Chapter 7

The Stock Market, the Theory of Rational Expectations, and the Efficient Market Hypothesis

PREVIEW

Rarely does a day go by that the stock market isn't a major news item. We have witnessed huge swings in the stock market in recent years. The 1990s were an extraordinary decade for stocks: the Dow Jones and S&P 500 indexes increased more than 400%, while the tech-laden NASDAQ index rose more than 1,000%. By early 2000, both indexes had reached record highs. Unfortunately, the good times did not last, and many investors lost their shirts. Starting in early 2000, the stock market began to decline: the NASDAQ crashed, falling by over 50%, while the Dow Jones and S&P 500 indexes fell by 30% through January 2003.

Because so many people invest in the stock market and the price of stocks affects the ability of people to retire comfortably, the market for stocks is undoubtedly the financial market that receives the most attention and scrutiny. In this chapter, we look at how this important market works.

We begin by discussing the fundamental theories that underlie the valuation of stocks. These theories are critical to understanding the forces that cause the value of stocks to rise and fall minute by minute and day by day. Once we have learned the methods for stock valuation, we need to explore how expectations about the market affect its behavior. We do so by examining the *theory of rational expectations*. When this theory is applied to financial markets, the outcome is the *efficient market hypothesis*. The theory of rational expectations is also central to debates about the conduct of monetary policy, to be discussed in Chapter 28.

Theoretically, the theory of rational expectations should be a powerful tool for analyzing behavior. But to establish that it is *in reality* a useful tool, we must compare the outcomes predicted by the theory with empirical evidence. Although the evidence is mixed and controversial, it indicates that for many purposes, the theory of rational expectations is a good starting point for analyzing expectations.

Computing the Price of Common Stock

Common stock is the principal way that corporations raise equity capital. Holders of common stock own an interest in the corporation consistent with the percentage of outstanding shares owned. This ownership interest gives **stockholders**—those who hold stock in a corporation—a bundle of rights. The most important are the right to vote and to be the **residual claimant** of all funds flowing into the firm (known as **cash flows**), meaning that the stockholder receives whatever remains after all other claims against the firm's assets have been satisfied. Stockholders are paid dividends from the net earnings of the corporation. **Dividends** are payments made periodically, usually every quarter, to stockholders. The board of directors of the firm sets the level of the dividend, usually upon the recommendation of management. In addition, the stockholder has the right to sell the stock.

http://stocks.tradingcharts.com

Access detailed stock quotes, charts, and historical stock data.

One basic principle of finance is that the value of any investment is found by computing the value today of all cash flows the investment will generate over its life. For example, a commercial building will sell for a price that reflects the net cash flows (rents – expenses) it is projected to have over its useful life. Similarly, we value common stock as the value in today's dollars of all future cash flows. The cash flows a stockholder might earn from stock are dividends, the sales price, or both.

To develop the theory of stock valuation, we begin with the simplest possible scenario: You buy the stock, hold it for one period to get a dividend, then sell the stock. We call this the *one-period valuation model*.

The One-Period Valuation Model

Suppose that you have some extra money to invest for one year. After a year, you will need to sell your investment to pay tuition. After watching *CNBC* or *Wall Street Week* on TV, you decide that you want to buy Intel Corp. stock. You call your broker and find that Intel is currently selling for $50 per share and pays $0.16 per year in dividends. The analyst on *Wall Street Week* predicts that the stock will be selling for $60 in one year. Should you buy this stock?

To answer this question, you need to determine whether the current price accurately reflects the analyst's forecast. To value the stock today, you need to find the present discounted value of the expected cash flows (future payments) using the formula in Equation 1 of Chapter 4. Note that in this equation, the discount factor used to discount the cash flows is the required return on investments in equity rather than the interest rate. The cash flows consist of one dividend payment plus a final sales price. When these cash flows are discounted back to the present, the following equation computes the current price of the stock:

$$P_0 = \frac{Div_1}{(1 + k_e)} + \frac{P_1}{(1 + k_e)} \quad (1)$$

where P_0 = the current price of the stock. The zero subscript refers to time period zero, or the present.
Div_1 = the dividend paid at the end of year 1.
k_e = the required return on investments in equity.
P_1 = the price at the end of the first period; the assumed sales price of the stock.

To see how Equation 1 works, let's compute the price of the Intel stock if, after careful consideration, you decide that you would be satisfied to earn a 12% return on the investment. If you have decided that $k_e = 0.12$, are told that Intel pays $0.16 per year in dividends ($Div_1 = 0.16$), and forecast the share price of $60 for next year ($P_1 = \$60$), you get the following from Equation 1:

$$P_0 = \frac{0.16}{1 + 0.12} + \frac{\$60}{1 + 0.12} = \$0.14 + \$53.57 = \$53.71$$

Based on your analysis, you find that the present value of all cash flows from the stock is \$53.71. Because the stock is currently priced at \$50 per share, you would choose to buy it. However, you should be aware that the stock may be selling for less than \$53.71, because other investors place a different risk on the cash flows or estimate the cash flows to be less than you do.

The Generalized Dividend Valuation Model

Using the same concept, the one-period dividend valuation model can be extended to any number of periods: The value of stock is the present value of all future cash flows. The only cash flows that an investor will receive are dividends and a final sales price when the stock is ultimately sold in period n. The generalized multi-period formula for stock valuation can be written as:

$$P_0 = \frac{D_1}{(1 + k_e)^1} + \frac{D_2}{(1 + k_e)^2} + \ldots + \frac{D_n}{(1 + k_e)^n} + \frac{P_n}{(1 + k_e)^n} \tag{2}$$

If you tried to use Equation 2 to find the value of a share of stock, you would soon realize that you must first estimate the value the stock will have at some point in the future before you can estimate its value today. In other words, you must find P_n in order to find P_0. However, if P_n is far in the future, it will not affect P_0. For example, the present value of a share of stock that sells for \$50 seventy-five years from now using a 12% discount rate is just one cent [$\$50/(1.12^{75}) = \$0.01$]. This reasoning implies that the current value of a share of stock can be calculated as simply the present value of the future dividend stream. The **generalized dividend model** is rewritten in Equation 3 without the final sales price:

$$P_0 = \sum_{t=1}^{\infty} \frac{D_t}{(1 + k_e)^t} \tag{3}$$

Consider the implications of Equation 3 for a moment. The generalized dividend model says that the price of stock is determined only by the present value of the dividends and that nothing else matters. Many stocks do not pay dividends, so how is it that these stocks have value? *Buyers of the stock expect that the firm will pay dividends someday.* Most of the time a firm institutes dividends as soon as it has completed the rapid growth phase of its life cycle.

The generalized dividend valuation model requires that we compute the present value of an infinite stream of dividends, a process that could be difficult, to say the least. Therefore, simplified models have been developed to make the calculations easier. One such model is the **Gordon growth model**, which assumes constant dividend growth.

The Gordon Growth Model

Many firms strive to increase their dividends at a constant rate each year. Equation 4 rewrites Equation 3 to reflect this constant growth in dividends:

$$P_0 = \frac{D_0 \times (1 + g)^1}{(1 + k_e)^1} + \frac{D_0 \times (1 + g)^2}{(1 + k_e)^2} + \cdots + \frac{D_0 \times (1 + g)^\infty}{(1 + k_e)^\infty} \tag{4}$$

where D_0 = the most recent dividend paid
g = the expected constant growth rate in dividends
k_e = the required return on an investment in equity

Equation 4 has been simplified using algebra to obtain Equation 5.[1]

$$P_0 = \frac{D_0 \times (1 + g)}{(k_e - g)} = \frac{D_1}{(k_e - g)} \tag{5}$$

This model is useful for finding the value of stock, given a few assumptions:

1. *Dividends are assumed to continue growing at a constant rate forever.* Actually, as long as they are expected to grow at a constant rate for an extended period of time, the model should yield reasonable results. This is because errors about distant cash flows become small when discounted to the present.
2. *The growth rate is assumed to be less than the required return on equity, k_e.* Myron Gordon, in his development of the model, demonstrated that this is a reasonable assumption. In theory, if the growth rate were faster than the rate demanded by holders of the firm's equity, in the long run the firm would grow impossibly large.

How the Market Sets Security Prices

Suppose you went to an auto auction. The cars are available for inspection before the auction begins, and you find a little Mazda Miata that you like. You test-drive it in the parking lot and notice that it makes a few strange noises, but you decide that you would still like the car. You decide $5,000 would be a fair price that would allow you to pay some repair bills should the noises turn out to be serious. You see that the auction is ready to begin, so you go in and wait for the Miata to enter.

Suppose there is another buyer who also spots the Miata. He test-drives the car and recognizes that the noises are simply the result of worn brake pads that he can fix himself at a nominal cost. He decides that the car is worth $7,000. He also goes in and waits for the Miata to enter.

Who will buy the car and for how much? Suppose only the two of you are interested in the Miata. You begin the bidding at $4,000. He ups your bid to $4,500. You

---

[1]To generate Equation 5 from Equation 4, first multiply both sides of Equation 4 by $(1 + k_e)/(1 + g)$ and subtract Equation 4 from the result. This yields:

$$\frac{P_0 \times (1 + k_e)}{(1 + g)} - P_0 = D_0 - \frac{D_0 \times (1 + g)^\infty}{(1 + k_e)^\infty}$$

Assuming that k_e is greater than g, the term on the far right will approach zero and can be dropped. Thus, after factoring P_0 out of the left-hand side:

$$P_0 \times \left[\frac{1 + k_e}{1 + g} - 1\right] = D_0$$

Next, simplify by combining terms to:

$$P_0 \times \frac{(1 + k_e) - (1 + g)}{1 + g} = D_0$$

$$P_0 = \frac{D_0 \times (1 + g)}{k_e - g} = \frac{D_1}{k_e - g}$$

bid your top price of $5,000. He counters with $5,100. The price is now higher than you are willing to pay, so you stop bidding. The car is sold to the more informed buyer for $5,100.

This simple example raises a number of points. First, the price is set by the buyer willing to pay the highest price. The price is not necessarily the highest price the asset could fetch, but it is incrementally greater than what any other buyer is willing to pay.

Second, the market price will be set by the buyer who can take best advantage of the asset. The buyer who purchased the car knew that he could fix the noise easily and cheaply. Because of this he was willing to pay more for the car than you were. The same concept holds for other assets. For example, a piece of property or a building will sell to the buyer who can put the asset to the most productive use.

Finally, the example shows the role played by information in asset pricing. Superior information about an asset can increase its value by reducing its risk. When you consider buying a stock, there are many unknowns about the future cash flows. The buyer who has the best information about these cash flows will discount them at a lower interest rate than will a buyer who is very uncertain.

Now let us apply these ideas to stock valuation. Suppose that you are considering the purchase of stock expected to pay a $2 dividend next year. Market analysts expect the firm to grow at 3% indefinitely. You are *uncertain* about both the constancy of the dividend stream and the accuracy of the estimated growth rate. To compensate yourself for this uncertainty (risk), you require a return of 15%.

Now suppose Jennifer, another investor, has spoken with industry insiders and feels more confident about the projected cash flows. Jennifer requires only a 12% return because her perceived risk is lower than yours. Bud, on the other hand, is dating the CEO of the company. He knows with more certainty what the future of the firm actually is, and thus requires only a 10% return.

What are the values each investor will give to the stock? Applying the Gordon growth model yields the following stock prices:

Investor	Discount Rate	Stock Price
You	15%	$16.67
Jennifer	12%	$22.22
Bud	10%	$28.57

You are willing to pay $16.67 for the stock. Jennifer would pay up to $22.22, and Bud would pay $28.57. The investor with the lowest perceived risk is willing to pay the most for the stock. If there were no other traders but these three, the market price would be between $22.22 and $28.57. If you already held the stock, you would sell it to Bud.

We thus see that the players in the market, bidding against each other, establish the market price. When new information is released about a firm, expectations change and with them, prices change. New information can cause changes in expectations about the level of future dividends or the risk of those dividends. Since market participants are constantly receiving new information and revising their expectations, it is reasonable that stock prices are constantly changing as well.

Application

Monetary Policy and Stock Prices

Stock market analysts tend to hang on every word that the Chairman of the Federal Reserve utters because they know that an important determinant of stock prices is monetary policy. But how does monetary policy affect stock prices?

The Gordon growth model in Equation 5 tells us how. Monetary policy can affect stock prices in two ways. First, when the Fed lowers interest rates, the return on bonds (an alternative asset to stocks) declines, and investors are likely to accept a lower required rate of return on an investment in equity (k_e). The resulting decline in k_e would lower the denominator in the Gordon growth model (Equation 5), lead to a higher value of P_0, and raise stock prices. Furthermore, a lowering of interest rates is likely to stimulate the economy, so that the growth rate in dividends, g, is likely to be somewhat higher. This rise in g also causes the denominator in Equation 5 to fall, which also leads to a higher P_0 and a rise in stock prices.

As we will see in Chapter 26, the impact of monetary policy on stock prices is one of the key ways in which monetary policy affects the economy.

Application

The September 11 Terrorist Attacks, the Enron Scandal, and the Stock Market

In 2001, two big shocks hit the stock market: the September 11 terrorist attacks and the Enron scandal. Our analysis of stock price evaluation, again using the Gordon growth model, can help us understand how these events affected stock prices.

The September 11 terrorist attacks raised the possibility that terrorism against the United States would paralyze the country. These fears led to a downward revision of the growth prospects for U.S. companies, thus lowering the dividend growth rate (g) in the Gordon model. The resulting rise in the denominator in Equation 5 would lead to a decline in P_0 and hence a decline in stock prices.

Increased uncertainty for the U.S. economy would also raise the required return on investment in equity. A higher k_e also leads to a rise in the denominator in Equation 5, a decline in P_0, and a general fall in stock prices. As the Gordon model predicts, the stock market fell by over 10% immediately after September 11.

Subsequently, the U.S. successes against the Taliban in Afghanistan and the absence of further terrorist attacks reduced market fears and uncertainty, causing g to recover and k_e to fall. The denominator in Equation 5 then fell, leading to a recovery in P_0 and a rebound in the stock market in October and November. However, by the beginning of 2002, the Enron scandal and disclosures that many companies had overstated their earnings caused many investors to doubt the formerly rosy forecast of earnings and dividend growth for corporations. The resulting revision of g downward, and the rise in k_e because of increased uncertainty about the quality of accounting information, would lead to a rise in the denominator in the Gordon Equation 5, thereby lowering P_0 for many companies and hence the overall stock market. As predicted by our analysis, this is exactly what happened. The stock market recovery was aborted, and the market began a downward slide.

The Theory of Rational Expectations

The analysis of stock price evaluation we have outlined in the previous section depends on people's expectations—especially of cash flows. Indeed, it is difficult to think of any sector in the economy in which expectations are not crucial; this is why it is important to examine how expectations are formed. We do so by outlining the *theory of rational expectations*, currently the most widely used theory to describe the formation of business and consumer expectations.

In the 1950s and 1960s, economists regularly viewed expectations as formed from past experience only. Expectations of inflation, for example, were typically viewed as being an average of past inflation rates. This view of expectation formation, called **adaptive expectations**, suggests that changes in expectations will occur slowly over time as past data change.[2] So if inflation had formerly been steady at a 5% rate, expectations of future inflation would be 5% too. If inflation rose to a steady rate of 10%, expectations of future inflation would rise toward 10%, but slowly: In the first year, expected inflation might rise only to 6%; in the second year, to 7%; and so on.

Adaptive expectations have been faulted on the grounds that people use more information than just past data on a single variable to form their expectations of that variable. Their expectations of inflation will almost surely be affected by their predictions of future monetary policy as well as by current and past monetary policy. In addition, people often change their expectations quickly in the light of new information. To meet these objections to adaptive expectations, John Muth developed an alternative theory of expectations, called **rational expectations**, which can be stated as follows: ***Expectations will be identical to optimal forecasts (the best guess of the future) using all available information.***[3]

What exactly does this mean? To explain it more clearly, let's use the theory of rational expectations to examine how expectations are formed in a situation that most of us encounter at some point in our lifetime: our drive to work. Suppose that when Joe Commuter travels when it is not rush hour, it takes an average of 30 minutes for

[2]More specifically, adaptive expectations—say, of inflation—are written as a weighted average of past inflation rates:

$$\pi_t^e = (1 - \lambda)\sum_{j=0}^{\infty}\lambda^j \pi_{t-j}$$

where π_t^e = adaptive expectation of inflation at time t
π_{t-j} = inflation at time $t - j$
λ = a constant between the values of 0 and 1

[3]John Muth, "Rational Expectations and the Theory of Price Movements," *Econometrica* 29 (1961): 315–335.

his trip. Sometimes it takes him 35 minutes, other times 25 minutes, but the average non-rush-hour driving time is 30 minutes. If, however, Joe leaves for work during the rush hour, it takes him, on average, an additional 10 minutes to get to work. Given that he leaves for work during the rush hour, the best guess of the driving time—the **optimal forecast**—is 40 minutes.

If the only information available to Joe before he leaves for work that would have a potential effect on his driving time is that he is leaving during the rush hour, what does rational expectations theory allow you to predict about Joe's expectations of his driving time? Since the best guess of his driving time using all available information is 40 minutes, Joe's expectation should also be the same. Clearly, an expectation of 35 minutes would not be rational, because it is not equal to the optimal forecast, the best guess of the driving time.

Suppose that the next day, given the same conditions and the same expectations, it takes Joe 45 minutes to drive because he hits an abnormally large number of red lights, and the day after that he hits all the lights right and it takes him only 35 minutes. Do these variations mean that Joe's 40-minute expectation is irrational? No, an expectation of 40 minutes' driving time is still a rational expectation. In both cases, the forecast is off by 5 minutes, so the expectation has not been perfectly accurate. However, the forecast does not have to be perfectly accurate to be rational—it need only be the *best possible* given the available information; that is, it has to be correct *on average,* and the 40-minute expectation meets this requirement. Since there is bound to be some randomness in Joe's driving time regardless of driving conditions, an optimal forecast will never be completely accurate.

The example makes the following important point about rational expectations: ***Even though a rational expectation equals the optimal forecast using all available information, a prediction based on it may not always be perfectly accurate.***

What if an item of information relevant to predicting driving time is unavailable or ignored? Suppose that on Joe's usual route to work there is an accident that causes a two-hour traffic jam. If Joe has no way of ascertaining this information, his rush-hour expectation of 40 minutes' driving time is still rational, because the accident information is not available to him for incorporation into his optimal forecast. However, if there was a radio or TV traffic report about the accident that Joe did not bother to listen to or heard but ignored, his 40-minute expectation is no longer rational. In light of the availability of this information, Joe's optimal forecast should have been two hours and 40 minutes.

Accordingly, there are two reasons why an expectation may fail to be rational:

1. People might be aware of all available information but find it takes too much effort to make their expectation the best guess possible.
2. People might be unaware of some available relevant information, so their best guess of the future will not be accurate.

Nonetheless, it is important to recognize that if an additional factor is important but information about it is not available, an expectation that does not take account of it can still be rational.

Formal Statement of the Theory

We can state the theory of rational expectations somewhat more formally. If X stands for the variable that is being forecast (in our example, Joe Commuter's driving time), X^e for the expectation of this variable (Joe's expectation of his driving time), and X^{of}

for the optimal forecast of X using all available information (the best guess possible of his driving time), the theory of rational expectations then simply says:

$$X^e = X^{of} \tag{6}$$

That is, the expectation of X equals the optimal forecast using all available information.

Rationale Behind the Theory

Why do people try to make their expectations match their best possible guess of the future using all available information? The simplest explanation is that it is costly for people not to do so. Joe Commuter has a strong incentive to make his expectation of the time it takes him to drive to work as accurate as possible. If he underpredicts his driving time, he will often be late to work and risk being fired. If he overpredicts, he will, on average, get to work too early and will have given up sleep or leisure time unnecessarily. Accurate expectations are desirable, and there are strong incentives for people to try to make them equal to optimal forecasts by using all available information.

The same principle applies to businesses. Suppose that an appliance manufacturer—say, General Electric—knows that interest-rate movements are important to the sales of appliances. If GE makes poor forecasts of interest rates, it will earn less profit, because it might produce either too many appliances or too few. There are strong incentives for GE to acquire all available information to help it forecast interest rates and use the information to make the best possible guess of future interest-rate movements.

The incentives for equating expectations with optimal forecasts are especially strong in financial markets. In these markets, people with better forecasts of the future get rich. The application of the theory of rational expectations to financial markets (where it is called the **efficient market hypothesis** or the **theory of efficient capital markets**) is thus particularly useful.

Implications of the Theory

Rational expectations theory leads to two commonsense implications for the forming of expectations that are important in the analysis of the aggregate economy:

1. ***If there is a change in the way a variable moves, the way in which expectations of this variable are formed will change as well.*** This tenet of rational expectations theory can be most easily understood through a concrete example. Suppose that interest rates move in such a way that they tend to return to a "normal" level in the future. If today's interest rate is high relative to the normal level, an optimal forecast of the interest rate in the future is that it will decline to the normal level. Rational expectations theory would imply that when today's interest rate is high, the expectation is that it will fall in the future.

Suppose now that the way in which the interest rate moves changes so that when the interest rate is high, it stays high. In this case, when today's interest rate is high, the optimal forecast of the future interest rate, and hence the rational expectation, is that it will stay high. Expectations of the future interest rate will no longer indicate that the interest rate will fall. The change in the way the interest-rate variable moves has therefore led to a change in the way that expectations of future interest rates are formed. The rational expectations analysis here is generalizable to expectations of any variable. Hence when there is a change in the way any variable moves, the way in which expectations of this variable are formed will change too.

2. ***The forecast errors of expectations will on average be zero and cannot be predicted ahead of time.*** The forecast error of an expectation is $X - X^e$, the difference between the realization of a variable X and the expectation of the variable; that is, if Joe Commuter's driving time on a particular day is 45 minutes and his expectation of the driving time is 40 minutes, the forecast error is 5 minutes.

Suppose that in violation of the rational expectations tenet, Joe's forecast error is not, on average, equal to zero; instead, it equals 5 minutes. The forecast error is now predictable ahead of time because Joe will soon notice that he is, on average, 5 minutes late for work and can improve his forecast by increasing it by 5 minutes. Rational expectations theory implies that this is exactly what Joe will do because he will want his forecast to be the best guess possible. When Joe has revised his forecast upward by 5 minutes, on average, the forecast error will equal zero so that it cannot be predicted ahead of time. Rational expectations theory implies that forecast errors of expectations cannot be predicted.

The Efficient Market Hypothesis: Rational Expectations in Financial Markets

www.investorhome.com/emh.htm

Learn more about the efficient market hypothesis.

While the theory of rational expectations was being developed by monetary economists, financial economists were developing a parallel theory of expectation formation in financial markets. It led them to the same conclusion as that of the rational expectations theorists: Expectations in financial markets are equal to optimal forecasts using all available information.[4] Although financial economists gave their theory another name, calling it *the efficient market hypothesis,* in fact their theory is just an application of rational expectations to the pricing of securities.

The efficient market hypothesis is based on the assumption that prices of securities in financial markets fully reflect all available information. You may recall from Chapter 4 that the rate of return from holding a security equals the sum of the capital gain on the security (the change in the price), plus any cash payments, divided by the initial purchase price of the security:

$$R = \frac{P_{t+1} - P_t + C}{P_t} \tag{7}$$

where R = rate of return on the security held from time t to $t + 1$ (say, the end of 2000 to the end of 2001)

P_{t+1} = price of the security at time $t + 1$, the end of the holding period

P_t = price of the security at time t, the beginning of the holding period

C = cash payment (coupon or dividend payments) made in the period t to $t + 1$

Let's look at the expectation of this return at time t, the beginning of the holding period. Because the current price P_t and the cash payment C are known at the beginning, the only variable in the definition of the return that is uncertain is the price next

[4]The development of the efficient market hypothesis was not wholly independent of the development of rational expectations theory, in that financial economists were aware of Muth's work.

period, P_{t+1}.[5] Denoting the expectation of the security's price at the end of the holding period as P^e_{t+1}, the expected return R^e is:

$$R^e = \frac{P^e_{t+1} - P_t + C}{P_t}$$

The efficient market hypothesis also views expectations of future prices as equal to optimal forecasts using all currently available information. In other words, the market's expectations of future securities prices are rational, so that:

$$P^e_{t+1} = P^{of}_{t+1}$$

which in turn implies that the expected return on the security will equal the optimal forecast of the return:

$$R^e = R^{of} \tag{8}$$

Unfortunately, we cannot observe either R^e or P^e_{t+1}, so the rational expectations equations by themselves do not tell us much about how the financial market behaves. However, if we can devise some way to measure the value of R^e, these equations will have important implications for how prices of securities change in financial markets.

The supply and demand analysis of the bond market developed in Chapter 5 shows us that the expected return on a security (the interest rate, in the case of the bond examined) will have a tendency to head toward the equilibrium return that equates the quantity demanded to the quantity supplied. Supply and demand analysis enables us to determine the expected return on a security with the following equilibrium condition: The expected return on a security R^e equals the equilibrium return R^*, which equates the quantity of the security demanded to the quantity supplied; that is,

$$R^e = R^* \tag{9}$$

The academic field of finance explores the factors (risk and liquidity, for example) that influence the equilibrium returns on securities. For our purposes, it is sufficient to know that we can determine the equilibrium return and thus determine the expected return with the equilibrium condition.

We can derive an equation to describe pricing behavior in an efficient market by using the equilibrium condition to replace R^e with R^* in the rational expectations equation (Equation 8). In this way, we obtain:

$$R^{of} = R^* \tag{10}$$

This equation tells us that ***current prices in a financial market will be set so that the optimal forecast of a security's return using all available information equals the security's equilibrium return***. Financial economists state it more simply: In an efficient market, a security's price fully reflects all available information.

Rationale Behind the Theory

Let's see what the efficient markets condition means in practice and why it is a sensible characterization of pricing behavior. Suppose that the equilibrium return on a security—say, Exxon common stock—is 10% at an annual rate, and its current price

[5]There are cases where *C* might not be known at the beginning of the period, but that does not make a substantial difference to the analysis. We would in that case assume that not only price expectations but also the expectations of *C* are optimal forecasts using all available information.

P_t is lower than the optimal forecast of tomorrow's price P^{of}_{t+1} so that the optimal forecast of the return at an annual rate is 50%, which is greater than the equilibrium return of 10%. We are now able to predict that, on average, Exxon's return would be abnormally high. This situation is called an **unexploited profit opportunity** because, on average, people would be earning more than they should, given the characteristics of that security. Knowing that, on average, you can earn such an abnormally high rate of return on Exxon because $R^{of} > R^*$, you would buy more, which would in turn drive up its current price P_t relative to the expected future price P^{of}_{t+1}, thereby lowering R^{of}. When the current price had risen sufficiently so that R^{of} equals R^* and the efficient markets condition (Equation 10) is satisfied, the buying of Exxon will stop, and the unexploited profit opportunity will have disappeared.

Similarly, a security for which the optimal forecast of the return is −5% and the equilibrium return is 10% ($R^{of} < R^*$) would be a poor investment, because, on average, it earns less than the equilibrium return. In such a case, you would sell the security and drive down its current price relative to the expected future price until R^{of} rose to the level of R^* and the efficient markets condition is again satisfied. What we have shown can be summarized as follows:

$$\begin{gathered} R^{of} > R^* \rightarrow P_t\uparrow \rightarrow R^{of}\downarrow \\ R^{of} < R^* \rightarrow P_t\downarrow \rightarrow R^{of}\uparrow \\ \text{until} \\ R^{of} = R^* \end{gathered}$$

Another way to state the efficient markets condition is this: ***In an efficient market, all unexploited profit opportunities will be eliminated.***

An extremely important factor in this reasoning is that ***not everyone in a financial market must be well informed about a security or have rational expectations for its price to be driven to the point at which the efficient markets condition holds***. Financial markets are structured so that many participants can play. As long as a few keep their eyes open for unexploited profit opportunities, they will eliminate the profit opportunities that appear, because in so doing, they make a profit. The efficient market hypothesis makes sense, because it does not require everyone in a market to be cognizant of what is happening to every security.

Stronger Version of the Efficient Market Hypothesis

Many financial economists take the efficient market hypothesis one step further in their analysis of financial markets. Not only do they define efficient markets as those in which expectations are rational—that is, equal to optimal forecasts using all available information—but they also add the condition that an efficient market is one in which prices reflect the true fundamental (intrinsic) value of the securities. Thus in an efficient market, all prices are always correct and reflect **market fundamentals** (items that have a direct impact on future income streams of the securities). This stronger view of market efficiency has several important implications in the academic field of finance. First, it implies that in an efficient capital market, one investment is as good as any other because the securities' prices are correct. Second, it implies that a security's price reflects all available information about the intrinsic value of the security. Third, it implies that security prices can be used by managers of both financial and nonfinancial firms to assess their cost of capital (cost of financing their investments) accurately and hence that security prices can be used to help them make the correct decisions about whether a specific investment is worth making or not. The stronger version of market efficiency is a basic tenet of much analysis in the finance field.

Evidence on the Efficient Market Hypothesis

Early evidence on the efficient market hypothesis was quite favorable to it, but in recent years, deeper analysis of the evidence suggests that the hypothesis may not always be entirely correct. Let's first look at the earlier evidence in favor of the hypothesis and then examine some of the more recent evidence that casts some doubt on it.

Evidence in Favor of Market Efficiency

Evidence in favor of market efficiency has examined the performance of investment analysts and mutual funds, whether stock prices reflect publicly available information, the random-walk behavior of stock prices, and the success of so-called technical analysis.

Performance of Investment Analysts and Mutual Funds. We have seen that one implication of the efficient market hypothesis is that when purchasing a security, you cannot expect to earn an abnormally high return, a return greater than the equilibrium return. This implies that it is impossible to beat the market. Many studies shed light on whether investment advisers and mutual funds (some of which charge steep sales commissions to people who purchase them) beat the market. One common test that has been performed is to take buy and sell recommendations from a group of advisers or mutual funds and compare the performance of the resulting selection of stocks with the market as a whole. Sometimes the advisers' choices have even been compared to a group of stocks chosen by throwing darts at a copy of the financial page of the newspaper tacked to a dartboard. The *Wall Street Journal*, for example, has a regular feature called "Investment Dartboard" that compares how well stocks picked by investment advisers do relative to stocks picked by throwing darts. Do the advisers win? To their embarrassment, the dartboard beats them as often as they beat the dartboard. Furthermore, even when the comparison includes only advisers who have been successful in the past in predicting the stock market, the advisers still don't regularly beat the dartboard.

Consistent with the efficient market hypothesis, mutual funds also do not beat the market. Not only do mutual funds not outperform the market on average, but when they are separated into groups according to whether they had the highest or lowest profits in a chosen period, the mutual funds that did well in the first period do not beat the market in the second period.[6]

The conclusion from the study of investment advisers and mutual fund performance is this: ***Having performed well in the past does not indicate that an investment adviser or a mutual fund will perform well in the future.*** This is not pleasing news to investment advisers, but it is exactly what the efficient market hypothesis predicts. It says that some advisers will be lucky and some will be unlucky. Being lucky does not mean that a forecaster actually has the ability to beat the market.

[6]An early study that found that mutual funds do not outperform the market is Michael C. Jensen, "The Performance of Mutual Funds in the Period 1945–64," *Journal of Finance* 23 (1968): 389–416. Further studies on mutual fund performance are Mark Grimblatt and Sheridan Titman, "Mutual Fund Performance: An Analysis of Quarterly Portfolio Holdings," *Journal of Business* 62 (1989): 393–416; R. A. Ippolito, "Efficiency with Costly Information: A Study of Mutual Fund Performance, 1965–84," *Quarterly Journal of Economics* 104 (1989): 1–23; J. Lakonishok, A. Shleifer, and R. Vishny, "The Structure and Performance of the Money Management Industry," *Brookings Papers on Economic Activity, Microeconomics* (1992); and B. Malkiel, "Returns from Investing in Equity Mutual Funds, 1971–1991," *Journal of Finance* 50 (1995): 549–72.

Do Stock Prices Reflect Publicly Available Information? The efficient market hypothesis predicts that stock prices will reflect all publicly available information. Thus if information is already publicly available, a positive announcement about a company will not, on average, raise the price of its stock because this information is already reflected in the stock price. Early empirical evidence also confirmed this conjecture from the efficient market hypothesis: Favorable earnings announcements or announcements of stock splits (a division of a share of stock into multiple shares, which is usually followed by higher earnings) do not, on average, cause stock prices to rise.[7]

Random-Walk Behavior of Stock Prices. The term **random walk** describes the movements of a variable whose future changes cannot be predicted (are random) because, given today's value, the variable is just as likely to fall as to rise. An important implication of the efficient market hypothesis is that stock prices should approximately follow a random walk; that is, ***future changes in stock prices should, for all practical purposes, be unpredictable***. The random-walk implication of the efficient market hypothesis is the one most commonly mentioned in the press, because it is the most readily comprehensible to the public. In fact, when people mention the "random-walk theory of stock prices," they are in reality referring to the efficient market hypothesis.

The case for random-walk stock prices can be demonstrated. Suppose that people could predict that the price of Happy Feet Corporation (HFC) stock would rise 1% in the coming week. The predicted rate of capital gains and rate of return on HFC stock would then be over 50% at an annual rate. Since this is very likely to be far higher than the equilibrium rate of return on HFC stock ($R^{of} > R^*$), the efficient markets hypothesis indicates that people would immediately buy this stock and bid up its current price. The action would stop only when the predictable change in the price dropped to near zero so that $R^{of} = R^*$.

Similarly, if people could predict that the price of HFC stock would fall by 1%, the predicted rate of return would be negative ($R^{of} < R^*$), and people would immediately sell. The current price would fall until the predictable change in the price rose back to near zero, where the efficient market condition again holds. The efficient market hypothesis suggests that the predictable change in stock prices will be near zero, leading to the conclusion that stock prices will generally follow a random walk.[8]

Financial economists have used two types of tests to explore the hypothesis that stock prices follow a random walk. In the first, they examine stock market records to see if changes in stock prices are systematically related to past changes and hence could have been predicted on that basis. The second type of test examines the data to see if publicly available information other than past stock prices could have been used to predict changes. These tests are somewhat more stringent because additional information (money supply growth, government spending, interest rates, corporate profits) might be used to help forecast stock returns. Early results from both types of tests

[7]Ray Ball and Philip Brown, "An Empirical Evaluation of Accounting Income Numbers," *Journal of Accounting Research* 6 (1968):159–178, and Eugene F. Fama, Lawrence Fisher, Michael C. Jensen, and Richard Roll, "The Adjustment of Stock Prices to New Information," *International Economic Review* 10 (1969): 1–21.

[8]Note that the random-walk behavior of stock prices is only an *approximation* derived from the efficient market hypothesis. It would hold exactly only for a stock for which an unchanged price leads to its having the equilibrium return. Then, when the predictable change in the stock price is exactly zero, $R^{of} = R^*$.

generally confirmed the efficient market view that stock prices are not predictable and follow a random walk.[9]

Technical Analysis. A popular technique used to predict stock prices, called *technical analysis*, is to study past stock price data and search for patterns such as trends and regular cycles. Rules for when to buy and sell stocks are then established on the basis of the patterns that emerge. The efficient market hypothesis suggests that technical analysis is a waste of time. The simplest way to understand why is to use the random-walk result derived from the efficient market hypothesis that holds that past stock price data cannot help predict changes. Therefore, technical analysis, which relies on such data to produce its forecasts, cannot successfully predict changes in stock prices.

Two types of tests bear directly on the value of technical analysis. The first performs the empirical analysis described earlier to evaluate the performance of any financial analyst, technical or otherwise. The results are exactly what the efficient market hypothesis predicts: Technical analysts fare no better than other financial analysts; on average, they do not outperform the market, and successful past forecasting does not imply that their forecasts will outperform the market in the future. The second type of test (first performed by Sidney Alexander) takes the rules developed in technical analysis for when to buy and sell stocks and applies them to new data.[10] The performance of these rules is then evaluated by the profits that would have been made using them. These tests also discredit technical analysis: It does not outperform the overall market.

Application

Should Foreign Exchange Rates Follow a Random Walk?

Although the efficient market hypothesis is usually applied to the stock market, it can also be used to show that foreign exchange rates, like stock prices, should generally follow a random walk. To see why this is the case, consider what would happen if people could predict that a currency would appreciate

[9]The first type of test, using only stock market data, is referred to as a test of *weak-form efficiency*, because the information that can be used to predict stock prices is restricted to past price data. The second type of test is referred to as a test of *semistrong-form efficiency*, because the information set is expanded to include all publicly available information, not just past stock prices. A third type of test is called a test of *strong-form efficiency*, because the information set includes insider information, known only to the managers (directors) of the corporation, as when they plan to declare a high dividend. Strong-form tests do sometimes indicate that insider information can be used to predict changes in stock prices. This finding does not contradict the efficient market hypothesis, because the information is not available to the market and hence cannot be reflected in market prices. In fact, there are strict laws against using insider information to trade in financial markets. For an early survey on the three forms of tests, see Eugene F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25 (1970): 383–416.

[10]Sidney Alexander, "Price Movements in Speculative Markets: Trends or Random Walks?" *Industrial Management Review,* May 1961, pp. 7–26, and Sidney Alexander, "Price Movements in Speculative Markets: Trends or Random Walks? No. 2," in *The Random Character of Stock Prices,* ed. Paul Cootner (Cambridge, Mass.: MIT Press, 1964), pp. 338–372. More recent evidence also seems to discredit technichal analysis; for example, F. Allen and R. Karjalainen, "Using Genetic Algorithms to Find Technical Trading Rules," *Journal of Financial Economics* 51 (1999): 245–271. However, some other research is more favorable to technical analysis: e.g., R. Sullivan, A. Timmerman, and H. White, "Data-Snooping, Technical Trading Rule Performance and the Bootstrap," Centre for Economic Policy Research Discussion Paper No. 1976, 1998.

by 1% in the coming week. By buying this currency, they could earn a greater than 50% return at an annual rate, which is likely to be far above the equilibrium return for holding a currency. As a result, people would immediately buy the currency and bid up its current price, thereby reducing the expected return. The process would stop only when the predictable change in the exchange rate dropped to near zero so that the optimal forecast of the return no longer differed from the equilibrium return. Likewise, if people could predict that the currency would depreciate by 1% in the coming week, they would sell it until the predictable change in the exchange rate was again near zero. The efficient market hypothesis therefore implies that future changes in exchange rates should, for all practical purposes, be unpredictable; in other words, exchange rates should follow random walks. This is exactly what empirical evidence finds.[11]

Evidence Against Market Efficiency

All the early evidence supporting the efficient market hypothesis appeared to be overwhelming, causing Eugene Fama, a prominent financial economist, to state in his famous 1970 survey of the empirical evidence on the efficient market hypothesis, "The evidence in support of the efficient markets model is extensive, and (somewhat uniquely in economics) contradictory evidence is sparse."[12] However, in recent years, the hypothesis has begun to show a few cracks, referred to as *anomalies*, and empirical evidence indicates that the efficient market hypothesis may not always be generally applicable.

Small-Firm Effect. One of the earliest reported anomalies in which the stock market did not appear to be efficient is called the *small-firm effect*. Many empirical studies have shown that small firms have earned abnormally high returns over long periods of time, even when the greater risk for these firms has been taken into account.[13] The small-firm effect seems to have diminished in recent years, but is still a challenge to the efficient market hypothesis. Various theories have been developed to explain the small-firm effect, suggesting that it may be due to rebalancing of portfolios by institutional investors, tax issues, low liquidity of small-firm stocks, large information costs in evaluating small firms, or an inappropriate measurement of risk for small-firm stocks.

January Effect. Over long periods of time, stock prices have tended to experience an abnormal price rise from December to January that is predictable and hence inconsistent with random-walk behavior. This so-called **January effect** seems to have diminished in recent years for shares of large companies but still occurs for shares of

---

[11] See Richard A. Meese and Kenneth Rogoff, "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14 (1983): 3–24.

[12] Eugene F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25 (1970): 383–416.

[13] For example, see Marc R. Reinganum, "The Anomalous Stock Market Behavior of Small Firms in January: Empirical Tests of Tax Loss Selling Effects," *Journal of Financial Economics* 12 (1983): 89–104; Jay R. Ritter, "The Buying and Selling Behavior of Individual Investors at the Turn of the Year," *Journal of Finance* 43 (1988): 701–717; and Richard Roll, "Vas Ist Das? The Turn-of-the-Year Effect: Anomaly or Risk Mismeasurement?" *Journal of Portfolio Management* 9 (1988): 18–28.

[14] For example, see Donald B. Keim, "The CAPM and Equity Return Regularities," *Financial Analysts Journal* 42 (May–June 1986): 19–34.

small companies.[14] Some financial economists argue that the January effect is due to tax issues. Investors have an incentive to sell stocks before the end of the year in December, because they can then take capital losses on their tax return and reduce their tax liability. Then when the new year starts in January, they can repurchase the stocks, driving up their prices and producing abnormally high returns. Although this explanation seems sensible, it does not explain why institutional investors such as private pension funds, which are not subject to income taxes, do not take advantage of the abnormal returns in January and buy stocks in December, thus bidding up their price and eliminating the abnormal returns.[15]

Market Overreaction. Recent research suggests that stock prices may overreact to news announcements and that the pricing errors are corrected only slowly.[16] When corporations announce a major change in earnings—say, a large decline—the stock price may overshoot, and after an initial large decline, it may rise back to more normal levels over a period of several weeks. This violates the efficient market hypothesis, because an investor could earn abnormally high returns, on average, by buying a stock immediately after a poor earnings announcement and then selling it after a couple of weeks when it has risen back to normal levels.

Excessive Volatility. A phenomenon closely related to market overreaction is that the stock market appears to display *excessive volatility*; that is, fluctuations in stock prices may be much greater than is warranted by fluctuations in their fundamental value. In an important paper, Robert Shiller of Yale University found that fluctuations in the S&P 500 stock index could not be justified by the subsequent fluctuations in the dividends of the stocks making up this index. There has been much subsequent technical work criticizing these results, but Shiller's work, along with research finding that there are smaller fluctuations in stock prices when stock markets are closed, has produced a consensus that stock market prices appear to be driven by factors other than fundamentals.[17]

Mean Reversion. Some researchers have also found that stock returns display **mean reversion**: Stocks with low returns today tend to have high returns in the future, and vice versa. Hence stocks that have done poorly in the past are more likely to do well in the future, because mean reversion indicates that there will be a predictable positive change in the future price, suggesting that stock prices are not a random walk. Other researchers have found that mean reversion is not nearly as strong in data after World

---

[15]Another anomaly that makes the stock market seem less than efficient is that the *Value Line Survey*, one of the most prominent investment advice newsletters, has produced stock recommendations that have yielded abnormally high returns on average. See Fischer Black, "Yes, Virginia, There Is Hope: Tests of the Value Line Ranking System," *Financial Analysts Journal* 29 (September–October 1973): 10–14, and Gur Huberman and Shmuel Kandel, "Market Efficiency and Value Line's Record," *Journal of Business* 63 (1990): 187–216. Whether the excellent performance of the *Value Line Survey* will continue in the future is, of course, a question mark.

[16]Werner De Bondt and Richard Thaler, "Further Evidence on Investor Overreaction and Stock Market Seasonality," *Journal of Finance* 62 (1987): 557–580.

[17]Robert Shiller, "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (1981): 421–436, and Kenneth R. French and Richard Roll, "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics* 17 (1986): 5–26.

War II and so have raised doubts about whether it is currently an important phenomenon. The evidence on mean reversion remains controversial.[18]

New Information Is Not Always Immediately Incorporated into Stock Prices. Although it is generally found that stock prices adjust rapidly to new information, as is suggested by the efficient market hypothesis, recent evidence suggests that, inconsistent with the efficient market hypothesis, stock prices do not instantaneously adjust to profit announcements. Instead, on average stock prices continue to rise for some time after the announcement of unexpectedly high profits, and they continue to fall after surprisingly low profit announcments.[19]

Overview of the Evidence on the Efficient Market Hypothesis

As you can see, the debate on the efficient market hypothesis is far from over. The evidence seems to suggest that the efficient market hypothesis may be a reasonable starting point for evaluating behavior in financial markets. However, there do seem to be important violations of market efficiency that suggest that the efficient market hypothesis may not be the whole story and so may not be generalizable to all behavior in financial markets.

Application

Practical Guide to Investing in the Stock Market

The efficient market hypothesis has numerous applications to the real world. It is especially valuable because it can be applied directly to an issue that concerns many of us: how to get rich (or at least not get poor) in the stock market. (The "Following the Financial News" box shows how stock prices are reported daily.) A practical guide to investing in the stock market, which we develop here, provides a better understanding of the use and implications of the efficient market hypothesis.

How Valuable Are Published Reports by Investment Advisers?

Suppose you have just read in the "Heard on the Street" column of the *Wall Street Journal* that investment advisers are predicting a boom in oil stocks because an oil shortage is developing. Should you proceed to withdraw all your hard-earned savings from the bank and invest it in oil stocks?

[18]Evidence for mean reversion has been reported by James M. Poterba and Lawrence H. Summers, "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics* 22 (1988): 27–59; Eugene F. Fama and Kenneth R. French, "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy* 96 (1988): 246–273; and Andrew W. Lo and A. Craig MacKinlay, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies* 1 (1988): 41–66. However, Myung Jig Kim, Charles R. Nelson, and Richard Startz, in "Mean Reversion in Stock Prices? A Reappraisal of the Evidence," *Review of Economic Studies* 58 (1991): 515–528, question whether some of these findings are valid. For an excellent summary of this evidence, see Charles Engel and Charles S. Morris, "Challenges to Stock Market Efficiency: Evidence from Mean Reversion Studies," *Federal Reserve Bank of Kansas City Economic Review,* September–October 1991, pp. 21–35. See also N. Jegadeesh and Sheridan Titman, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance* 48 (1993): 65–92, which shows that mean reversion also occurs for individual stocks.

[19]For example, see R. Ball and P. Brown, "An Empirical Evaluation of Accounting Income Numbers," *Journal of Accounting Research* 6 (1968): 159–178, L. Chan, N. Jegadeesh, and J. Lakonishok, "Momentum Strategies," *Journal of Finance* 51 (1996): 1681–1713, and Eugene Fama, "Market Efficiency, Long-Term Returns and Behavioral Finance," *Journal of Financial Economics* 49 (1998): 283–306.

Following the Financial News

Stock Prices

Stock prices are published daily, and in the *Wall Street Journal*, they are reported in the sections "NYSE—Composite Transactions," "Amex—Composite Transactions," and "NASDAQ National Market Issues." Stock prices are quoted in the following format:

YTD % Chg.	52-Week Hi	52-Week Lo	Stock (Sym.)	Div.	Yld. %	PE	Vol. 100s	Close	Net Chg.
0.6	23.85	15.50♣	IntAlum **IAL**	1.20	6.9	88	21	17.39	0.10
4.0	126.39	54.01	IBM **IBM**	.60	.7	29	76523	80.57	3.07
1.9	37.45	26.05	IntFlavor **IFF**	.60	1.7	21	5952	35.78	0.68
2.9	80.10	47.75♣	IntGameTch **IGT**		...	24	9427	78.15	2.23

Source: Wall Street Journal, January 3, 2003, p. C4.

The following information is included in each column. International Business Machines (IBM) common stock is used as an example.

YTD % Chg: The stock price percentage change for the calendar year to date, adjusted for stock splits and dividends over 10%
52 Weeks Hi: Highest price of a share in the past 52 weeks: 126.39 for IBM stock
52 Weeks Lo: Lowest price of a share in the past 52 weeks: 54.01 for IBM stock
Stock: Company name: IBM for International Business Machines
Sym: Symbol that identifies company: IBM
Div: Annual dividends: $0.60 for IBM
Yld %: Yield for stock expressed as annual dividends divided by today's closing price: 0.7% (= 0.6 ÷ 80.57) for IBM stock
PE: Price-earnings ratio; the stock price divided by the annual earnings per share: 29
Vol 100s: Number of shares (in hundreds) traded that day: 7,652,300 shares for IBM
Close: Closing price (last price) that day: 80.57
Net Chg: Change in the closing price from the previous day: 3.07

Prices quoted for shares traded over-the-counter (through dealers rather than on an organized exchange) are sometimes quoted with the same information, but in many cases only the bid price (the price the dealer is willing to pay for the stock) and the asked price (the price the dealer is willing to sell the stock for) are quoted.

The efficient market hypothesis tells us that when purchasing a security, we cannot expect to earn an abnormally high return, a return greater than the equilibrium return. Information in newspapers and in the published reports of investment advisers is readily available to many market participants and is already reflected in market prices. So acting on this information will not yield abnormally high returns, on average. As we have seen, the empirical evidence for the most part confirms that recommendations from investment advisers cannot help us outperform the general market. Indeed, as Box 1 suggests, human investment advisers in San Francisco do not on average even outperform an orangutan!

Probably no other conclusion is met with more skepticism by students than this one when they first hear it. We all know or have heard of somebody who has been successful in the stock market for a period of many years. We wonder, "How could someone be so consistently successful if he or she did not really know how to predict when returns would be abnormally high?"

Box 1

Should You Hire an Ape as Your Investment Adviser?

The *San Francisco Chronicle* came up with an amusing way of evaluating how successful investment advisers are at picking stocks. They asked eight analysts to pick five stocks at the beginning of the year and then compared the performance of their stock picks to those chosen by Jolyn, an orangutan living at Marine World/Africa USA in Vallejo, California. Consistent with the results found in the "Investment Dartboard" feature of the *Wall Street Journal*, Jolyn beat the investment advisers as often as they beat her. Given this result, you might be just as well off hiring an orangutan as your investment adviser as you would hiring a human being!

The following story, reported in the press, illustrates why such anecdotal evidence is not reliable.

A get-rich-quick artist invented a clever scam. Every week, he wrote two letters. In letter A, he would pick team A to win a particular football game, and in letter B, he would pick the opponent, team B. A mailing list would then be separated into two groups, and he would send letter A to the people in one group and letter B to the people in the other. The following week he would do the same thing but would send these letters only to the group who had received the first letter with the correct prediction. After doing this for ten games, he had a small cluster of people who had received letters predicting the correct winning team for every game. He then mailed a final letter to them, declaring that since he was obviously an expert predictor of the outcome of football games (he had picked winners ten weeks in a row) and since his predictions were profitable for the recipients who bet on the games, he would continue to send his predictions only if he were paid a substantial amount of money. When one of his clients figured out what he was up to, the con man was prosecuted and thrown in jail!

What is the lesson of the story? Even if no forecaster is an accurate predictor of the market, there will always be a group of consistent winners. A person who has done well regularly in the past cannot guarantee that he or she will do well in the future. Note that there will also be a group of persistent losers, but you rarely hear about them because no one brags about a poor forecasting record.

Should You Be Skeptical of Hot Tips?

Suppose your broker phones you with a hot tip to buy stock in the Happy Feet Corporation (HFC) because it has just developed a product that is completely effective in curing athlete's foot. The stock price is sure to go up. Should you follow this advice and buy HFC stock?

The efficient market hypothesis indicates that you should be skeptical of such news. If the stock market is efficient, it has already priced HFC stock so that its expected return will equal the equilibrium return. The hot tip is not particularly valuable and will not enable you to earn an abnormally high return.

You might wonder, though, if the hot tip is based on new information and would give you an edge on the rest of the market. If other market participants have gotten this information before you, the answer is no. As soon as the information hits the street, the unexploited profit opportunity it creates will be quickly eliminated. The stock's price will already reflect the information, and you should expect to realize only the equilibrium return. But if you are one of the first to gain the new information, it can do you some good. Only then can you be one of the lucky ones who, on average, will earn an abnormally high return by helping eliminate the profit opportunity by buying HFC stock.

Do Stock Prices Always Rise When There Is Good News?

If you follow the stock market, you might have noticed a puzzling phenomenon: When good news about a stock, such as a particularly favorable earnings report, is announced, the price of the stock frequently does not rise. The efficient market hypothesis and the random-walk behavior of stock prices explain this phenomenon.

Because changes in stock prices are unpredictable, when information is announced that has already been expected by the market, the stock price will remain unchanged. The announcement does not contain any new information that should lead to a change in stock prices. If this were not the case and the announcement led to a change in stock prices, it would mean that the change was predictable. Because that is ruled out in an efficient market, ***stock prices will respond to announcements only when the information being announced is new and unexpected***. If the news is expected, there will be no stock price response. This is exactly what the evidence we described earlier, which shows that stock prices reflect publicly available information, suggests will occur.

Sometimes an individual stock price declines when good news is announced. Although this seems somewhat peculiar, it is completely consistent with the workings of an efficient market. Suppose that although the announced news is good, it is not as good as expected. HFC's earnings may have risen 15%, but if the market expected earnings to rise by 20%, the new information is actually unfavorable, and the stock price declines.

Efficient Market Prescription for the Investor

What does the efficient market hypothesis recommend for investing in the stock market? It tells us that hot tips, investment advisers' published recommendations, and technical analysis—all of which make use of publicly available information—cannot help an investor outperform the market. Indeed, it indicates that anyone without better information than other market participants cannot expect to beat the market. So what is an investor to do?

The efficient market hypothesis leads to the conclusion that such an investor (and almost all of us fit into this category) should not try to outguess the market by constantly buying and selling securities. This process does nothing but boost the income of brokers, who earn commissions on each trade.[20] Instead, the investor should pursue a "buy and hold" strategy—purchase stocks and hold them for long periods of time. This will lead to the same returns, on average, but the investor's net profits will be higher, because fewer brokerage commissions will have to be paid.

---

[20]The investor may also have to pay Uncle Sam capital gains taxes on any profits that are realized when a security is sold—an additional reason why continual buying and selling does not make sense.

It is frequently a sensible strategy for a small investor, whose costs of managing a portfolio may be high relative to its size, to buy into a mutual fund rather than individual stocks. Because the efficient market hypothesis indicates that no mutual fund can consistently outperform the market, an investor should not buy into one that has high management fees or that pays sales commissions to brokers, but rather should purchase a no-load (commission-free) mutual fund that has low management fees.

As we have seen, the evidence indicates that it will not be easy to beat the prescription suggested here, although some of the anomalies to the efficient market hypothesis suggest that an extremely clever investor (which rules out most of us) may be able to outperform a buy-and-hold strategy.

Evidence on Rational Expectations in Other Markets

Evidence in other financial markets also supports the efficient market hypothesis and hence the rationality of expectations. For example, there is little evidence that financial analysts are able to outperform the bond market.[21] The returns on bonds appear to conform to the efficient markets condition of Equation 10.

Rationality of expectations is, however, much harder to test in markets other than financial markets, because price data that reflect expectations are not as readily available. The most common tests of rational expectations in these markets make use of survey data on the forecasts of market participants. For example, one well-known study by James Pesando used a survey of inflation expectations collected from prominent economists and inflation forecasters.[22] In that survey, these people were asked what they predicted the inflation rate would be over the next six months and over the next year. Because rational expectations theory implies that forecast errors should on average be zero and cannot be predicted, tests of the theory involve asking whether the forecast errors in a survey could be predicted ahead of time using publicly available information. The evidence from Pesando's and subsequent studies is mixed. Sometimes the forecast errors cannot be predicted, and at other times they can. The evidence is not as supportive of rational expectations theory as the evidence from financial markets.

Does the fact that forecast errors from surveys are often predictable suggest that we should reject rational expectations theory in these other markets? The answer is: not necessarily. One problem with this evidence is that the expectations data are obtained from surveys rather than from actual economic decisions of market participants. That is a serious criticism of this evidence. Survey responses are not always reliable, because there is little incentive for participants to tell the truth. For example, when people are asked in surveys how much television they watch, responses greatly underestimate the actual time spent. Neither are people very truthful about the shows

[21] See the discussion in Frederic S. Mishkin, "Efficient Markets Theory: Implications for Monetary Policy," *Brookings Papers on Economic Activity* 3 (1978): 707–768, of the results in Michael J. Prell, "How Well Do the Experts Forecast Interest Rates?" *Federal Reserve Bank of Kansas City Monthly Review*, September–October 1973, pp. 3–15.

[22] James Pesando, "A Note on the Rationality of the Livingston Price Expectations," *Journal of Political Economy* 83 (1975): 845–858.

they watch. They may say they watch ballet on public television, but we know they are actually watching Vanna White light up the letters on *Wheel of Fortune* instead, because it, not ballet, gets high Nielsen ratings. How many people will admit to being regular watchers of *Wheel of Fortune*?

A second problem with survey evidence is that a market's behavior may not be equally influenced by the expectations of all the survey participants, making survey evidence a poor guide to market behavior. For example, we have already seen that prices in financial markets often *behave* as if expectations are rational even though many of the market participants do not have rational expectations.[23]

Proof is not yet conclusive on the validity of rational expectations theory in markets other than financial markets. One important conclusion, however, that is supported by the survey evidence is that ***if there is a change in the way a variable moves, there will be a change in the way expectations of this variable are formed as well***.

Application

What Do the Black Monday Crash of 1987 and the Tech Crash of 2000 Tell Us About Rational Expectations and Efficient Markets?

On October 19, 1987, dubbed "Black Monday," the Dow Jones Industrial Average declined more than 20%, the largest one-day decline in U.S. history. The collapse of the high-tech companies' share prices from their peaks in March 2000 caused the heavily tech-laden NASDAQ index to fall from around 5,000 in March 2000 to around 1,500 in 2001 and 2002, for a decline of well over 60%. These two crashes have caused many economists to question the validity of efficient markets and rational expectations. They do not believe that a rational marketplace could have produced such a massive swing in share prices. To what degree should these stock market crashes make us doubt the validity of rational expectations and the efficient market hypothesis?

Nothing in rational expectations theory rules out large changes in stock prices. A large change in stock prices can result from new information that produces a dramatic decline in optimal forecasts of the future valuation of firms. However, economists are hard pressed to come up with fundamental changes in the economy that can explain the Black Monday and tech crashes. One lesson from these crashes is that factors other than market fundamentals probably have an effect on stock prices. Hence these crashes have convinced many economists that the stronger version of the efficient market hypothesis, which states that asset prices reflect the true fundamental (intrinsic) value of securities, is incorrect. They attribute a large role in determination of stock prices to market psychology and to the institutional structure of the marketplace. However, nothing in this view contradicts the basic reasoning behind rational expectations or the efficient market hypothesis—that market participants eliminate unexploited profit opportunities. Even though stock market prices may not always solely reflect

[23]There is some fairly strong evidence for this proposition. For example, Frederic S. Mishkin, "Are Market Forecasts Rational?" *American Economic Review* 71 (1981): 295–306, finds that although survey forecasts of short-term interest rates are not rational, the bond market *behaves* as if the expectations of these interest rates are rational.

market fundamentals, this does not mean that rational expectations do not hold. As long as stock market crashes are unpredictable, the basic lessons of the theory of rational expectations hold.

Some economists have come up with theories of what they call *rational bubbles* to explain stock market crashes. A **bubble** is a situation in which the price of an asset differs from its fundamental market value. In a rational bubble, investors can have rational expectations that a bubble is occurring because the asset price is above its fundamental value but continue to hold the asset anyway. They might do this because they believe that someone else will buy the asset for a higher price in the future. In a rational bubble, asset prices can therefore deviate from their fundamental value for a long time because the bursting of the bubble cannot be predicted and so there are no unexploited profit opportunities.

However, other economists believe that the Black Monday crash of 1987 and the tech crash of 2000 suggest that there may be unexploited profit opportunities and that the theory of rational expectations and the efficient market hypothesis might be fundamentally flawed. The controversy over whether capital markets are efficient or expectations are rational continues.

Summary

1. Stocks are valued as the present value of future dividends. Unfortunately, we do not know very precisely what these dividends will be. This introduces a great deal of error to the valuation process. The Gordon growth model is a simplified method of computing stock value that depends on the assumption that the dividends are growing at a constant rate forever. Given our uncertainty regarding future dividends, this assumption is often the best we can do.
2. The interaction among traders in the market is what actually sets prices on a day-to-day basis. The trader that values the security the most (either because of less uncertainty about the cash flows or because of greater estimated cash flows) will be willing to pay the most. As new information is released, investors will revise their estimates of the true value of the security and will either buy or sell it depending upon how the market price compares to their estimated valuation. Because small changes in estimated growth rates or required return result in large changes in price, it is not surprising that the markets are often volatile.
3. The efficient market hypothesis states that current security prices will fully reflect all available information, because in an efficient market, all unexploited profit opportunities are eliminated. The elimination of unexploited profit opportunities necessary for a financial market to be efficient does not require that all market participants be well informed.
4. The evidence on the efficient market hypothesis is quite mixed. Early evidence on the performance of investment analysts and mutual funds, whether stock prices reflect publicly available information, the random-walk behavior of stock prices, and the success of so-called technical analysis was quite favorable to the efficient market hypothesis. However, in recent years, evidence on the small-firm effect, the January effect, market overreaction, excessive volatility, mean reversion, and new information is not always incorporated into stock prices, suggesting that the hypothesis may not always be entirely correct. The evidence seems to suggest that the efficient market hypothesis may be a reasonable starting point for evaluating behavior in financial markets but may not be generalizable to all behavior in financial markets.
5. The efficient market hypothesis indicates that hot tips, investment advisers' published recommendations, and technical analysis cannot help an investor out-perform the market. The prescription for investors is to pursue a

buy-and-hold strategy—purchase stocks and hold them for long periods of time. Empirical evidence generally supports these implications of the efficient market hypothesis in the stock market.

6. The stock market crash of 1987 and the tech crash of 2000 have convinced many financial economists that the stronger version of the efficient market hypothesis, which states that asset prices reflect the true fundamental (intrinsic) value of securities, is not correct. It is less clear that these crashes shows that the weaker version of the efficient market hypothesis is wrong. Even if the stock market was driven by factors other than fundamentals, these crashes do not clearly demonstrate that many of the basic lessons of the efficient market hypothesis are no longer valid, as long as these crashes could not have been predicted.

Key Terms

adaptive expectations, p. 147
bubble, p. 164
cash flows, p. 141
dividends, p. 142
efficient market hypothesis, p. 149
generalized dividend model, p. 143
Gordon growth model, p. 143
January effect, p. 156
market fundamentals, p. 152
mean reversion, p. 157
optimal forecast, p. 148
random walk, p. 154
rational expectations, p. 147
residual claimant, p. 141
stockholders, p. 141
theory of efficient capital markets, p. 149
unexploited profit opportunity, p. 152

Questions and Problems

Questions marked with an asterisk are answered at the end of the book in an appendix, "Answers to Selected Questions and Problems."

1. What basic principle of finance can be applied to the valuation of any investment asset?

*2. Identify the cash flows available to an investor in stock. How reliably can these cash flows be estimated? Compare the problem of estimating stock cash flows to estimating bond cash flows. Which security would you predict to be more volatile?

3. Compute the price of a share of stock that pays a $1 per year dividend and that you expect to be able to sell in one year for $20, assuming you require a 15% return.

*4. After careful analysis, you have determined that a firm's dividends should grow at 7% on average in the foreseeable future. Its last dividend was $3. Compute the current price of this stock, assuming the required return is 18%.

5. Some economists think that the central banks should try to prick bubbles in the stock market before they get out of hand and cause later damage when they burst. How can monetary policy be used to prick a bubble? Explain how it can do this using the Gordon growth model.

*6. "Forecasters' predictions of inflation are notoriously inaccurate, so their expectations of inflation cannot be rational." Is this statement true, false, or uncertain? Explain your answer.

7. "Whenever it is snowing when Joe Commuter gets up in the morning, he misjudges how long it will take him to drive to work. Otherwise, his expectations of the driving time are perfectly accurate. Considering that it snows only once every ten years where Joe lives, Joe's expectations are almost always perfectly accurate." Are Joe's expectations rational? Why or why not?

*8. If a forecaster spends hours every day studying data to forecast interest rates but his expectations are not as accurate as predicting that tomorrow's interest rates will be identical to today's interest rate, are his expectations rational?

9. "If stock prices did not follow a random walk, there would be unexploited profit opportunities in the market." Is this statement true, false, or uncertain? Explain your answer.

*10. Suppose that increases in the money supply lead to a rise in stock prices. Does this mean that when you see that the money supply has had a sharp rise in the past week, you should go out and buy stocks? Why or why not?

11. If the public expects a corporation to lose $5 a share this quarter and it actually loses $4, which is still the largest loss in the history of the company, what does the efficient market hypothesis say will happen to the price of the stock when the $4 loss is announced?

*12. If I read in the *Wall Street Journal* that the "smart money" on Wall Street expects stock prices to fall, should I follow that lead and sell all my stocks?

13. If my broker has been right in her five previous buy and sell recommendations, should I continue listening to her advice?

*14. Can a person with rational expectations expect the price of IBM to rise by 10% in the next month?

15. "If most participants in the stock market do not follow what is happening to the monetary aggregates, prices of common stocks will not fully reflect information about them." Is this statement true, false, or uncertain? Explain your answer.

*16. "An efficient market is one in which no one ever profits from having better information than the rest." Is this statement true, false, or uncertain? Explain your answer.

17. If higher money growth is associated with higher future inflation and if announced money growth turns out to be extremely high but is still less than the market expected, what do you think would happen to long-term bond prices?

*18. "Foreign exchange rates, like stock prices, should follow a random walk." Is this statement true, false, or uncertain? Explain your answer.

19. Can we expect the value of the dollar to rise by 2% next week if our expectations are rational?

*20. "Human fear is the source of stock market crashes, so these crashes indicate that expectations in the stock market cannot be rational." Is this statement true, false, or uncertain? Explain your answer.

Web Exercises

1. Visit www.forecasts.org/data/index.htm. Click on "Stock Index" at the very top of the page. Now choose "U.S. Stock Indices-monthly." Review the indices for the DJIA, the S&P 500, and the NASDAQ composite. Which index appears most volatile? In which index would you have rather invested in 1985 if the investment had been allowed to compound until now?

2. The Internet is a great source of information on stock prices and stock price movements. There are many sites that provide up-to-the minute data on stock market indices. One of the best is found at http://finance.lycos.com/home/livecharts. This site provides free real-time streaming of stock market data. Click on the $indu to have the chart display the Dow Jones Industrial Average. Look at the stock trend over various intervals by adjusting the update frequency (click on "INT" at the top of the chart). Have stock prices been going up or down over the last day, week, month, and year?