b and 5.4d) rise in the few months before the split but then plummet in the few months following the split, when the anticipated dividend increase is not forthcoming. These split stocks with poor dividend performance on the average perform poorly in each of the 12 months following the split, but their period of poorest performance is in the few months immediately after the split, when the improved dividend, if it were coming at all, would most likely be declared. The hypothesis is further reinforced by the observation that when a year has passed after the split, the cumulative average residual has fallen to about where it was five months prior to the split, which is probably about the earliest time reliable information concerning a possible split is likely to reach the market. Thus, by the time it becomes clear that the anticipated dividend increase is not forthcoming, the apparent effects of the split seem to be completely wiped away, and the stock's returns revert to their normal relationship with market returns. In sum, FFJR suggest that once the information effects of associated dividend changes are properly considered, a split per se has no net effect on common stock returns.

Finally, and most important, although the behavior of post-split returns is very different depending on whether or not dividend "increases" occur, and despite the fact that a substantial majority of split securities do experience dividend "increases," when all splits are examined together (Figures 5.3a–b), the average residuals are randomly distributed about 0 during the year after the split, so that there is no net movement either up or down in the cumulative average residuals. Thus, the market apparently makes unbiased forecasts of the implications of a split for future dividends, and these forecasts are fully reflected in the price of the security by the end of the split month. After considerably more data analysis than we can summarize here, FFJR conclude that their results are consistent with the hypothesis that the stock market is efficient, at least with respect to its ability to adjust to the information implicit in a split.

One point from the remainder of the FFJR analysis should be mentioned. FFJR especially emphasize that the persistent upward drift of the cumulative average residuals in the months preceding the split is not a phenomenon that could be used to increase expected trading profits. The reason is that the behavior of the average residuals is not representative of the behavior of the residuals for individual securities. In months prior to the split, successive sample residuals for individual securities seem to be independent. But in most cases, there are a few months in which the residuals are abnormally large and positive. The months of large residuals differ from security to security, however, and the differences in timing explain why the signs of the average residuals are uniformly positive for many months preceding the split.

Since one purpose of this book is to encourage the reader to develop a critical eye for discussions of empirical work, some comments about the FFJR analysis are relevant. First, FFJR are somewhat "aggressive" in interpreting their empirical results. In their view, the unusual behavior of the returns on a splitting security in the months immediately preceding a split reflects the information content of the dividend change that usually accompanies a split. There is, however, no direct evidence in their data that dividends or splits convey real information to the market about the future prospects of a firm. For example, an alternative view, completely consistent with their empirical results, is that dividends are a passive variable in the whole process. That is, companies tend to increase dividends when earnings increase and to decrease dividends when earnings decrease. In this view, the FFJR data suggest that splits tend to occur when firms have experienced unusual increases in earnings, which accounts for the positive average residuals of splitting shares in the months preceding the split. As chance will have it, however, the good times do not persist for all firms. Some of them experience earnings declines in the year after the split, which in the FFJR data show up as decreased dividends. Thus, the behavior of dividends is merely a proxy for the behavior of earnings, and neither dividend changes nor splits are a source of information.

It is still the case, however, that in this alternative view the FFJR evidence is consistent with the hypothesis that the market is efficient. Thus, about 30 percent of the firms will come on relatively bad times (decreased earnings) subsequent to splitting their shares, and this will be reflected in decreased dividends. If the market is efficient when adjusting security prices to the high earnings for the period preceding the split, it will take full account of the chances of good and bad times in the period following the split, so that splitting shares will not, on average, experience unusually high or low returns in the period following the split. In Figures 5.3a–b the behavior of the average residuals in the years after the split is consistent with this implication of market efficiency.

OTHER STUDIES OF PUBLIC ANNOUNCEMENTS

Variants of the method of residual analysis developed by FFJR have been used by others to study the effects of different kinds of public announcements, and all of these studies are in most respects consistent with the hypothesis that the market is efficient.

For example, using data on 261 major firms for the period 1946-1966, Ball and Brown (1968) apply the method to study the effects of annual earnings announcements. They use the residuals from a time series regression of the annual earnings of a firm on the average earnings of all their firms to classify the firm's earnings for a given year as having "increased" or "decreased" relative to the market. Residuals from estimates of the market model obtained from monthly data are then used to compute cumulative average return residuals separately for those earnings that "increased" and those that "decreased." The cumulative average return residuals rise throughout the year in advance of the announcement for the "increased" earnings category, and fall for the "decreased" earnings category. Ball and Brown conclude that no more than about 10–15 percent of the information in the annual earnings announcement has not been anticipated by the month of the announcement.

Further evidence consistent with the hypothesis that the market is efficient is provided in the work of Scholes (1972) on large secondary offerings of common stock, that is, large underwritten sales of existing common stocks by individuals and institutions. He finds that, on average, large secondary issues are associated with a decline of between 1 and 2 percent in the cumulative average residual returns for the corresponding common stocks. Since the magnitude of the price adjustment is unrelated to the size of the issue, Scholes concludes that the adjustment is not due to "selling pressure," as is commonly believed, but rather results from negative information implicit in the fact that somebody is trying to sell a large block of a firm's stock. Moreover, he presents evidence that the value of the information in a secondary offering depends to some extent on the vendor. As might be expected, by far the largest negative cumulative average residuals occur where the vendor is the corporation itself or one of its officers, with investment companies a distant second. The identity of the vendor is not generally known at the time of a secondary offering, however, and corporate insiders need only report their transactions in their company's stock to the Securities and Exchange Commission within six days after a sale. By this time, the market on average has fully adjusted to the information in the secondary, as indicated by the fact that the average residuals behave randomly thereafter.

To avoid giving a falsely monolithic appearance to the evidence consistent with the hypothesis that the market is efficient, we should note that although Scholes's work indicates that prices adjust efficiently to the public information in a secondary, his work is also evidence that corporate insiders at least sometimes have important information about their firms that is not yet reflected in prices. This is evidence against market efficiency, since it says that prices do not fully reflect *all* available information.* Moreover, other evidence of the same sort is offered by Neiderhoffer and Osborne (1966), who point out that specialists on the NYSE apparently use their monopolistic access to information concerning unfilled limit orders (orders to buy and sell at given prices) to generate monopoly profits.

Like any null hypothesis, however, the hypothesis that the market is efficient is not likely to be a completely accurate view of the world. We might look at the various tests as providing the evidence that helps us to judge the extent to which the market is efficient and the extent to which it is inefficient. The evidence discussed so far is consistent with market efficiency in the sense that prices fully reflect publicly available information, such as past prices, splits, earnings announcements, etc., but there is also evidence that the market is not completely efficient, since corporate insiders and NYSE specialists apparently have access to information that is not fully reflected in prices. In practical terms, the evidence suggests that if an investor or investment counselor only has access to publicly available information, then the hypothesis that the market is efficient is an appropriate approximation to the world. If prices fully reflect publicly available information, then such information cannot be used to beat the market. On the other hand, market efficiency is an inappropriate view of the world for a corporate insider or an NYSE specialist, since they sometimes have access to and can trade on information that is not fully reflected in prices.

D. *Returns Conform to a Risk-Return Relationship*

The most recent tests of market efficiency make use of a model of market equilibrium in which the market sets prices at any time $t-1$ so that there is a positive relationship between the expected return on a security from time $t-1$ to time t and the risk of the security. For example, one such study, by Mandelker (1974), is concerned with the adjustment of prices to the announcement that two firms will merge. Another, by Jaffe (1974), is concerned with the adjustment of prices to any information implicit in insider trading.

We cannot do justice to tests of market efficiency based on risk-return models of market equilibrium until we consider these models in some detail. This is the topic of Chapters 7-9. Tests of market efficiency that are based on these risk-return models are discussed in Chapter 9.

*Evidence that insiders have monopolistic access to information about their firm is also to be found in the work of Lorie and Neiderhoffer (1968) and Jaffe (1974). Jaffe's work is discussed in Chapter 9.

IV. Conclusions and Some Fine Points of the Theory

In the model of price formation presented in this chapter, at any time $t-1$ the "market" assesses a joint distribution for security prices at time $t, f_m(p_{1t}, \ldots, p_{nt} | \phi^m_{t-1})$. The characteristics of this distribution, along with some propositions about the nature of market equilibrium (for example, equilibrium expected returns are positive), are then the basis of the equilibrium prices of securities, $p_{1,t-1}, \ldots, p_{n,t-1}$, set at $t-1$. This is clearly a simplified view of the world, and we now discuss some of the ways in which it is not completely realistic.

First, in the description of the process of price formation given above, the "market" assesses probability distributions and the "market" sets prices. This can only be a completely accurate view of the world if all the individual participants in the market (a) have the same information and (b) agree on its implications for the joint distribution of future prices. Neither of these conditions is completely descriptive. Nor is it completely realistic to presume that when market prices are determined, they result from a conscious assessment of the joint distribution of security prices by all or most or even many investors.

Pushing this line of attack even further, the two-step process of price formation assumed in this chapter masks some even stronger assumptions about the analytical capabilities of investors. Thus, prices set at $t-1$ result from an assessment of the joint distribution of prices for time t. But the world is not presumed to end at time t, so the prices that turn up at t must themselves be the consequences of a market equilibrium. That is, pushing the two-step process of price formation one period ahead, prices at time t will be set on the basis of characteristics of the joint distribution assessed at t on prices for $t+1$. And the process will be repeated at each future point in time. Thus, when at time $t-1$ the market assesses a joint distribution on prices for t, it must assess what the state of the world at $t-1$ implies about the likelihoods of different states at t, and it must assume something about how it will respond to different states in setting security prices at t. To do this, it must in turn make assessments about the likelihoods of different states of the world at $t+1$ and how it will respond to them in setting prices and so forth. In short, the discussion of a two-step process of price formation in the simple model glosses over the fact that the first step, assessment of the joint distribution of prices for time t, also implies assessments of the joint distributions of prices at each future point in time, with all of the judgments about future