CHAPTER ELEVEN

The Efficient Market Hypothesis

ONE OF THE early applications of computers in economics in the 1950s was to analyze economic time series. Business cycle theorists felt that tracing the evolution of several economic variables over time would clarify and predict the progress of the economy through boom and bust periods. A natural candidate for analysis was the behavior of stock market prices over time. Assuming that stock prices reflect the prospects of the firm, recurrent patterns of peaks and troughs in economic performance ought to show up in those prices.

Maurice Kendall examined this proposition in 1953.[1] He found to his great surprise that he could identify no predictable patterns in stock prices. Prices seemed to evolve randomly. They were as likely to go up as they were to go down on any particular day, regardless of past performance. The data provided no way to predict price movements.

At first blush, Kendall's results were disturbing to some financial economists. They seemed to imply that the stock market is dominated by erratic market psychology, or "animal spirits"—that it follows no logical rules. In short, the results appeared to confirm the irrationality of the market. On further reflection, however, economists came to reverse their interpretation of Kendall's study.

It soon became apparent that random price movements indicated a well-functioning or efficient market, not an irrational one. In this chapter we explore the reasoning behind what may seem a surprising conclusion. We show how competition among analysts leads naturally to market efficiency, and we examine the implications of the efficient market hypothesis for investment policy. We also consider empirical evidence that supports and contradicts the notion of market efficiency.

[1]Maurice Kendall, "The Analysis of Economic Time Series, Part I: Prices," *Journal of the Royal Statistical Society* 96 (1953).

11.1 Random Walks and the Efficient Market Hypothesis

Suppose Kendall had discovered that stock price changes are predictable. What a gold mine this would have been. If they could use Kendall's equations to predict stock prices, investors would reap unending profits simply by purchasing stocks that the computer model implied were about to increase in price and by selling those stocks about to fall in price.

A moment's reflection should be enough to convince yourself that this situation could not persist for long. For example, suppose that the model predicts with great confidence that XYZ stock price, currently at $100 per share, will rise dramatically in 3 days to $110. What would all investors with access to the model's prediction do today? Obviously, they would place a great wave of immediate buy orders to cash in on the prospective increase in stock price. No one holding XYZ, however, would be willing to sell. The net effect would be an *immediate* jump in the stock price to $110. The forecast of a future price increase will lead instead to an immediate price increase. In other words, the stock price will immediately reflect the "good news" implicit in the model's forecast.

This simple example illustrates why Kendall's attempt to find recurrent patterns in stock price movements was likely to fail. A forecast about favorable *future* performance leads instead to favorable *current* performance, as market participants all try to get in on the action before the price jump.

More generally, one might say that any information that could be used to predict stock performance should already be reflected in stock prices. As soon as there is any information indicating that a stock is underpriced and therefore offers a profit opportunity, investors flock to buy the stock and immediately bid up its price to a fair level, where only ordinary rates of return can be expected. These "ordinary rates" are simply rates of return commensurate with the risk of the stock.

However, if prices are bid immediately to fair levels, given all available information, it must be that they increase or decrease only in response to new information. New information, by definition, must be unpredictable; if it could be predicted, then the prediction would be part of today's information. Thus stock prices that change in response to new (that is, previously unpredicted) information also must move unpredictably.

This is the essence of the argument that stock prices should follow a **random walk,** that is, that price changes should be random and unpredictable.[2] Far from a proof of market irrationality, randomly evolving stock prices would be the necessary consequence of intelligent investors competing to discover relevant information on which to buy or sell stocks before the rest of the market becomes aware of that information.

Don't confuse randomness in price *changes* with irrationality in the *level* of prices. If prices are determined rationally, then only new information will cause them to change. Therefore, a random walk would be the natural result of prices that always reflect all current knowledge. Indeed, if stock price movements were predictable, that would be damning evidence of stock market inefficiency, because the ability to predict prices would indicate that

[2]Actually, we are being a little loose with terminology here. Strictly speaking, we should characterize stock prices as following a submartingale, meaning that the expected change in the price can be positive, presumably as compensation for the time value of money and systematic risk. Moreover, the expected return may change over time as risk factors change. A random walk is more restrictive in that it constrains successive stock returns to be independent *and* identically distributed. Nevertheless, the term "random walk" is commonly used in the looser sense that price changes are essentially unpredictable. We will follow this convention.

all available information was not already reflected in stock prices. Therefore, the notion that stocks already reflect all available information is referred to as the **efficient market hypothesis** (EMH).[3]

Figure 11.1 illustrates the response of stock prices to new information in an efficient market. The graph plots the price response of a sample of firms that were targets of takeover attempts. In most takeovers, the acquiring firm pays a substantial premium over current market prices. Therefore, announcement of a takeover attempt should cause the stock price to jump. The figure shows that stock prices jump dramatically on the day the news becomes public. However, there is no further drift in prices *after* the announcement date, suggesting that prices reflect the new information, including the likely magnitude of the takeover premium, by the end of the trading day.

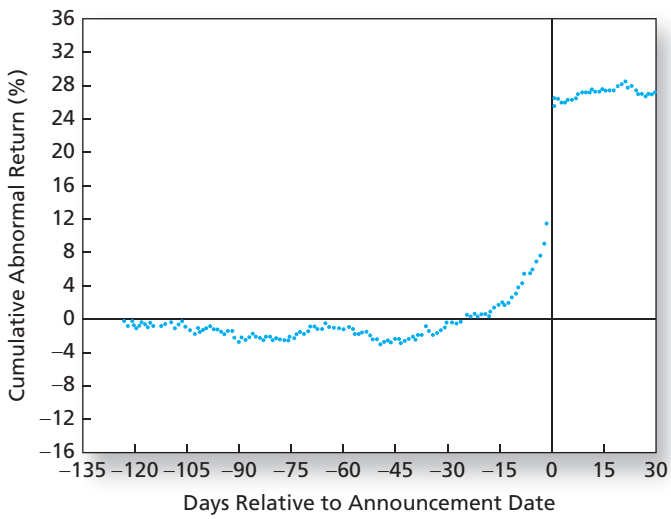

Figure 11.1 Cumulative abnormal returns before takeover attempts: target companies

Source: Arthur Keown and John Pinkerton, "Merger Announcements and Insider Trading Activity," *Journal of Finance* 36 (September 1981). Used with permission of John Wiley and Sons, via Copyright Clearance Center. Updates courtesy of Jinghua Yan.

Even more dramatic evidence of rapid response to new information may be found in intraday prices. For example, Patell and Wolfson show that most of the stock price response to corporate dividend or earnings announcements occurs within 10 minutes of the announcement.[4] A nice illustration of such rapid adjustment is provided in a study by Busse and Green, who track minute-by-minute stock prices of firms that are featured on CNBC's "Morning" or "Midday Call" segments.[5] Minute 0 in Figure 11.2 is the time at which the stock is mentioned on the midday show. The top line is the average price movement of stocks that receive positive reports, while the bottom line reports returns on stocks with negative reports. Notice that the top line levels off, indicating that the market has fully digested the news within 5 minutes of the report. The bottom line levels off within about 12 minutes.

Competition as the Source of Efficiency

Why should we expect stock prices to reflect "all available information"? After all, if you are willing to spend time and money on gathering information, it might seem reasonable that you could turn up something that has been overlooked by the rest of the investment community. When information is costly to uncover and analyze, one would expect investment analysis calling for such expenditures to result in an increased expected return.

[3]Market efficiency should not be confused with the idea of efficient portfolios introduced in Chapter 7. An informationally efficient *market* is one in which information is rapidly disseminated and reflected in prices. An efficient *portfolio* is one with the highest expected return for a given level of risk.

[4]J. M. Patell and M. A. Wolfson, "The Intraday Speed of Adjustment of Stock Prices to Earnings and Dividend Announcements," *Journal of Financial Economics* 13 (June 1984), pp. 223–52.

[5]J. A. Busse and T. C. Green, "Market Efficiency in Real Time," *Journal of Financial Economics* 65 (2002), pp. 415–37. You can find an intraday movie version of this figure at **www.bus.emory.edu/cgreen/docs/cnbc/cnbc.html**.

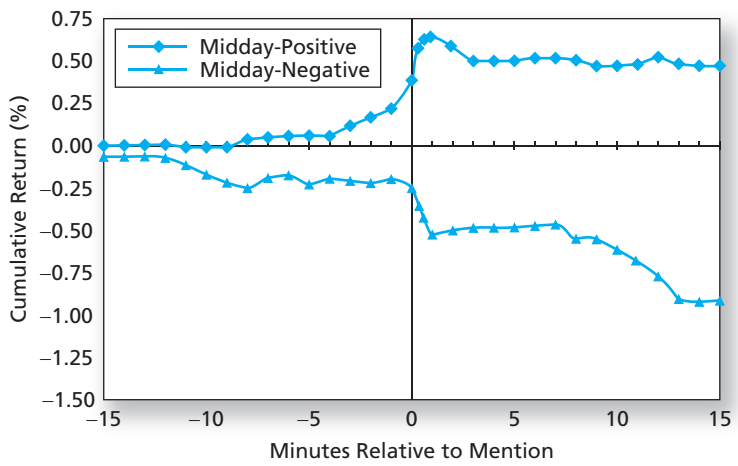

Figure 11.2 Stock price reaction to CNBC reports. The figure shows the reaction of stock prices to on-air stock reports during the "Midday Call" segment on CNBC. The chart plots cumulative returns beginning 15 minutes before the stock report.

Source: Reprinted from J. A. Busse and T. C. Green, "Market Efficiency in Real Time," *Journal of Financial Economics* 65 (2002), p. 422. Copyright 2002, with permission from Elsevier.

This point has been stressed by Grossman and Stiglitz.[6] They argued that investors will have an incentive to spend time and resources to analyze and uncover new information only if such activity is likely to generate higher investment returns. Thus, in market equilibrium, efficient information-gathering activity should be fruitful. Moreover, it would not be surprising to find that the degree of efficiency differs across various markets. For example, emerging markets that are less intensively analyzed than U.S. markets or in which accounting disclosure requirements are less rigorous may be less efficient than U.S. markets. Small stocks that receive relatively little coverage by Wall Street analysts may be less efficiently priced than large ones. Still, while we would not go so far as to say that you absolutely cannot come up with new information, it makes sense to consider and respect your competition.

Example 11.1 Rewards for Incremental Performance

Consider an investment management fund currently managing a $5 billion portfolio. Suppose that the fund manager can devise a research program that could increase the portfolio rate of return by one-tenth of 1% per year, a seemingly modest amount. This program would increase the dollar return to the portfolio by $5 billion × .001, or $5 million. Therefore, the fund would be willing to spend up to $5 million per year on research to increase stock returns by a mere tenth of 1% per year. With such large rewards for such small increases in investment performance, it should not be surprising that professional portfolio managers are willing to spend large sums on industry analysts, computer support, and research effort, and therefore that price changes are, generally speaking, difficult to predict.

With so many well-backed analysts willing to spend considerable resources on research, easy pickings in the market are rare. Moreover, the incremental rates of return on research activity may be so small that only managers of the largest portfolios will find them worth pursuing.

Although it may not literally be true that "all" relevant information will be uncovered, it is virtually certain that there are many investigators hot on the trail of most leads that seem likely to improve investment performance. Competition among these many well-backed,

[6]Sanford J. Grossman and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review* 70 (June 1980).

Matchmakers for the Information Age

The most precious commodity on Wall Street is information, and informed players can charge handsomely for providing it. An industry of so-called *expert network providers* has emerged for selling access to experts with unique insights about a wide variety of firms and industries to investors who need that information to make decisions. These firms have been dubbed matchmakers for the information age. Experts can range from doctors who help predict the release of blockbuster drugs to meteorologists who forecast weather that can affect commodity prices to business executives who can provide specialized insight about companies and industries.

But some of those experts have peddled prohibited inside information. In 2011, Winifred Jiau, a consultant for Primary Global Research, was convicted of selling information about Nvidia and Marvell Technologies to the hedge fund SAC Capital Advisors. Several other employees of Primary Global also were charged with insider trading.

Expert firms are supposed to provide only public information, along with the expert's insights and perspective. But the temptation to hire experts with inside information and charge handsomely for access to them is obvious. The SEC has raised concerns about the boundary between legitimate and illegal services, and several hedge funds in 2011 shut down after raids searched for evidence of such illicit activity.

In the wake of increased scrutiny, compliance efforts of both buyers and sellers of expert information have mushroomed. The largest network firm is Gerson Lehrman Group with a stable of 300,000 experts. It now maintains records down to the minute of which of its experts talks to whom and the topics they have discussed.[7] These records could be turned over to authorities in the event of an insider trading investigation. For their part, some hedge funds have simply ceased working with expert-network firms or have promulgated clearer rules for when their employees may talk with consultants.

Even with these safeguards, however, there remains room for trouble. For example, an investor may meet an expert through a legitimate network and then the two may establish a consulting relationship on their own. This legal matchmaking becomes the precursor to the illegal selling of insider tips. Where there is a will to cheat, there usually will be a way.

WORDS FROM THE STREET

highly paid, aggressive analysts ensures that, as a general rule, stock prices ought to reflect available information regarding their proper levels.

Information is often said to be the most precious commodity on Wall Street, and the competition for it is intense. Sometimes the quest for a competitive advantage can tip over into a search for illegal inside information. In 2011, Raj Rajaratnam, the head of the Galleon Group hedge fund which once managed $6.5 billion, was convicted on insider trading charges for soliciting tips from a network of corporate insiders and traders. Rajaratnam's was only one of several major insider trading cases working their way through the courts in 2011. While Galleon's practices were egregious, drawing a clear line separating legitimate and prohibited sources of information often can be difficult. For example, a large industry of *expert network* firms has emerged in the last decade to connect (for a fee) investors to industry experts who can provide unique perspective on a company. As the nearby box discusses, this sort of arrangement can easily cross the line into insider trading.

Versions of the Efficient Market Hypothesis

It is common to distinguish among three versions of the EMH: the weak, semistrong, and strong forms of the hypothesis. These versions differ by their notions of what is meant by the term "all available information."

The **weak-form** hypothesis asserts that stock prices already reflect all information that can be derived by examining market trading data such as the history of past prices, trading volume, or short interest. This version of the hypothesis implies that trend analysis is fruitless. Past stock price data are publicly available and virtually costless to obtain. The weak-form hypothesis holds that if such data ever conveyed reliable signals about future performance, all investors already would have learned to exploit the signals. Ultimately,

[7]"Expert Networks Are the Matchmakers for the Information Age," *The Economist,* June 16, 2011.

the signals lose their value as they become widely known because a buy signal, for instance, would result in an immediate price increase.

The **semistrong-form** hypothesis states that all publicly available information regarding the prospects of a firm must be reflected already in the stock price. Such information includes, in addition to past prices, fundamental data on the firm's product line, quality of management, balance sheet composition, patents held, earning forecasts, and accounting practices. Again, if investors have access to such information from publicly available sources, one would expect it to be reflected in stock prices.

Finally, the **strong-form** version of the efficient market hypothesis states that stock prices reflect all information relevant to the firm, even including information available only to company insiders. This version of the hypothesis is quite extreme. Few would argue with the proposition that corporate officers have access to pertinent information long enough before public release to enable them to profit from trading on that information. Indeed, much of the activity of the Securities and Exchange Commission is directed toward preventing insiders from profiting by exploiting their privileged situation. Rule 10b-5 of the Security Exchange Act of 1934 sets limits on trading by corporate officers, directors, and substantial owners, requiring them to report trades to the SEC. These insiders, their relatives, and any associates who trade on information supplied by insiders are considered in violation of the law.

Defining insider trading is not always easy, however. After all, stock analysts are in the business of uncovering information not already widely known to market participants. As we saw in Chapter 3 as well as in the nearby box, the distinction between private and inside information is sometimes murky.

Notice one thing that all versions of the EMH have in common: They all assert that prices should reflect *available* information. We do not expect traders to be superhuman or market prices to always be right. We will always wish for more information about a company's prospects than will be available. Sometimes market prices will turn out in retrospect to have been outrageously high, at other times absurdly low. The EMH asserts only that at the given time, using current information, we cannot be sure if today's prices will ultimately prove themselves to have been too high or too low. If markets are rational, however, we can expect them to be correct on average.

CONCEPT CHECK 11.1

a. Suppose you observed that high-level managers make superior returns on investments in their company's stock. Would this be a violation of weak-form market efficiency? Would it be a violation of strong-form market efficiency?

b. If the weak form of the efficient market hypothesis is valid, must the strong form also hold? Conversely, does strong-form efficiency imply weak-form efficiency?

11.2 Implications of the EMH

Technical Analysis

Technical analysis is essentially the search for recurrent and predictable patterns in stock prices. Although technicians recognize the value of information regarding future economic prospects of the firm, they believe that such information is not necessary for a successful trading strategy. This is because whatever the fundamental reason for a change in stock price, if the stock price responds slowly enough, the analyst will be able to identify a trend

that can be exploited during the adjustment period. The key to successful technical analysis is a sluggish response of stock prices to fundamental supply-and-demand factors. This prerequisite, of course, is diametrically opposed to the notion of an efficient market.

Technical analysts are sometimes called *chartists* because they study records or charts of past stock prices, hoping to find patterns they can exploit to make a profit. As an example of technical analysis, consider the *relative strength* approach. The chartist compares stock performance over a recent period to performance of the market or other stocks in the same industry. A simple version of relative strength takes the ratio of the stock price to a market indicator such as the S&P 500 index. If the ratio increases over time, the stock is said to exhibit relative strength because its price performance is better than that of the broad market. Such strength presumably may continue for a long enough period of time to offer profit opportunities.

One of the most commonly heard components of technical analysis is the notion of **resistance levels** or **support levels.** These values are said to be price levels above which it is difficult for stock prices to rise, or below which it is unlikely for them to fall, and they are believed to be levels determined by market psychology.

Example 11.2 Resistance Levels

Consider stock XYZ, which traded for several months at a price of $72 and then declined to $65. If the stock eventually begins to increase in price, $72 is considered a resistance level (according to this theory) because investors who bought originally at $72 will be eager to sell their shares as soon as they can break even on their investment. Therefore, at prices near $72 a wave of selling pressure would exist. Such activity imparts a type of "memory" to the market that allows past price history to influence current stock prospects.

The efficient market hypothesis implies that technical analysis is without merit. The past history of prices and trading volume is publicly available at minimal cost. Therefore, any information that was ever available from analyzing past prices has already been reflected in stock prices. As investors compete to exploit their common knowledge of a stock's price history, they necessarily drive stock prices to levels where expected rates of return are exactly commensurate with risk. At those levels one cannot expect abnormal returns.

As an example of how this process works, consider what would happen if the market believed that a level of $72 truly was a resistance level for stock XYZ in Example 11.2. No one would be willing to purchase the stock at a price of $71.50, because it would have almost no room to increase in price, but ample room to fall. However, if no one would buy it at $71.50, then $71.50 would become a resistance level. But then, using a similar analysis, no one would buy it at $71, or $70, and so on. The notion of a resistance level is a logical conundrum. Its simple resolution is the recognition that if the stock is ever to sell at $71.50, investors *must* believe that the price can as easily increase as fall. The fact that investors are willing to purchase (or even hold) the stock at $71.50 is evidence of their belief that they can earn a fair expected rate of return at that price.

CONCEPT CHECK 11.2

If everyone in the market believes in resistance levels, why do these beliefs not become self-fulfilling prophecies?

An interesting question is whether a technical rule that seems to work will continue to work in the future once it becomes widely recognized. A clever analyst may occasionally uncover a profitable trading rule, but the real test of efficient markets is whether the rule itself becomes reflected in stock prices once its value is discovered. Once a useful

technical rule (or price pattern) is discovered, it ought to be invalidated when the mass of traders attempts to exploit it. In this sense, price patterns ought to be *self-destructing.*

Thus the market dynamic is one of a continual search for profitable trading rules, followed by destruction by overuse of those rules found to be successful, followed by more searching for yet-undiscovered rules.

Fundamental Analysis

Fundamental analysis uses earnings and dividend prospects of the firm, expectations of future interest rates, and risk evaluation of the firm to determine proper stock prices. Ultimately, it represents an attempt to determine the present discounted value of all the payments a stockholder will receive from each share of stock. If that value exceeds the stock price, the fundamental analyst would recommend purchasing the stock.

Fundamental analysts usually start with a study of past earnings and an examination of company balance sheets. They supplement this analysis with further detailed economic analysis, ordinarily including an evaluation of the quality of the firm's management, the firm's standing within its industry, and the prospects for the industry as a whole. The hope is to attain insight into future performance of the firm that is not yet recognized by the rest of the market. Chapters 17 through 19 provide a detailed discussion of the types of analyses that underlie fundamental analysis.

Once again, the efficient market hypothesis predicts that *most* fundamental analysis also is doomed to failure. If the analyst relies on publicly available earnings and industry information, his or her evaluation of the firm's prospects is not likely to be significantly more accurate than those of rival analysts. Many well-informed, well-financed firms conduct such market research, and in the face of such competition it will be difficult to uncover data not also available to other analysts. Only analysts with a unique insight will be rewarded.

Fundamental analysis is much more difficult than merely identifying well-run firms with good prospects. Discovery of good firms does an investor no good in and of itself if the rest of the market also knows those firms are good. If the knowledge is already public, the investor will be forced to pay a high price for those firms and will not realize a superior rate of return.

The trick is not to identify firms that are good, but to find firms that are *better* than everyone else's estimate. Similarly, poorly run firms can be great bargains if they are not quite as bad as their stock prices suggest.

This is why fundamental analysis is difficult. It is not enough to do a good analysis of a firm; you can make money only if your analysis is better than that of your competitors because the market price will already reflect all commonly recognized information.

Active versus Passive Portfolio Management

By now it is apparent that casual efforts to pick stocks are not likely to pay off. Competition among investors ensures that any easily implemented stock evaluation technique will be used widely enough so that any insights derived will be reflected in stock prices. Only serious analysis and uncommon techniques are likely to generate the *differential* insight necessary to yield trading profits.

Moreover, these techniques are economically feasible only for managers of large portfolios. If you have only $100,000 to invest, even a 1% per year improvement in performance generates only $1,000 per year, hardly enough to justify herculean efforts. The billion-dollar manager, however, reaps extra income of $10 million annually from the same 1% increment.

If small investors are not in a favored position to conduct active portfolio management, what are their choices? The small investor probably is better off investing in mutual funds. By pooling resources in this way, small investors can gain from economies of scale.

More difficult decisions remain, though. Can investors be sure that even large mutual funds have the ability or resources to uncover mispriced stocks? Furthermore, will any mispricing be sufficiently large to repay the costs entailed in active portfolio management?

Proponents of the efficient market hypothesis believe that active management is largely wasted effort and unlikely to justify the expenses incurred. Therefore, they advocate a **passive investment strategy** that makes no attempt to outsmart the market. A passive strategy aims only at establishing a well-diversified portfolio of securities without attempting to find under- or overvalued stocks. Passive management is usually characterized by a buy-and-hold strategy. Because the efficient market theory indicates that stock prices are at fair levels, given all available information, it makes no sense to buy and sell securities frequently, which generates large trading costs without increasing expected performance.

One common strategy for passive management is to create an **index fund,** which is a fund designed to replicate the performance of a broad-based index of stocks. For example, Vanguard's 500 Index Fund holds stocks in direct proportion to their weight in the Standard & Poor's 500 stock price index. The performance of the 500 Index Fund therefore replicates the performance of the S&P 500. Investors in this fund obtain broad diversification with relatively low management fees. The fees can be kept to a minimum because Vanguard does not need to pay analysts to assess stock prospects and does not incur transaction costs from high portfolio turnover. Indeed, while the typical annual charge for an actively managed equity fund is around 1% of assets, the expense ratio of the 500 Index Fund is only .17%. Vanguard's 500 Index Fund is among the largest equity mutual funds with over $100 billion of assets in 2012, and about 15% of assets in equity funds are indexed.

Indexing need not be limited to the S&P 500, however. For example, some of the funds offered by the Vanguard Group track the broader-based CRSP[8] index of the total U.S. equity market, the Barclays U.S. Aggregate Bond Index, the CRSP index of small-capitalization U.S. companies, and the *Financial Times* indexes of the European and Pacific Basin equity markets. Several other mutual fund complexes offer indexed portfolios, but Vanguard dominates the retail market for indexed products.

Exchange-traded funds, or ETFs, are a close (and often lower-expense) alternative to indexed mutual funds. As noted in Chapter 4, these are shares in diversified portfolios that can be bought or sold just like shares of individual stock. ETFs matching several broad stock market indexes such as the S&P 500 or CRSP indexes and dozens of international and industry stock indexes are available to investors who want to hold a diversified sector of a market without attempting active security selection.

CONCEPT CHECK 11.3

What would happen to market efficiency if *all* investors attempted to follow a passive strategy?

The Role of Portfolio Management in an Efficient Market

If the market is efficient, why not pick stocks by throwing darts at *The Wall Street Journal* instead of trying rationally to choose a stock portfolio? This is a tempting conclusion to draw from the notion that security prices are fairly set, but it is far too facile. There is a role for rational portfolio management, even in perfectly efficient markets.

You have learned that a basic principle in portfolio selection is diversification. Even if all stocks are priced fairly, each still poses firm-specific risk that can be eliminated through diversification. Therefore, rational security selection, even in an efficient market, calls for the selection of a well-diversified portfolio providing the systematic risk level that the investor wants.

Rational investment policy also requires that tax considerations be reflected in security choice. High-tax-bracket investors generally will not want the same securities that low

[8]CRSP is the Center for Research in Security Prices at the University of Chicago.

bracket investors find favorable. At an obvious level, high-bracket investors find it advantageous to buy tax-exempt municipal bonds despite their relatively low pretax yields, whereas those same bonds are unattractive to low-tax-bracket or tax-exempt investors. At a more subtle level, high-bracket investors might want to tilt their portfolios in the direction of capital gains as opposed to interest income, because capital gains are taxed less heavily and because the option to defer the realization of capital gains income is more valuable the higher the current tax bracket. Hence these investors may prefer stocks that yield low dividends yet offer greater expected capital gains income. They also will be more attracted to investment opportunities for which returns are sensitive to tax benefits, such as real estate ventures.

A third argument for rational portfolio management relates to the particular risk profile of the investor. For example, a Toyota executive whose annual bonus depends on Toyota's profits generally should not invest additional amounts in auto stocks. To the extent that his or her compensation already depends on Toyota's well-being, the executive is already overinvested in Toyota and should not exacerbate the lack of diversification. This lesson was learned with considerable pain in September 2008 by Lehman Brothers employees who were famously invested in their own firm when the company failed. Roughly 30% of the shares in the firm were owned by its 24,000 employees, and their losses on those shares totaled around $10 billion.

Investors of varying ages also might warrant different portfolio policies with regard to risk bearing. For example, older investors who are essentially living off savings might choose to avoid long-term bonds whose market values fluctuate dramatically with changes in interest rates (discussed in Part Four). Because these investors are living off accumulated savings, they require conservation of principal. In contrast, younger investors might be more inclined toward long-term inflation-indexed bonds. The steady flow of real income over long periods of time that is locked in with these bonds can be more important than preservation of principal to those with long life expectancies.

In conclusion, there is a role for portfolio management even in an efficient market. Investors' optimal positions will vary according to factors such as age, tax bracket, risk aversion, and employment. The role of the portfolio manager in an efficient market is to tailor the portfolio to these needs, rather than to beat the market.

Resource Allocation

We've focused so far on the investment implications of the efficient market hypothesis. Deviations from efficiency may offer profit opportunities to better-informed traders at the expense of less-informed ones.

However, deviations from informational efficiency would also result in a large cost that will be borne by all citizens, namely, inefficient resource allocation. Recall that in a capitalist economy, investments in *real* assets such as plant, equipment, and know-how are guided in large part by the prices of financial assets. For example, if the value of telecommunication capacity reflected in stock market prices exceeds the cost of installing such capacity, managers might justifiably conclude that telecom investments seem to have positive net present value. In this manner, capital market prices guide allocation of real resources.

If markets were inefficient and securities commonly mispriced, then resources would be systematically misallocated. Corporations with overpriced securities would be able to obtain capital too cheaply, and corporations with undervalued securities might forgo investment opportunities because the cost of raising capital would be too high. Therefore, inefficient capital markets would diminish one of the most potent benefits of a market economy. As an example of what can go wrong, consider the dot-com bubble of the late 1990s, which sent a strong but, as it turned out, wildly overoptimistic signal about prospects for Internet and telecommunication firms and ultimately led to substantial overinvestment in those industries.

Before writing off markets as a means to guide resource allocation, however, one has to be reasonable about what can be expected from market forecasts. In particular, you shouldn't confuse an efficient market, where all available information is reflected in prices, with a perfect-foresight market. As we said earlier, "all available information" is still far from complete information, and generally rational market forecasts will sometimes be wrong; sometimes, in fact, they will be very wrong.

11.3 Event Studies

The notion of informationally efficient markets leads to a powerful research methodology. If security prices reflect all currently available information, then price changes must reflect new information. Therefore, it seems that one should be able to measure the importance of an event of interest by examining price changes during the period in which the event occurs.

An **event study** describes a technique of empirical financial research that enables an observer to assess the impact of a particular event on a firm's stock price. A stock market analyst might want to study the impact of dividend changes on stock prices, for example. An event study would quantify the relationship between dividend changes and stock returns.

Analyzing the impact of any particular event is more difficult than it might at first appear. On any day, stock prices respond to a wide range of economic news such as updated forecasts for GDP, inflation rates, interest rates, or corporate profitability. Isolating the part of a stock price movement that is attributable to a specific event is not a trivial exercise.

The general approach starts with a proxy for what the stock's return would have been in the absence of the event. The **abnormal return** due to the event is estimated as the difference between the stock's actual return and this benchmark. Several methodologies for estimating the benchmark return are used in practice. For example, a very simple approach measures the stock's abnormal return as its return minus that of a broad market index. An obvious refinement is to compare the stock's return to those of other stocks matched according to criteria such as firm size, beta, recent performance, or ratio of price to book value per share. Another approach estimates normal returns using an asset pricing model such as the CAPM or one of its multifactor generalizations such as the Fama-French three-factor model.

Many researchers have used a "market model" to estimate abnormal returns. This approach is based on the index models we introduced in Chapter 9. Recall that a single-index model holds that stock returns are determined by a market factor and a firm-specific factor. The stock return, r_t, during a given period t, would be expressed mathematically as

$$r_t = a + br_{Mt} + e_t \tag{11.1}$$

where r_{Mt} is the market's rate of return during the period and e_t is the part of a security's return resulting from firm-specific events. The parameter b measures sensitivity to the market return, and a is the average rate of return the stock would realize in a period with a zero market return.[9] Equation 11.1 therefore provides a decomposition of r_t into market and firm-specific factors. The firm-specific or abnormal return may be interpreted as the unexpected return that results from the event.

Determination of the abnormal return in a given period requires an estimate of e_t. Therefore, we rewrite Equation 11.1:

$$e_t = r_t - (a - br_{Mt}) \tag{11.2}$$

[9]We know from Chapter 9 that the CAPM implies that the intercept a in Equation 11.1 should equal $r_f(1 - \beta)$. Nevertheless, it is customary to estimate the intercept in this equation empirically rather than imposing the CAPM value. One justification for this practice is that empirically fitted security market lines seem flatter than predicted by the CAPM (see Chapter 13), which would make the intercept implied by the CAPM too small.