CHAPTER 1

INTRODUCTION

We all dream of beating the market and being super investors and spend an inordinate amount of time and resources in this endeavor. Consequently, we are easy prey for the magic bullets and the secret formulae offered by eager salespeople pushing their wares. In spite of our best efforts, most of us fail in our attempts to be more than "average" investors. Nonetheless, we keep trying, hoping that we can be more like the investing legends – another Warren Buffett or Peter Lynch. We read the words written by and about successful investors, hoping to find in them the key to their stock-picking abilities, so that we can replicate them and become wealthy quickly.

In our search, though, we are whipsawed by contradictions and anomalies. In one corner of the investment townsquare, stands one advisor, yelling to us to buy businesses with solid cash flows and liquid assets because that's what worked for Buffett. In another corner, another investment expert cautions us that this approach worked only in the old world, and that in the new world of technology, we have to bet on companies with solid growth prospects. In yet another corner, stands a silver tongued salesperson with vivid charts and presents you with evidence of his capacity to get you in and out of markets at exactly the right times. It is not surprising that facing this cacophony of claims and counterclaims that we end up more confused than ever.

In this chapter, we present the argument that to be successful with any investment strategy, you have to begin with an investment philosophy that is consistent at its core and which matches not only the markets you choose to invest in but your individual characteristics. In other words, the key to success in investing may lie not in knowing what makes Peter Lynch successful but in finding out more about yourself.

What is an investment philosophy?

An investment philosophy is a coherent way of thinking about markets, how they work (and sometimes do not) and the types of mistakes that you believe consistently underlie investor behavior. Why do we need to make assumptions about investor mistakes? As we will argue, most investment strategies are designed to take advantage of errors made by some or all investors in pricing stocks. Those mistakes themselves are driven by far more basic assumptions about human behavior. To provide an illustration, the rational or irrational tendency of human beings to join crowds can result in price momentum – stocks that have gone up the most in the recent past are more likely to go up in the near future. Let us consider, therefore, the ingredients of an investment philosophy.

Human Frailty

Underlying all investment philosophies is a view about human behavior. In fact, one weakness of conventional finance and valuation has been the short shrift given to human behavior. It is not that we (in conventional finance) assume that all investors are rational, but that we assume that irrationalities are random and cancel out. Thus, for every investor who tends to follow the crowd too much (a momentum investor), we assume an investor who goes in the opposite direction (a contrarian), and that their push and pull in prices will ultimately result in a rational price. While this may, in fact, be a reasonable assumption for the very long term, it may not be a realistic one for the short term.

Academics and practitioners in finance who have long viewed the rational investor assumption with skepticism have developed a new branch of finance called behavioral finance which draws on psychology, sociology and finance to try to explain both why investors behave the way they do and the consequences for investment strategies. As we go through this book, examining different investment philosophies, we will try at the outset of each philosophy to explore the assumptions about human behavior that represent its base.

Market Efficiency

A closely related second ingredient of an investment philosophy is the view of market efficiency or its absence that you need for the philosophy to be a successful one. While all active investment philosophies make the assumption that markets are inefficient, they differ in their views on what parts of the market the inefficiencies are most likely to show up and how long they will last. Some investment philosophies assume that markets are correct most of the time but that they overreact when new and large pieces of information are released about individual firms – they go up too much on good news and down too much on bad news. Other investment strategies are founded on the belief that markets can make mistakes in the aggregate – the entire market can be under or overvalued – and that some investors (mutual fund managers, for example) are more likely to make these mistakes than others. Still other investment strategies may be based on the assumption that while markets do a good job of pricing stocks where there is a substantial amount of information – financial statements, analyst reports and financial press coverage –they systematically misprice stocks on which such information is not available.

Tactics and Strategies

Once you have an investment philosophy in place, you develop investment strategies that build on the core philosophy. Consider, for instance, the views on market efficiency expounded in the last section. The first investor, who believes that markets over react to news, may develop a strategy of buying stocks after large negative earnings surprises (where the announced earnings come in well below expectations) and selling stocks after positive earnings surprises. The second investor who believes that markets make mistakes in the aggregate may look at technical indicators (such as mutual fund cash positions and short sales ratios) to find out whether the market is over bought or over sold and take a contrary position. The third investor who believes that market mistakes are more likely when information is absent may look for stocks that are not followed by analysts or owned by institutional investors.

It is worth noting that the same investment philosophy can spawn multiple investment strategies. Thus, a belief that investors consistently overestimate the value of growth and under estimate the value of existing assets can manifest itself in a number of different strategies ranging from a passive one of buying low PE ratio stocks to a more active one of buying such companies and attempting to liquidate them for their assets. In other words, the number of investment strategies will vastly outnumber the number of investment philosophies.

Why do you need an investment philosophy?

Most investors have no investment philosophy, and the same can be said about many money managers and professional investment advisors. They adopt investment strategies that seem to work (for other investors) and abandon them when they do not. Why, if this is possible, you might ask, do you need an investment philosophy? The answer is simple. In the absence of an investment philosophy, you will tend to shift from strategy to strategy simply based upon a strong sales pitch from a proponent or perceived recent success. There are two negative consequences for your portfolio:

a. Lacking a rudder or a core set of beliefs, you will be easy prey for charlatans and pretenders, with each one claiming to have found the magic strategy that beats the market.
b. As you switch from strategy to strategy, you will have to change your portfolio, resulting in high transactions costs and you will pay more in taxes.
c. While there may be strategies that do work for some investors, they may not be appropriate for you, given your objectives, risk aversion and personal characteristics. In addition to having a portfolio that under performs the market, you are likely to find yourself with an ulcer or worse.

With a strong sense of core beliefs, you will have far more control over your destiny. Not only will you be able to reject strategies that do not fit your core beliefs about markets but also to tailor investment strategies to your needs. In addition, you will be able to get much more of a big picture view of what it is that is truly different across strategies and what they have in common.

The Big Picture of Investing

To see where the different investment philosophies fit into investing, let us begin by looking at the process of creating an investment portfolio. Note that this is a process that we all follow – amateur as well as professional investors - though it may be simpler for an individual constructing his or her own portfolio than it is for a pension fund manager with a varied and demanding clientele.

Step 1: Understanding the Client

The process always starts with the investor and understanding his or her needs and preferences. For a portfolio manager, the investor is a client, and the first and often most significant part of the investment process is understanding the client's needs, the client's tax status and most importantly, his or her risk preferences. For an individual investor constructing his or her own portfolio, this may seem simpler, but understanding one's own needs and preferences is just as important a first step as it is for the portfolio manager.

Step 2: Portfolio Construction

The next part of the process is the actual construction of the portfolio, which we divide into three sub-parts.

- The first of these is the decision on how to allocate the portfolio across different asset classes defined broadly as equities, fixed income securities and real assets (such as real estate, commodities and other assets). This asset allocation decision can also be framed in terms of investments in domestic assets versus foreign assets, and the factors driving this decision.
- The second component is the asset selection decision, where individual assets are picked within each asset class to make up the portfolio. In practical terms, this is the step where the stocks that make up the equity component, the bonds that make up the fixed income component and the real assets that make up the real asset component are selected.
- The final component is execution, where the portfolio is actually put together. Here investors must weigh the costs of trading against their perceived needs to trade quickly. While the importance of execution will vary across investment strategies, there are many investors who fail at this stage in the process.

Step 3: Evaluate portfolio performance

The final part of the process, and often the most painful one for professional money managers, is performance evaluation. Investing is after all focused on one objective and one objective alone, which is to make the most money you can, given your particular risk preferences. Investors are not forgiving of failure and unwilling to accept even the best of excuses, and loyalty to money managers is not a commonly found trait. By the same token, performance evaluation is just as important to the individual investor who constructs his or her own portfolio, since the feedback from it should largely determine how that investor approaches investing in the future.

These parts of the process are summarized in Figure 1.1, and we will return to this figure to emphasize the steps in the process as we consider different investment philosophies. As you will see, while all investment philosophies may have the same end objective of beating the market, each philosophy will emphasize a different component of the overall process and require different skills for success.

Figure 1.1: The Investment Process

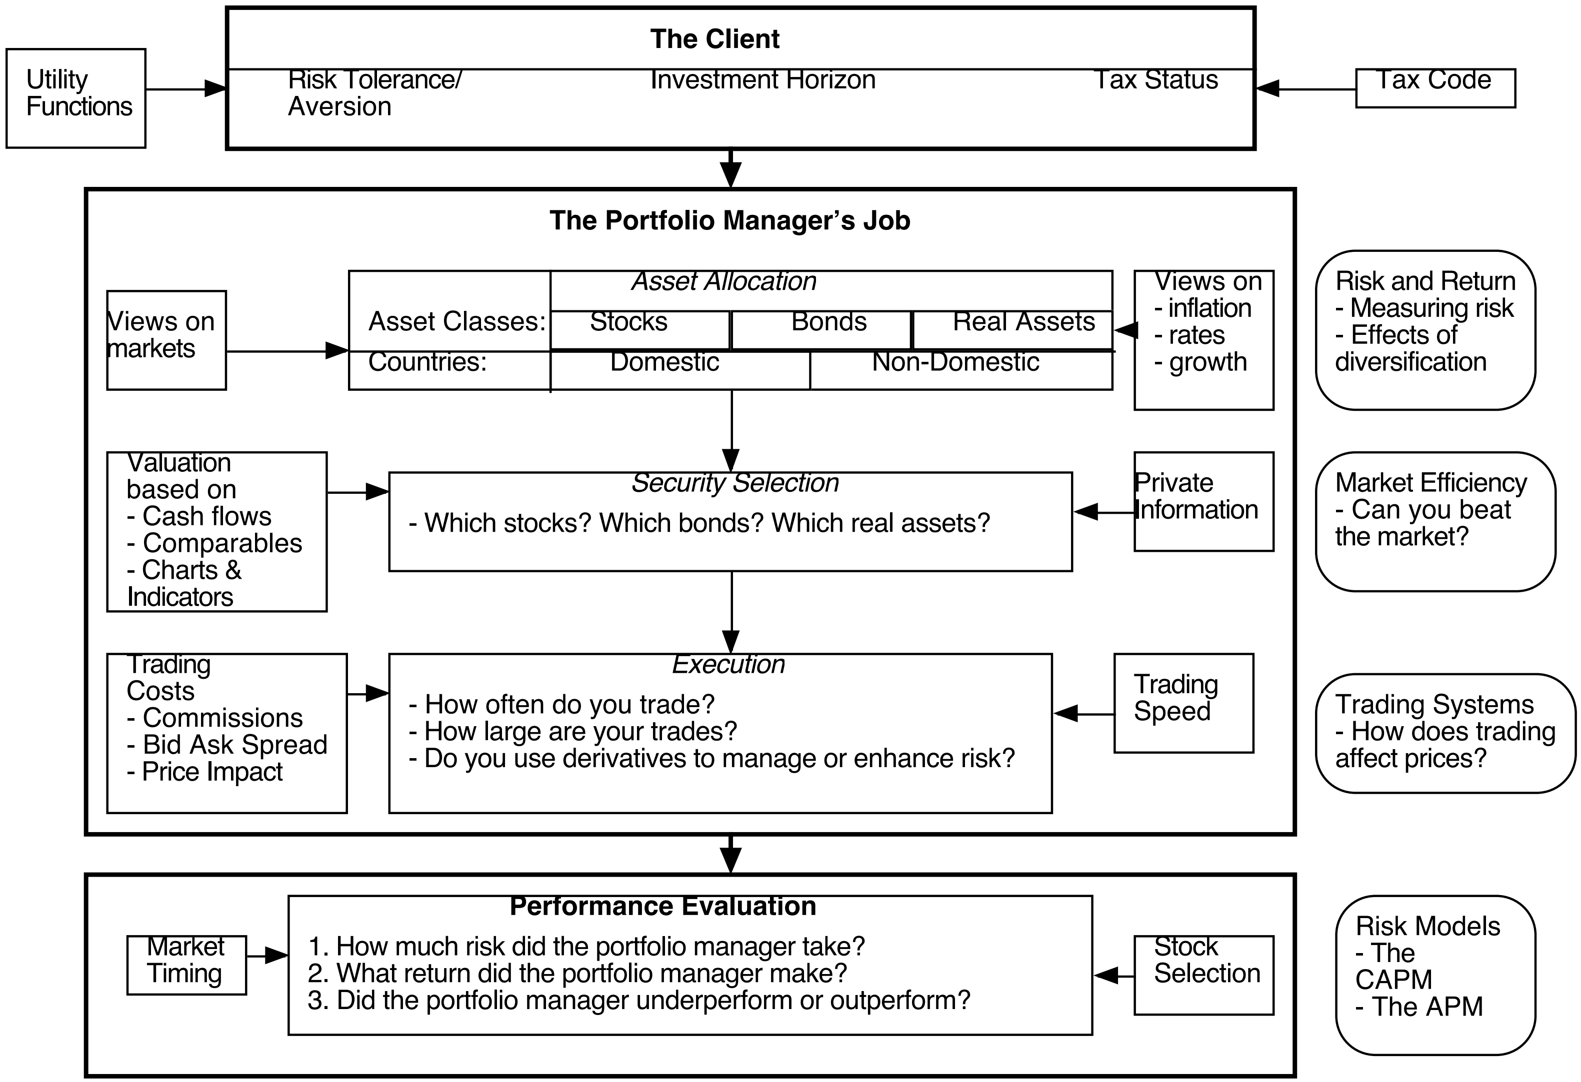

Categorizing Investment Philosophies

We will present the range of investment philosophies in this section, using the investment process to illustrate each philosophy. While we will leave much of the detail for later chapters, we will attempt to present at least the core of each philosophy here.

Market Timing versus Asset Selection

The broadest categorization of investment philosophies is on whether they are based upon timing overall markets or finding individual assets that are mispriced. The first set of philosophies can be categorized as *market timing* philosophies, while the second can be viewed as *security selection* philosophies.

Within each, though, are numerous strands that take very different views about markets. Consider market timing first. While most of us consider market timing only in the context of the stock market, there are investors who consider market timing to include a much broader range of markets – currency markets, bond markets and real estate come to mind. The range of choices among security selection philosophies is even wider and can span charting and technical indicators, fundamentals (earnings, cashflows or growth) and information (earnings reports, acquisition announcements).

While market timing has allure to all of us (because it pays off so well when you are right), it is difficult to succeed at for exactly that reason. There are all too often too many investors attempting to time markets, and succeeding consistently is very difficult to do. If you decide to pick stocks, how do you choose whether you pick them based upon charts, fundamentals or growth potential? The answer, as we will see, in the next section will depend not only on your views of the market and empirical evidence but also on your personal characteristics.

Activist versus Passive Investing

At the broadest level, investment philosophies can also be categorized as active or passive strategies. In a *passive strategy*, you invest in a stock or company and wait for your investment to pay off. Assuming that your strategy is successful, this will come from the market recognizing and correcting a misvaluation. Thus, a portfolio manager who buys stocks with low price earnings ratios and stable earnings is following a passive strategy. So is an index fund manager, who essentially buys all stocks in the index. In an *activist strategy*, you invest in a company and then try to change the way the company is run to make it more valuable. Venture capitalists can be categorized as activist investors since they not only take positions in promising companies but they also provide significant inputs into how these firms are run. In recent years, we have seen investors like Michael Price and the

California State pension fund (Calpers) bring this activist philosophy to publicly traded companies, using the clout of large positions to change the way companies are run. We should hasten to draw a contrast between activist investing and active investing. Any investor who tries to beat the market by picking stocks is viewed as an active investor. Thus, active investors can adopt passive strategies or activist strategies.

Time Horizon

Different investment philosophies require different time horizons. A philosophy based upon the assumption that markets overreact to new information may generate short term strategies. For instance, you may buy stocks right after a bad earnings announcement, hold a few weeks and sell (hopefully at a higher price, as the market corrects its over reaction). In contrast, a philosophy of buying neglected companies (stocks that are not followed by analysts or held by institutional investors) may require much longer time horizons.

One factor that will determine the time horizon of an investment philosophy is the nature of the adjustment that has to occur for you to reap the rewards of a successful strategy. Passive value investors who buy stocks in companies that they believe are under valued may have to wait years for the market correction to occur, even if they are right. Investors who trade ahead or after earnings reports, because they believe that markets do not respond correctly to such reports, may hold the stock for only a few days. At the extreme, investors who see the same (or very similar) assets being priced differently in two markets may buy the cheaper one and sell the more expensive one, locking in "arbitrage" profits in a few minutes.

Coexistence of Contradictory Strategies

One of the most fascinating aspects of investment philosophy is the coexistence of investment philosophies based upon contradictory views of the markets. Thus, you can have market timers who trade on *price momentum* (suggesting that investors are slow to learn from information) and market timers who are *contrarians* (which is based on the belief that markets over react). Among security selectors who use fundamentals, you can have *value investors* who buy value stocks, because they believe markets overprice growth, and *growth investors* who buy growth stocks using exactly the opposite justification. The coexistence of these contradictory impulses for investing may strike some as irrational, but it is healthy and may actually be responsible for keeping the market in balance. In addition, you can have investors with contradictory philosophies co-existing in the market because of their different time horizons, views on risk and tax status. For instance, tax exempt investors may find stocks that pay large dividends a bargain, while taxable investors may reject these same stocks because dividends are taxed at the ordinary tax rate.

Investment Philosophies in Context

We can consider the differences between investment philosophies in the context of the investment process, described in figure 1.1. Market timing strategies primarily affect the asset allocation decision. Thus, investors who believe that stocks are under valued will invest more of their portfolios in stocks than would be justified given their risk preferences. Security selection strategies in all their forms – technical analysis, fundamentals or private information – all center on the security selection component of the portfolio management process. You could argue that strategies that are not based upon grand visions of market efficiency but are designed to take advantage of momentary mispricing of assets in markets (such as arbitrage) revolve around the execution segment of portfolio management. It is not surprising that the success of such opportunistic strategies depend upon trading quickly to take advantage of pricing errors, and keeping transactions costs low. Figure 1.2 presents the different investment philosophies.

Figure 1.2: Investment Philosophies

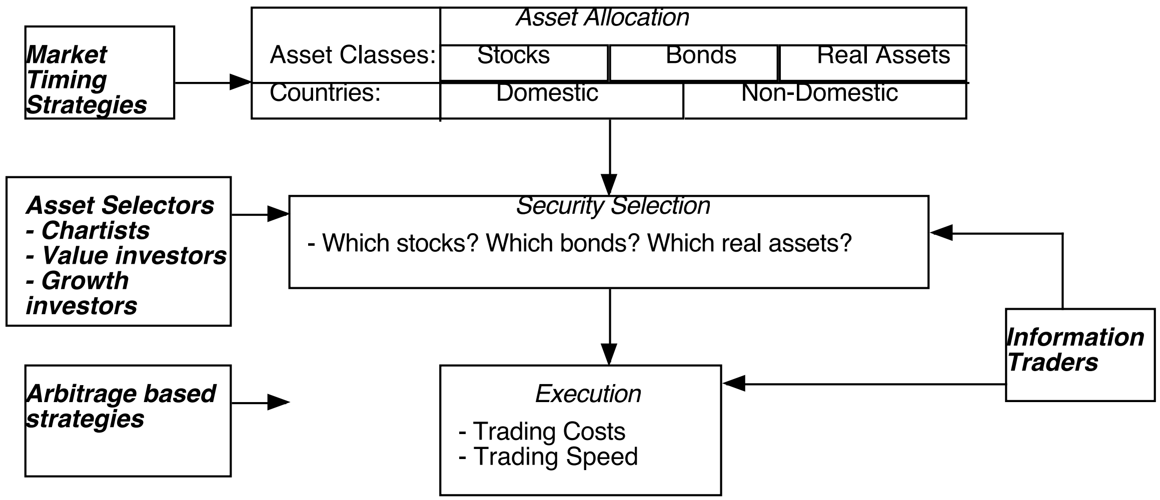

Developing an Investment Philosophy: The Step

If every investor needs an investment philosophy, what is the process that you go through to come up with such a philosophy? While this entire book is about the process, we can lay out the three steps involved in this section.

Step 1: Understand the fundamentals of risk and valuation

Before you embark on the journey of finding an investment philosophy, you need to get your financial toolkit ready. At the minimum, you should understand

- how to measure the risk in an investment and relate it to expected returns.
- how to value an asset, whether it be a bond, stock or a business
- the ingredients of trading costs, and the trade off between the speed of trading and the cost of trading

We would hasten to add that you do not need to be a mathematical wizard to do any of these and we will begin this book with a section dedicated to providing these basic tools.

Step 2: Develop a point of view about how markets work and where they might break down

Every investment philosophy is grounded in a point of view about human behavior (and irrationality). While personal experience often determines how we view our fellow human beings, we should expand this to consider broader evidence from markets on how investors act before we make our final judgments.

Over the last few decades, it has become easy to test different investment strategies as data becomes more accessible. There now exists a substantial body of research on the investment strategies that have beaten the market over time. For instance, researchers have found convincing evidence that stocks with low price to book value ratios have earned significantly higher returns than stocks of equivalent risk but higher price to book value ratios. It would be foolhardy not to review this evidence in the process of developing your investment philosophy. At the same time, though, you should keep in mind three caveats about this research:

- Since they are based upon the past, they represent a look in the rearview mirror. Strategies that earned substantial returns in the 1990s may no longer be viable strategies now. In fact, as successful strategies get publicized either directly (in books and articles) or indirectly (by portfolio managers trading on them), you should expect to see them become less effective.
- Much of the research is based upon constructing hypothetical portfolios, where you buy and sell stocks at historical prices and little or no attention is paid to

transactions costs. To the extent that trading can cause prices to move, the actual returns on strategies can be very different from the returns on the hypothetical portfolio.

- A test of an investment strategy is almost always a joint test of both the strategy and a model for risk. To see why, consider the evidence that stocks with low price to book value ratios earn higher returns than stocks with high price to book value ratios, with similar risk (at least as measured by the models we use). To the extent that we mismeasure risk or ignore a key component of risk, it is entirely possible that the higher returns are just a reward for the greater risk associated with low price to book value stocks.

Since understanding whether a strategy beats the market is such a critical component of investing, we will consider the approaches that are used to test a strategy, some basic rules that need to be followed in doing these tests and common errors that are made (unintentionally or intentionally) when running such tests. As we look at each investment philosophy, we will review the evidence that is available on strategies that emerge from that philosophy.

Step 3: Find the philosophy that provides the best fit for you

Once you understand the basics of investing, form your views on human foibles and behavior and review the evidence accumulated on each of the different investment philosophies, you are ready to make your choice. In our view, there is potential for success with almost every investment philosophy (yes, even charting) but the prerequisites for success can vary. In particular, success may rest on:

- *Your risk aversion*: Some strategies are inherently riskier than others. For instance, venture capital or private equity investing, where you invest your funds in small, private businesses that show promise is inherently more risky than buying value stocks – equity in large, stable, publicly traded companies. The returns are also likely to be higher. However, more risk averse investors should avoid the first strategy and focus on the second. Picking an investment philosophy (and strategy) that requires you to take on more risk than you feel comfortable taking can be hazardous to your health and your portfolio.
- *The size of your portfolio*: Some strategies require larger portfolios for success whereas others work only on a smaller scale. For instance, it is very difficult to be an activist value investor if you have only $ 100,000 in your portfolio, since firms are unlikely to listen to your complaints. On the other hand, a portfolio manager with $ 100 billion to invest may not be able to adopt a strategy that requires buying small,

neglected companies. With such a large portfolio, she would very quickly end up becoming the dominant stockholder in each of the companies and affecting the price every time she trade.

- *Your time horizon*: Some investment philosophies are predicated on a long time horizon, whereas others require much shorter time horizons. If you are investing your own funds, your time horizon is determined by your personal characteristics – some of us are more patient than others – and your needs for cash – the greater the need for liquidity, the shorter your time horizon has to be. If you are a professional (an investment adviser or portfolio manager), managing the funds of others, it is your clients time horizon and cash needs that will drive your choice of investment philosophies and strategies.
- *Your tax status*: Since such a significant portion of your money ends up going to the tax collectors, they have a strong influence on your investment strategies and perhaps even the investment philosophy you adopt. In some cases, you may have to abandon strategies that you find attractive on a pre-tax basis because of the tax bite that they expose you to.

Thus, the right investment philosophy for you will reflect your particular strengths and weaknesses. It should come as no surprise, then, that investment philosophies that work for some investors do not work for others. Consequently, there can be no one investment philosophy that can be labeled "best" for all investors.

Conclusion

An investment philosophy represents a set of core beliefs about how investors behave and markets work. To be a successful investor, you not only have to consider the evidence from markets but you also have to examine your own strengths and weaknesses to come up with an investment philosophy that best fits you. Investors without core beliefs tend to wander from strategy to strategy, drawn by the anecdotal evidence or recent success, creating transactions costs and incurring losses as a consequence. Investors with clearly defined investment philosophies tend to be more consistent and disciplined in their investment choices.

In this chapter, we considered a broad range of investment philosophies from market timing to arbitrage and placed each of them in the broad framework of portfolio management. We also examined the three steps in the path to an investment philosophy, beginning with the understanding of the tools of investing – risk, trading costs and valuation – continuing with an evaluation of the empirical evidence on whether, when and how markets break down and concluding with a self-assessment, to find the investment philosophy that best matches your time horizon, risk preferences and portfolio characteristics.

CHAPTER 2

UPSIDE, DOWNSIDE: UNDERSTANDING RISK

Risk is part of investing and understanding what it is and how it is measured is essential to developing an investment philosophy. In this chapter, we will lay the foundations for analyzing risk in investments. We present alternative models for measuring risk and converting these risk measures into an expected return. We will also consider ways in an investor can measure his or her risk aversion.

We begin with a discussion of risk and present our analysis in three steps. In the first step, we define risk in terms of uncertainty about future returns. The greater this uncertainty, the more risky an investment is perceived to be. The next step, which we believe is the central one, is to decompose this risk into risk that can be diversified away by investors and risk that cannot. In the third step, we look at how different risk and return models in finance attempt to measure this non-diversifiable risk. We compare and contrast the most widely used model, the capital asset pricing model, with other models, and explain how and why they diverge in their measures of risk and the implications for the equity risk premium. In the second part of this chapter, we consider default risk and how it is measured by ratings agencies. In addition, we discuss the determinants of the default spread and why it might change over time.

What is risk?

Risk, for most of us, refers to the likelihood that in life's games of chance, we will receive an outcome that we will not like. For instance, the risk of driving a car too fast is getting a speeding ticket, or worse still, getting into an accident. Webster's dictionary, in fact, defines risk as "exposing to danger or hazard". Thus, risk is perceived almost entirely in negative terms.

In finance, our definition of risk is both different and broader. Risk, as we see it, refers to the likelihood that we will receive a return on an investment that is different from the return we expected to make. Thus, risk includes not only the bad outcomes, i.e, returns that are lower than expected, but also good outcomes, i.e., returns that are higher than expected. In fact, we can refer to the former as downside risk and the latter is upside risk; but we consider both when measuring risk. In fact, the spirit of our definition of risk in finance is captured best by the Chinese symbols for risk, which are reproduced below:

危機

The first symbol is the symbol for "danger", while the second is the symbol for "opportunity", making risk a mix of danger and opportunity. It illustrates very clearly the tradeoff that every investor and business has to make – between the higher rewards that come with the opportunity and the higher risk that has to be borne as a consequence of the danger.

Much of this chapter can be viewed as an attempt to come up with a model that best measures the "danger" in any investment and then attempts to convert this into the "opportunity" that we would need to compensate for the danger. In financial terms, we term the danger to be "risk" and the opportunity to be "expected return".

Equity Risk and Expected Return

To demonstrate how risk is viewed in finance, we will present risk analysis in three steps. First, we will define risk . Second, we will differentiate between risk that is specific to one or a few investments and risk that affects a much wider cross section of investments. We will argue that in a market where investors are well diversified, it is only the latter risk, called market ris*k* that will be rewarded. Third, we will look at alternative models for measuring this market risk and the expected returns that go with it.

I. Defining Risk

Investors who buy assets expect to earn returns over the time horizon that they hold the asset. Their actual returns over this holding period may be very different from the expected returns and it is this difference between actual and expected returns that is source of risk. For example, assume that you are an investor with a 1-year time horizon buying a 1-year Treasury bill (or any other default-free one-year bond) with a 5% expected return. At the end of the 1-year holding period, the actual return on this investment will be 5%, which is equal to the expected return. The return distribution for this investment is shown in Figure 2.1.

Figure 2.1: Probability Distribution for Riskfree Investment

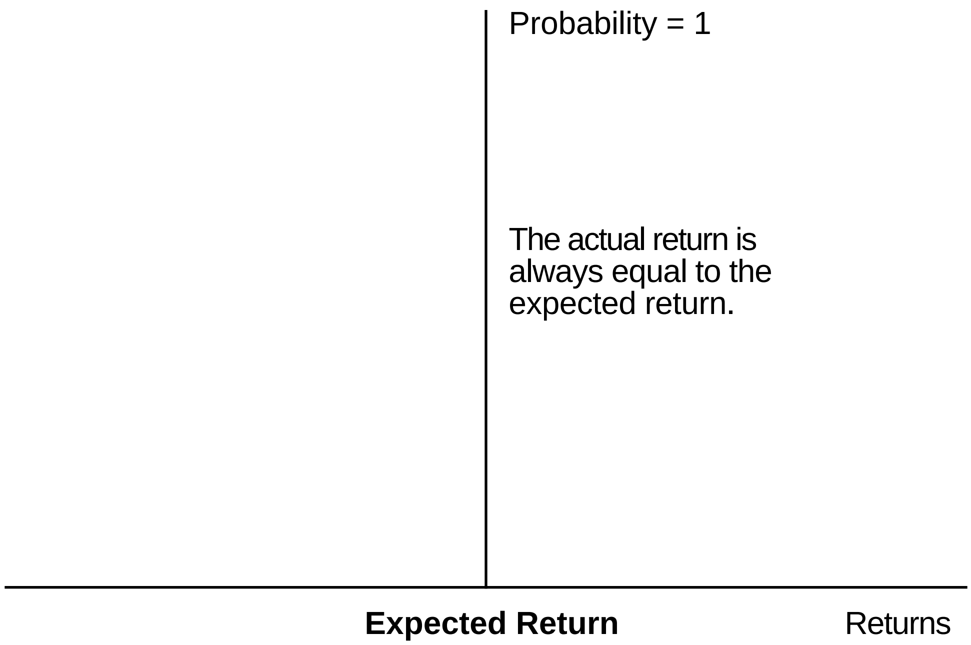

This is a riskless investment.

To provide a contrast to the riskless investment, consider an investor who buys stock in a company like Cisco. This investor, having done her research, may conclude that she can make an expected return of 30% on Cisco over her 1-year holding period. The actual return over this period will almost certainly not be equal to 30%; it might be much greater or much lower. The distribution of returns on this investment is illustrated in Figure 2.2.

Figure 2.2: Probability Distribution for Risky Investment

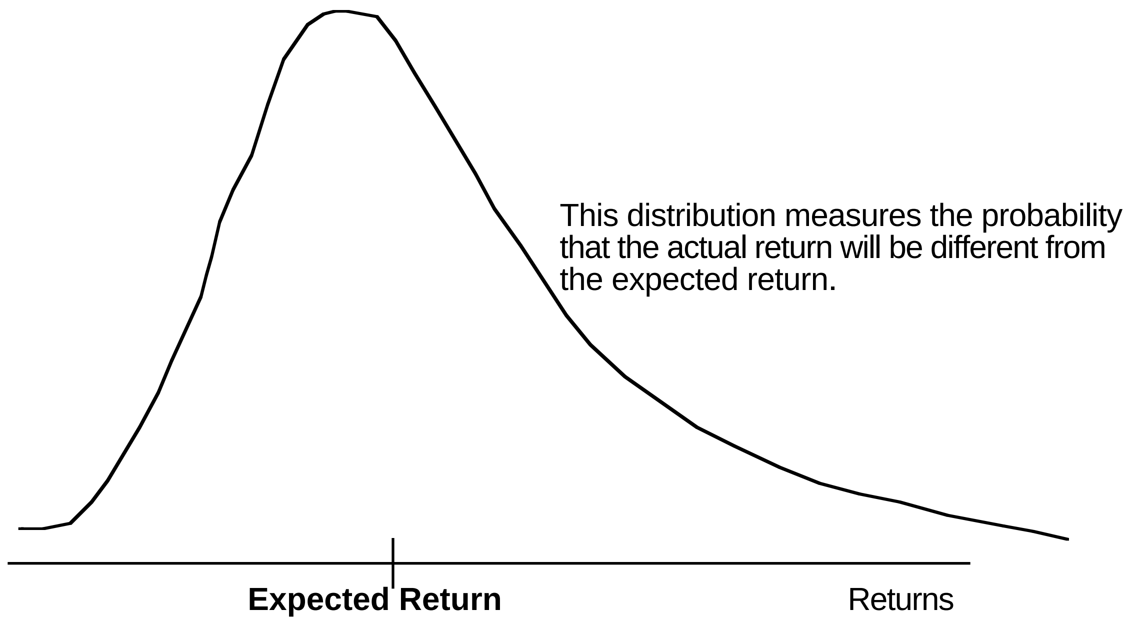

In addition to the expected return, an investor has to note that the actual returns, in this case, are different from the expected return. The spread of the actual returns around the expected return is measured by the variance or standard deviation of the distribution; the greater the deviation of the actual returns from expected returns, the greater the variance.

One of the limitations of variance is that it considers all variation from the expected return to be risk. Thus, the potential that you will earn a 60% return on Cisco (30% more than the expected return of 30%) affects the variance exactly as much as the potential that you will earn 0% (30% less than the expected return). In other words, you do not distinguish between downside and upside risk. This is justified by arguing that risk is symmetric – upside risk must inevitably create the potential for downside risk.[1] If you are bothered by this assumption, you could compute a modified version of the variance, called the *semi-variance*, where you consider only the returns that fall below the expected return.

The Most and Least Volatile Stocks: Take a look at the 50 most and 50 least volatile stocks traded in the United States, based upon 5 years of weekly data.

It is true that measuring risk with variance or semi-variance can provide too limited a view of risk, and there are some investors who use simpler stand-ins (proxies) for risk. For instance, you may consider stocks in some sectors (such as technology) to be riskier than

---

[1] In statistical terms, this is the equivalent of assuming that the distribution of returns is close to normal.

stocks in other sectors (say, food processing). Others prefer ranking or categorization systems, where you put firms into risk classes, rather than trying to measure its risk in units. Thus, Value Line ranks firms into five classes, based upon risk.

There is one final point that needs to be made about how variances and semi-variances are estimated for most stocks. Analysts usually look at the past – stock prices over the last 2 or 5 years- to make these estimates. This may be appropriate for firms that have not changed their fundamental characteristics – business or leverage – over the period. For firms that have changed significantly over time, variances from the past may provide a very misleading view of betas in the future.

II. Diversifiable and Non-diversifiable Risk

Although there are many reasons that actual returns may differ from expected returns, we can group the reasons into two categories: firm-specific and market-wide. The risks that arise from firm-specific actions affect one or a few investments, while the risk arising from market-wide reasons affect many or all investments. This distinction is critical to the way we assess risk in finance.

The Components of Risk

When an investor buys stock or takes an equity position in a firm, he or she is exposed to many risks. Some risk may affect only one or a few firms and it is this risk that we categorize as *firm-specific risk*. Within this category, we would consider a wide range of risks, starting with the risk that a firm may have misjudged the demand for a product from its customers; we call this *project risk*. For instance, consider an investment by Boeing in a new larger capacity airplane that we will call the Super Jumbo. This investment is based on the assumption that airlines want a larger airplane and will be willing to pay a higher price for it. If Boeing has misjudged this demand, it will clearly have an impact on Boeing's earnings and value, but it should not have a significant effect on other firms in the market. The risk could also arise from competitors proving to be stronger or weaker than anticipated; we call this *competitive risk*. For instance, assume that Boeing and Airbus are competing for an order from Qantas, the Australian airline. The possibility that Airbus may win the bid is a potential source of risk to Boeing and perhaps a few of its suppliers. But again, only a handful of firms in the market will be affected by it. Similarly, the Home Depot recently launched an online store to sell its home improvement products. Whether it succeeds or not is clearly important to the Home Depot and its competitors, but it is unlikely to have an impact on the rest of the market. In fact, we would extend our risk measures to include risks that may affect an entire sector but are restricted to that sector; we call this *sector risk*. For instance, a cut in the defense budget in the United States will adversely affect all firms in the defense business, including Boeing, but there should be no significant impact on other sectors, such as food and apparel. What is common across the three risks described above – project, competitive and sector risk – is that they affect only a small sub-set of firms.

There is other risk that is much more pervasive and affects many if not all investments. For instance, when interest rates increase, all investments are negatively affected, albeit to different degrees. Similarly, when the economy weakens, all firms feel the effects, though cyclical firms (such as automobiles, steel and housing) may feel it more. We term this risk *market risk*.

Finally, there are risks that fall in a gray area, depending upon how many assets they affect. For instance, when the dollar strengthens against other currencies, it has a significant impact on the earnings and values of firms with international operations. If most firms in the market have significant international operations, it could well be categorized as market risk. If only a few do, it would be closer to firm-specific risk. Figure 2.3 summarizes the break down or the spectrum of firm specific and market risks.

Figure 2.3: A Break Down of Risk

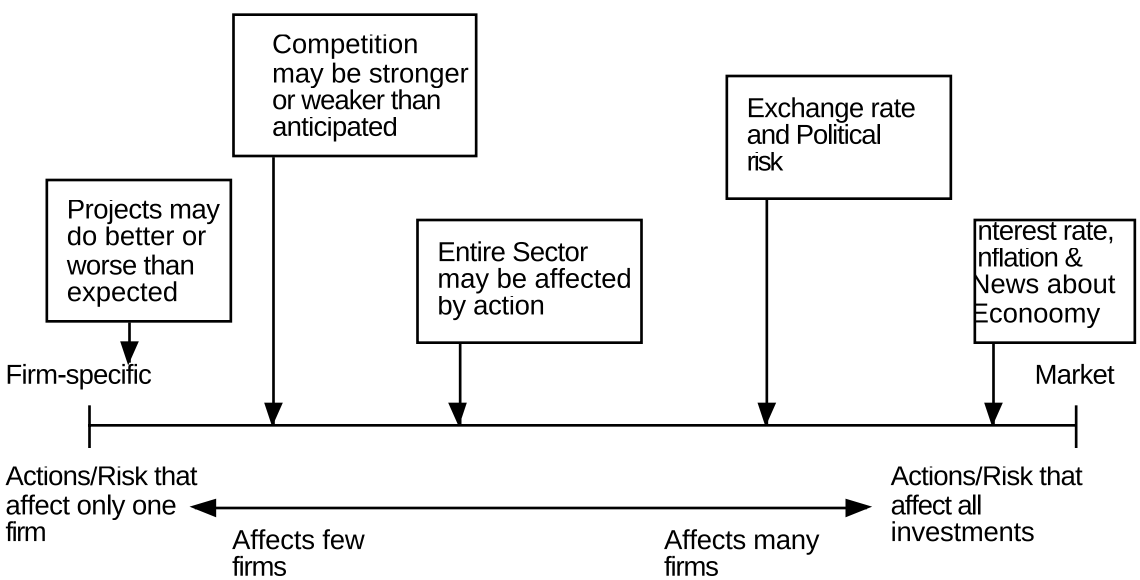

Why Diversification reduces or eliminates Firm-specific Risk: An Intuitive Explanation

As an investor, you could invest your entire portfolio in one stock, say Boeing. If you do so, you are exposed to both firm specific and market risk. If, however, you expand your portfolio to include other assets or stocks, you are diversifying, and by doing so, you can reduce your exposure to firm-specific risk. There are two reasons why diversification reduces or, at the limit, eliminates firm specific risk. The first is that each investment in a diversified portfolio is a much smaller percentage of that portfolio than would be the case if you were not diversified. Thus, any action that increases or decreases the value of only that investment or a small group of investments will have only a small impact on your overall portfolio, whereas undiversified investors are much more exposed to changes in the values of the investments in their portfolios. The second reason is that the effects of firm-specific actions on the prices of individual assets in a portfolio can be either positive or negative for each asset in any period. Thus, in very large portfolios, this risk will average out to zero and will not affect the overall value of the portfolio.

In contrast, the effects of market-wide movements are likely to be in the same direction for most or all investments in a portfolio, though some assets may be affected more than others. For instance, other things being equal, an increase in interest rates will lower the values of most assets in a portfolio. Being more diversified does not eliminate this risk.

Highest R-squared companies: Take a look at the 50 companies with the highest proportion of market risk using the last 5 years or weekly data.

One of the simplest ways of measuring how much risk in a firm is firm specific is to look at the proportion of the firm's price movements that are explained by the market. This is called the *R-squared* and it should range between zero and one can be stated as a percentage; it measures the proportion of the firm's risk that comes from the market. A firm with an R-squared of zero has 100% firm specific risk whereas a firm with an R-squared of 0% has no firm specific risk.

*Why is the marginal investor assumed to be diversified***?**

The argument that diversification reduces an investor's exposure to risk is clear both intuitively and statistically, but risk and return models in finance go further. The models look at risk through the eyes of the investor most likely to be trading on the investment at any point in time, i.e. the marginal investor. They argue that this investor, who sets prices for investments, is well diversified; thus, the only risk that he or she cares about is the risk added on to a diversified portfolio or market risk. This argument can be justified simply. The risk in an investment will always be perceived to be higher for an undiversified investor than for a diversified one, since the latter does not shoulder any firm-specific risk and the former does. If both investors have the same expectations about future earnings and cash flows on an asset, the diversified investor will be willing to pay a higher price for that asset because of his or her perception of lower risk. Consequently, the asset, over time, will end up being held by diversified investors.

This argument is powerful, especially in markets where assets can be traded easily and at low cost. Thus, it works well for a stock traded in the United States, since investors can become diversified at fairly low cost. In addition, a significant proportion of the trading in US stocks is done by institutional investors, who tend to be well diversified. It becomes a more difficult argument to sustain when assets cannot be easily traded, or the costs of trading are high. In these markets, the marginal investor may well be undiversified and firm-specific risk may therefore continue to matter when looking at individual investments. For instance, real estate in most countries is still held by investors who are undiversified and have the bulk of their wealth tied up in these investments.

III. Models Measuring Market Risk

While most risk and return models in use in finance agree on the first two steps of the risk analysis process, i.e., that risk comes from the distribution of actual returns around the expected return and that risk should be measured from the perspective of a marginal investor who is well diversified, they part ways when it comes to measuring non-diversifiable or market risk. In this section, we will discuss the different models that exist in finance for measuring market risk and why they differ. We will begin with what still is the standard model for measuring market risk in finance – the capital asset pricing model (CAPM) – and then discuss the alternatives to this model that have developed over the last two decades. While we will emphasize the differences, we will also look at what they have in common.

A. The Capital Asset Pricing Model (CAPM)

The risk and return model that has been in use the longest and is still the standard in most real world analyses is the capital asset pricing model (CAPM). In this section, we will examine the assumptions made by the model and the measures of market risk that emerge from these assumptions.

Assumptions

While diversification reduces the exposure of investors to firm specific risk, most investors limit their diversification to holding only a few assets. Even large mutual funds rarely hold more than a few hundred stocks and many of them hold as few as ten to twenty. There are two reasons why investors stop diversifying. One is that an investor or mutual fund manager can obtain most of the benefits of diversification from a relatively small portfolio, because the marginal benefits of diversification become smaller as the portfolio gets more diversified. Consequently, these benefits may not cover the marginal costs of diversification, which include transactions and monitoring costs. Another reason for limiting diversification is that many investors (and funds) believe they can find under valued assets and thus choose not to hold those assets that they believe to be fairly or over valued.

The capital asset pricing model assumes that there are no transactions costs and that all assets are traded. It also assumes that everyone has access to the same information and that investors therefore cannot find under or over valued assets in the market place. Making these assumptions allows investors to keep diversifying without additional cost. At the limit, each investor's will include every traded asset in the market held in proportion to its market value. The fact that this diversified portfolio includes all traded assets in the market is the reason it is called the *market portfolio*, which should not be a surprising result, given the benefits of diversification and the absence of transactions costs in the capital asset pricing model. If diversification reduces exposure to firm-specific risk and there are no costs associated with adding more assets to the portfolio, the logical limit to diversification is to hold a small proportion of every traded asset in the market. If this seems abstract, consider the market portfolio to be an extremely well diversified mutual fund that holds stocks and real assets, and treasury bills as the riskless asset. In the CAPM, all investors will hold combinations of treasury bills and the same mutual fund[2].

Investor Portfolios in the CAPM

If every investor in the market holds the identical market portfolio, how exactly do investors reflect their risk aversion in their investments? In the capital asset pricing model, investors adjust for their risk preferences in their allocation decision, where they decide how much to invest in a riskless asset and how much in the market portfolio. Investors who are risk averse might choose to put much or even all of their wealth in the riskless asset. Investors who want to take more risk will invest the bulk or even all of their wealth in the market portfolio. Investors, who invest all their wealth in the market portfolio and are still desirous of taking on more risk, would do so by borrowing at the riskless rate and investing more in the same market portfolio as everyone else.

These results are predicated on two additional assumptions. First, there exists a riskless asset, where the expected returns are known with certainty. Second, investors can lend and borrow at the same riskless rate to arrive at their optimal allocations. While lending at the riskless rate can be accomplished fairly simply by buying treasury bills or bonds, borrowing at the riskless rate might be more difficult to do for individuals. There are variations of the CAPM that allow these assumptions to be relaxed and still arrive at the conclusions that are consistent with the model.

---

[2] The significance of introducing the riskless asset into the choice mix, and the implications for portfolio choice were first noted in Sharpe (1964) and Lintner (1965). Hence, the model is sometimes called the Sharpe-Lintner model.

Measuring the Market Risk of an Individual Asset

The risk of any asset to an investor is the risk added by that asset to the investor's overall portfolio. In the CAPM world, where all investors hold the market portfolio, the risk to an investor of an individual asset will be the risk that this asset adds on to that portfolio. Intuitively, if an asset moves independently of the market portfolio, it will not add much risk to the market portfolio. In other words, most of the risk in this asset is firm-specific and can be diversified away. In contrast, if an asset tends to move up when the market portfolio moves up and down when it moves down, it will add risk to the market portfolio. This asset has more market risk and less firm-specific risk. Statistically, this added risk is measured by the *covariance* of the asset with the market portfolio.

> *Highest and Lowest Beta Stocks*: Take a look at the 50 highest beta and 50 lowest beta stocks traded in the United States, based upon 5 years of weekly data.

The covariance is a percentage value and it is difficult to pass judgment on the relative risk of an investment by looking at this value. In other words, knowing that the covariance of Boeing with the Market Portfolio is 55% does not provide us a clue as to whether Boeing is riskier or safer than the average asset. We therefore standardize the risk measure by dividing the covariance of each asset with the market portfolio by the variance of the market portfolio. This yields a risk measure called the **beta** of the asset:

$$\text{Beta of an asset } = \frac{\text{Covariance of asset with Market Portfolio}}{\text{Variance of the Market Portfolio}}$$

The beta of the market portfolio, and by extension, the average asset in it, is one. Assets that are riskier than average (using this measure of risk) will have betas that are greater than 1 and assets that are less risky than average will have betas that are less than 1. The riskless asset will have a beta of 0.

Getting Expected Returns

Once you accept the assumptions that lead to all investors holding the market portfolio and measure the risk of an asset with beta, the return you can expect to make can be written as a function of the risk-free rate and the beta of that asset.

Expected Return on an investment = Riskfree Rate + Beta (Risk Premium for buying the average risk investment)

Consider the three components that go into the expected return.

a. Riskless Rate: The return you can make on a riskfree investment becomes the base from which you build expected returns. Essentially, you are assuming that if you can make 5%

investing in treasury bills or bonds, you would not settle for less than this as an expected return for investing in a riskier asset. Generally speaking, we use the interest rate on government securities to estimate the riskfree rate, assuming that such securities have no default risk. While this may be a safe assumption in the United States and other developed markets, it may be inappropriate in many emerging markets, where governments themselves are viewed as capable of defaulting. In such cases, the government bond rate will include a premium for default risk and this premium will have to be removed to arrive at a riskfree rate. [3]

b. *The beta of the investment*: The beta is the only component in this model which varies from investment to investment, with investments that add more risk to the market portfolio having higher betas. But where do betas come from? Since the beta measures the risk added to a market portfolio by an individual stock, it is usually estimated by running a regression of past returns on the stock against returns on a market index.

Figure 2.4: Beta Estimate for Cisco: S&P 500

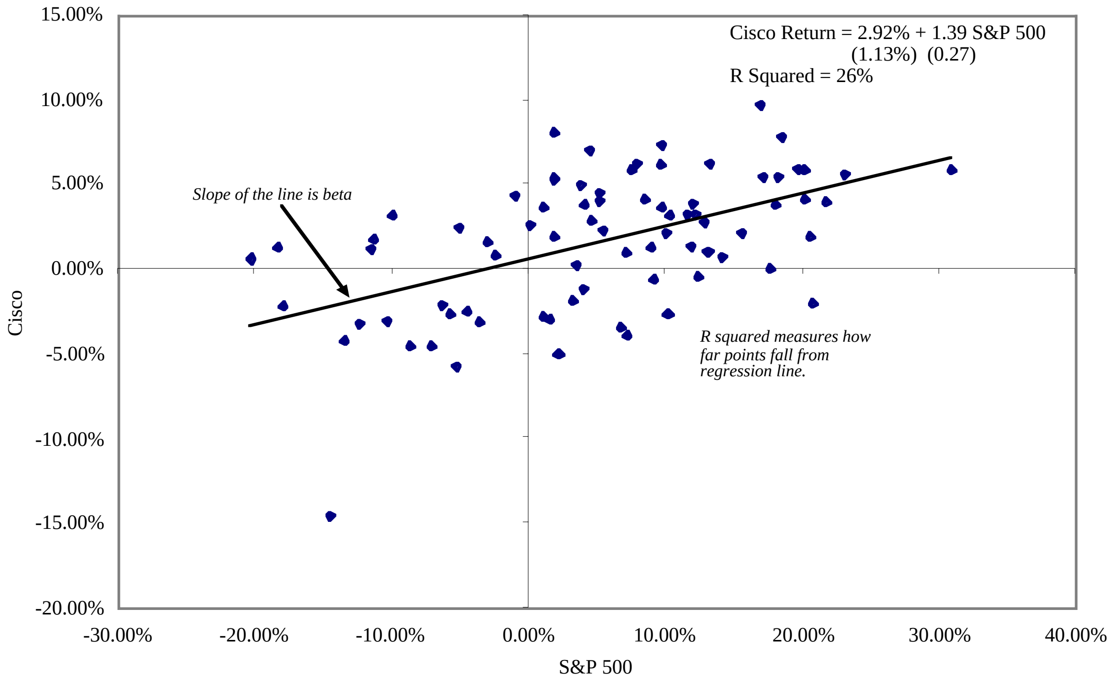

---

[3] Consider, for example, a government bond issued by the Brazilian government. Denominated in Brazilian Real, this bond has an interest rate of 17%. The Brazilian government is viewed as having default risk on this bond and is rated BB by Standard and Poor's. If we subtract the typical default spread earned by BB rated country bonds (about 5%) from 17%, we end up with a riskless rate in Brazilian Real of 12%.

The slope of the regression captures how sensitive a stock is to market movements and is the beta of the stock. In the regression above, for instance, the beta of Cisco would be 1.39. There are, however, two problems with regression betas. One is that the beta comes with estimation error – the standard error in the estimate is 0.27. Thus, the true beta for Cisco could be anywhere from .85 to 1.93 – this range is estimated by adding and subtracting two standard errors to the beta estimate. The other is that firms change over time and we are looking backwards rather than looking forwards. A better way to estimate betas is to look at the average beta for publicly traded firms in the business or businesses Cisco operates in. While these betas come from regressions as well, the average beta is always more precise than any one firm's beta estimate.

Risk Premium for the United States: Take a look at the equity risk premium implied in the U.S. stock market from 1960 through the most recent year.

c. *The risk premium for buying the average risk investment*: You can view this as the premium you would demand for investing in equities as a class as opposed to the riskless investment. Thus, if you require a return of 9% for investing in equities and the treasury bond rate is 5%, your risk premium is 4%. There are again two ways in which you can estimate this risk premium. One is to look at the past and look at the typical premium you would have earned investing in stocks as opposed to a riskless investment. This number is called a historical premium and yields about 5-7% for the United States. The other is to look at how stocks are priced today and to estimate the premium that investors must be demanding. This is called an implied premium and yields a value of about 4% for U.S. stocks in early 2002.

Bringing it all together, you could use the capital asset pricing model to estimate the expected return on a stock for Cisco for the future (assuming a treasury bond rate of 5%, the regression beta of 1.39 and a risk premium of 4%):

Expected return on Cisco = T. Bond Rate + Beta * Risk Premium

= 5% + 1.39 (4%) = 10.56%

What does this number imply? It does not mean that you will earn 10.56% every year from risk, but it does provide a benchmark that you will have to meet and beat if you are considering Cisco as an investment. For Cisco to be a good investment, you would have to expect it to make more than 10.56% as an annual return in the future.

In summary, in the capital asset pricing model, all the market risk is captured in the beta, measured relative to a market portfolio, which at least in theory should include all traded assets in the market place held in proportion to their market value.

Betas for Other Investments

Most services report betas for publicly traded stocks, but there is no reason why the concept cannot be extended to other investments. You could compute the beta of real estate, gold or even fine art as an investment, just as you computed the beta for Cisco. While analysts have done this and concluded that both real estate and gold are low beta investments (though not necessarily low variance investments), we would add a few cautionary notes. The first is that it is difficult to get traded prices on some alternative investments on a continuous basis. [4] The second is that many analysts continue to use the stock index as their measure of the market portfolio. Since the market portfolio in the capital asset pricing model is supposed to include all traded assets, this likely to give you betas that are biased downwards for non-equity investments.

If you modify the market portfolio to include other traded asset classes and compute betas for alternative investments, you may even find some that have negative betas. While, on the face of it, this may seem absurd, you can get negative betas for investments that reduce the risk (rather than add on to risk) of the market portfolio. Essentially, these investments act as insurance against some large component of market risk, going up as other investments in the portfolio go down. This is the reason why some analysts claim that gold as an investment should have a negative beta, because it tends to do well when inflation increases whereas financial investments are hurt.

B. Alternatives to the Capital Asset Pricing Model

The restrictive assumptions on transactions costs and private information in the capital asset pricing model and the model's dependence on the market portfolio have long been viewed with skepticism by both academics and practitioners. There are three alternatives to the CAPM that have been developed over time:

1. Arbitrage Pricing Model: To understand the arbitrage pricing model, we need to begin with a definition of arbitrage. The basic idea is a simple one. Two portfolios or assets with the same exposure to market risk should be priced to earn exactly the same expected returns. If they are not, you could buy the less expensive portfolio, sell the more expensive portfolio, have no risk exposure and earn a return that exceeds the riskless rate. This is arbitrage. If you assume that arbitrage is not possible and that investors are diversified, you can show that the expected return on an investment should be a function of its exposure to market risk. While this statement mirrors what was stated in the capital asset pricing model,

---

[4] Analysts have tried to get around this problem by using the prices of real estate investment trusts which are traded, but they represent a small fraction of all real estate investments.

the arbitrage pricing model does not make the restrictive assumptions about transactions costs and private information that lead to the conclusion that one beta can capture an investment's entire exposure to market risk. Instead, in the arbitrage pricing model, you can have multiples sources of market risk and different exposures to each (betas) and your expected return on an investment can be written as:

> Expected return = Riskfree rate + Beta for factor 1 (Risk premium for factor 1) + Beta for factor 2 (Risk premium for factor 2)….+ Beta for factor n (Risk premium for factor n)

The practical questions then become knowing how many factors there are that determine expected returns and what the betas for each investment are against these factors. The arbitrage model estimates both by examining historical data on stock returns for common patterns (since market risk affects most stocks) and estimating each stock's exposure to these patterns in a process called factor analysis. A factor analysis provides two output measures:

1. It specifies the number of common factors that affected the historical return data
2. It measures the beta of each investment relative to each of the common factors and provides an estimate of the actual risk premium earned by each factor.

The factor analysis does not, however, identify the factors in economic terms – the factors remain factor 1, factor etc. In summary, in the arbitrage pricing model, the market risk is measured relative to multiple unspecified macroeconomic variables, with the sensitivity of the investment relative to each factor being measured by a beta. The number of factors, the factor betas and factor risk premiums can all be estimated using the factor analysis.

2. Multi-factor Models for risk and return: The arbitrage pricing model's failure to identify the factors specifically in the model may be a statistical strength, but it is an intuitive weakness. The solution seems simple: Replace the unidentified statistical factors with specific economic factors and the resultant model should have an economic basis while still retaining much of the strength of the arbitrage pricing model. That is precisely what multi-factor models try to do. Multi-factor models generally are determined by historical data, rather than economic modeling. Once the number of factors has been identified in the arbitrage pricing model, their behavior over time can be extracted from the data. The behavior of the unnamed factors over time can then be compared to the behavior of macroeconomic variables over that same period to see whether any of the variables is correlated, over time, with the identified factors.

For instance, Chen, Roll, and Ross (1986) suggest that the following macroeconomic variables are highly correlated with the factors that come out of factor analysis: industrial production, changes in default premium, shifts in the term structure, unanticipated inflation, and changes in the real rate of return. These variables can then be correlated with returns to come up with a model of expected returns, with firm-specific betas calculated relative to each variable.

$$E(R) = R_f + \beta_{GNP}\left[E(R_{GNP}) - R_f\right] + \beta_I\left[E(R_I) - R_f\right] + \ldots + \beta_\partial\left[E(R_\partial) - R_f\right]$$

where

β_{GNP} = Beta relative to changes in industrial production

$E(R_{GNP})$ = Expected return on a portfolio with a beta of one on the industrial production factor and zero on all other factors

β_I = Beta relative to changes in inflation

$E(R_I)$ = Expected return on a portfolio with a beta of one on the inflation factor and zero on all other factors

The costs of going from the arbitrage pricing model to a macroeconomic multi-factor model can be traced directly to the errors that can be made in identifying the factors. The economic factors in the model can change over time, as will the risk premia associated with each one. For instance, oil price changes were a significant economic factor driving expected returns in the 1970s but are not as significant in other time periods. Using the wrong factor or missing a significant factor in a multi-factor model can lead to inferior estimates of expected return.

In summary, multi-factor models, like the arbitrage pricing model, assume that market risk can be captured best using multiple macro economic factors and betas relative to each. Unlike the arbitrage pricing model, multi factor models do attempt to identify the macro economic factors that drive market risk.

3. Regression or Proxy Models: All the models described so far begin by defining market risk in broad terms and then developing models that might best measure this market risk. All of them, however, extract their measures of market risk (betas) by looking at historical data. There is a final class of risk and return models that start with the returns and try to explain differences in returns across stocks over long time periods using characteristics such as a firm's market value or price multiples[5]. Proponents of these models argue that if some investments earn consistently higher returns than other investments, they must be riskier. Consequently, we could look at the characteristics that these high-return investments

---

[5] A price multiple is obtained by dividing the market price by its earnings or its book value. Studies indicate that stocks that have low price to earnings multiples or low price to book value multiples earn higher returns than other stocks.

have in common and consider these characteristics to be indirect measures or proxies for market risk.

Fama and French, in a highly influential study of the capital asset pricing model in the early 1990s, noted that actual returns between 1963 and 1990 have been highly correlated with book to price ratios[6] and size. High return investments, over this period, tended to be investments in companies with low market capitalization and high book to price ratios. Fama and French suggested that these measures be used as proxies for risk and report the following regression for monthly returns on stocks on the NYSE:

$$R_t = 1.77\% - 0.11 \ln(MV) + 0.35 \ln\left(\frac{BV}{MV}\right)$$

where

MV = Market Value of Equity

BV/MV = Book Value of Equity / Market Value of Equity

The values for market value of equity and book-price ratios for individual firms, when plugged into this regression, should yield expected monthly returns.

. A Composite of the CAPM and Proxy Models: Three Factor Models

The capital asset pricing model relates the expected return on an investment to its beta against a market portfolio. The proxy models find that there are other variables such as market capitalization and price to book ratios explain returns better than betas. There are composite models that attempt to blend the two and estimated expected returns as a function of betas, market capitalization and price to book ratios. These are also called factor models.

Will these composite models work better than the CAPM? Of course! Should we therefore use them instead of the CAPM? The answer is that it depends on what you are trying to do. If you are trying to explain the past performance of portfolio managers, it may make sense to use composite models, since failing to do so will make portfolio managers who invest in small cap stocks look much better than portfolio managers who invest in large cap stocks. If you are trying to estimate expected returns for the future, to make judgments on where to invest your money, you should be careful about going down this road, since it seems designed to lead the conclusion that everything is fairly priced. Consider why. If there are pockets of the market which are systematically mispriced – say small cap stocks with low price to book ratios – you want to buy these stocks and you will using a conventional risk and return model. If you use a composite model and include market capitalization and price to book ratios as factors, these same stocks will look fairly valued.

[6] The book to price ratio is the ratio of the book value of equity to the market value of equity.

A Comparative Analysis of Risk and Return Models

Figure 2.5 summarizes all the risk and return models in finance, noting their similarities in the first two steps and the differences in the way they define market risk.

Figure 2.5: Risk and Return Models in Finance

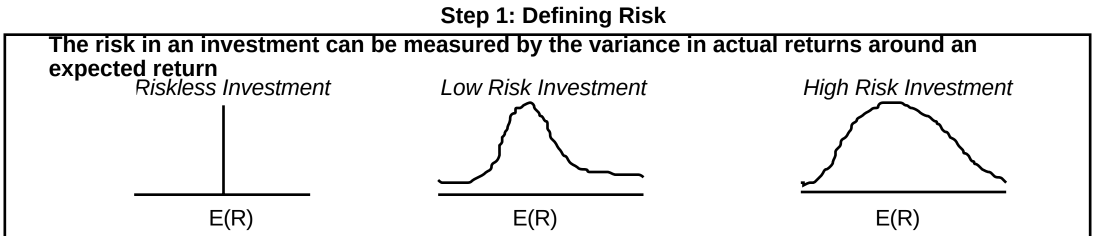

Step 2: Differentiating between Rewarded and Unrewarded Risk

Risk that is specific to investment (Firm Specific)
Can be diversified away in a diversified portfolio
1. each investment is a small proportion of portfolio
2. risk averages out across investments in portfolio

Risk that affects all investments (Market Risk)
Cannot be diversified away since most assets are affected by it.

The marginal investor is assumed to hold a "diversified" portfolio. Thus, only market risk will be rewarded and priced.

Step 3: Measuring Market Risk

The CAPM	The APM	Multi-Factor Models	Proxy Models
If there is 1. no private information 2. no transactions cost the optimal diversified portfolio includes every traded asset. Everyone will hold this market portfolio **Market Risk = Risk added by any investment to the market portfolio:**	If there are no arbitrage opportunities then the market risk of any asset must be captured by betas relative to factors that affect all investments. **Market Risk = Risk exposures of any asset to market factors**	Since market risk affects most or all investments, it must come from macro economic factors. **Market Risk = Risk exposures of any asset to macro economic factors.**	In an efficient market, differences in returns across long periods must be due to market risk differences. Looking for variables correlated with returns should then give us proxies for this risk. **Market Risk = Captured by the Proxy Variable(s)**
Beta of asset relative to Market portfolio (from a regression)	Betas of asset relative to unspecified market factors (from a factor analysis)	Betas of assets relative to specified macro economic factors (from a regression)	Equation relating returns to proxy variables (from a regression)

As noted in Figure 2.5, all the risk and return models developed in this chapter make some assumptions in common. They all assume that only market risk is rewarded and they derive the expected return as a function of measures of this risk. The capital asset pricing model makes the most restrictive assumptions about how markets work but arrives at the simplest model, with only one factor driving risk and requiring estimation. The arbitrage pricing model makes fewer assumptions but arrives at a more complicated model, at least in terms of the parameters that require estimation. The capital asset pricing model can be considered a specialized case of the arbitrage pricing model, where there is only one underlying factor and it is completely measured by the market index. In general, the CAPM has the advantage of being a simpler model to estimate and to use, but it will underperform the richer multi-factor models when an investment is sensitive to economic factors not well represented in the market index. For instance, oil company stocks, which derive most of their risk from oil price movements, tend to have low CAPM betas and low expected returns. An arbitrage pricing model, where one of the factors may measure oil and other commodity price movements, will yield a better estimate of risk and higher expected return for these firms[7].

Which of these models works the best? Is beta a good proxy for risk and is it correlated with expected returns? The answers to these questions have been debated widely in the last two decades. The first tests of the CAPM suggested that betas and returns were positively related, though other measures of risk (such as variance) continued to explain differences in actual returns. This discrepancy was attributed to limitations in the testing techniques. In 1977, Roll, in a seminal critique of the model's tests, suggested that since the market portfolio could never be observed, the CAPM could never be tested, and all tests of the CAPM were therefore joint tests of both the model and the market portfolio used in the tests. In other words, all that any test of the CAPM could show was that the model worked (or did not) given the proxy used for the market portfolio. It could therefore be argued that in any empirical test that claimed to reject the CAPM, the rejection could be of the proxy used for the market portfolio rather than of the model itself. Roll noted that there was no way to ever prove that the CAPM worked and thus no empirical basis for using the model.

Fama and French (1992) examined the relationship between betas and returns between 1963 and 1990 and concluded that there is no relationship. These results have been contested on three fronts. First, Amihud, Christensen, and Mendelson (1992), used the same data, performed different statistical tests and showed that differences in betas did, in fact, explain differences in returns during the time period. Second, Kothari and Shanken (1995) estimated betas using annual data, instead of the shorter intervals used in many tests, and concluded that betas do explain a significant proportion of the differences in returns across investments. Third, Chan and Lakonishok (1993) looked at a much longer time series of returns from 1926 to 1991 and found that the positive relationship between betas and returns broke down only in the period after 1982. They also find that betas are a useful guide to risk in extreme market conditions, with the riskiest firms (the 10% with highest betas) performing far worse than the market as a whole, in the ten worst months for the market between 1926 and 1991 (See Figure 2.6).

[7] Weston and Copeland used both approaches to estimate the cost of equity for oil companies in 1989 and came up with 14.4% with the CAPM and 19.1% using the arbitrage pricing model.

FIGURE 2.6: Returns and Betas: Ten Worst Months between 1926 and 1991

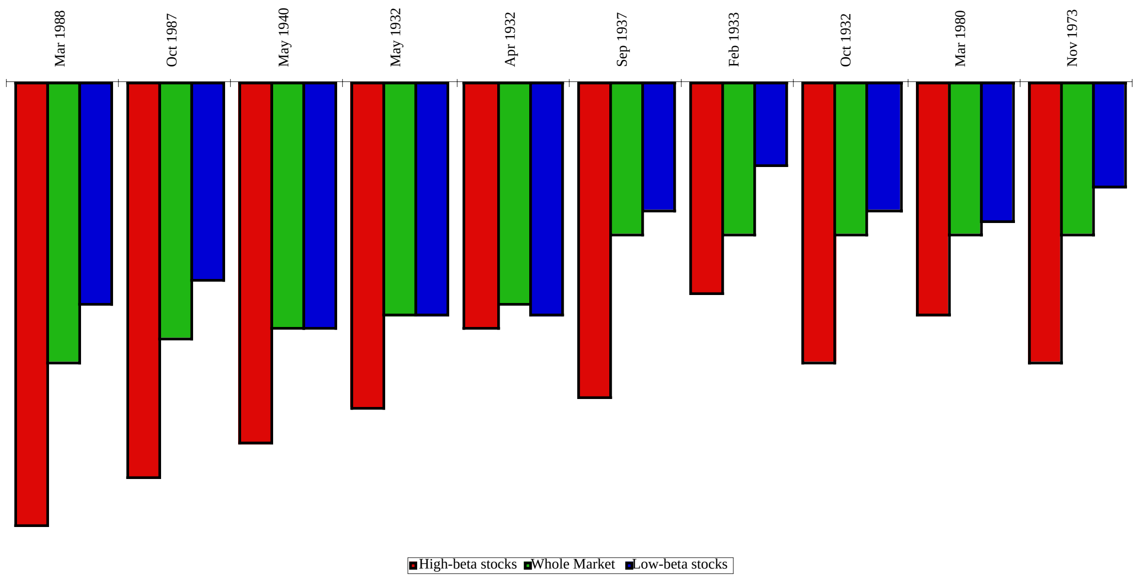

While the initial tests of the APM suggested that they might provide more promise in terms of explaining differences in returns, a distinction has to be drawn between the use of these models to explain differences in past returns and their use to predict expected returns in the future. The competitors to the CAPM clearly do a much better job at explaining past returns since they do not constrain themselves to one factor, as the CAPM does. This extension to multiple factors does become more of a problem when we try to project expected returns into the future, since the betas and premiums of each of these factors now have to be estimated. Because the factor premiums and betas are themselves volatile, the estimation error may eliminate the benefits that could be gained by moving from the CAPM to more complex models. The regression models that were offered as an alternative also have an estimation problem, since the variables that work best as proxies for market risk in one period (such as market capitalization) may not be the ones that work in the next period.

Ultimately, the survival of the capital asset pricing model as the default model for risk in real world applications is a testament to both its intuitive appeal and the failure of more complex models to deliver significant improvement in terms of estimating expected returns. We would argue that a judicious use of the capital asset pricing model, without an over reliance on historical data, is still the most effective way of dealing with risk in modern corporate finance.

Models of Default Risk

The risk that we have discussed hitherto in this chapter relates to cash flows on investments being different from expected cash flows. There are some investments, however, in which the cash flows are promised when the investment is made. This is the case, for instance, when you lend to a business or buy a corporate bond. However, the borrower may default on interest and principal payments on the borrowing. Generally speaking, borrowers with higher default risk should pay higher interest rates on their borrowing than those with lower default risk. This section examines the measurement of default risk and the relationship of default risk to interest rates on borrowing.

In contrast to the general risk and return models for equity, which evaluate the effects of market risk on expected returns, models of default risk measure the consequences of firm-specific default risk on promised returns. While diversification can be used to explain why firm-specific risk will not be priced into expected returns for equities, the same rationale cannot be applied to securities that have limited upside potential and much greater downside potential from firm-specific events. To see what we mean by limited upside potential, consider investing in the bond issued by a company. The coupons are fixed at the time of the issue and these coupons represent the promised cash flow on the bond. The best case scenario for you as an investor is that you receive the promised cash flows; you are not entitled to more than these cash flows even if the company is wildly successful. All other scenarios contain only bad news, though in varying degrees, with the delivered cash flows being less than the promised cash flows. Consequently, the expected return on a corporate bond is likely to reflect the firm-specific default risk of the firm issuing the bond.

The Determinants of Default Risk

The default risk of a firm is a function of two variables. The first is the firm's capacity to generate cash flows from operations and the second is its financial obligations – including interest and principal payments[8]. Firms that generate high cash flows relative to their financial obligations should have lower default risk than firms that generate low cash flows relative to their financial obligations. Thus, firms with significant existing investments, which generate relatively high cash flows, will have lower default risk than firms that do not.

In addition to the magnitude of a firm's cash flows, the default risk is also affected by the volatility in these cash flows. The more stability there is in cash flows the lower the

[8] Financial obligation refers to any payment that the firm has legally obligated itself to make, such as interest and principal payments. It does not include discretionary cash flows, such as dividend payments or new capital expenditures, which can be deferred or delayed, without legal consequences, though there may be economic consequences.

default risk in the firm. Firms that operate in predictable and stable businesses will have lower default risk than will other similar firms that operate in cyclical or volatile businesses.

Most models of default risk use financial ratios to measure the cash flow coverage (i.e., the magnitude of cash flows relative to obligations) and control for industry effects to evaluate the variability in cash flows.

Bond Ratings and Interest rates

The most widely used measure of a firm's default risk is its bond rating, which is generally assigned by an independent ratings agency. The two best known are Standard and Poor’s and Moody’s. Thousands of companies are rated by these two agencies and their views carry significant weight with financial markets.

The Ratings Process

The process of rating a bond usually starts when the issuing company requests a rating from a bond ratings agency. The ratings agency then collects information from both publicly available sources, such as financial statements, and the company itself and makes a decision on the rating. If the company disagrees with the rating, it is given the opportunity to present additional information. This process is presented schematically for one ratings agency, Standard and Poors (S&P), in Figure 2.7.

Figure 2.7: The Ratings Process

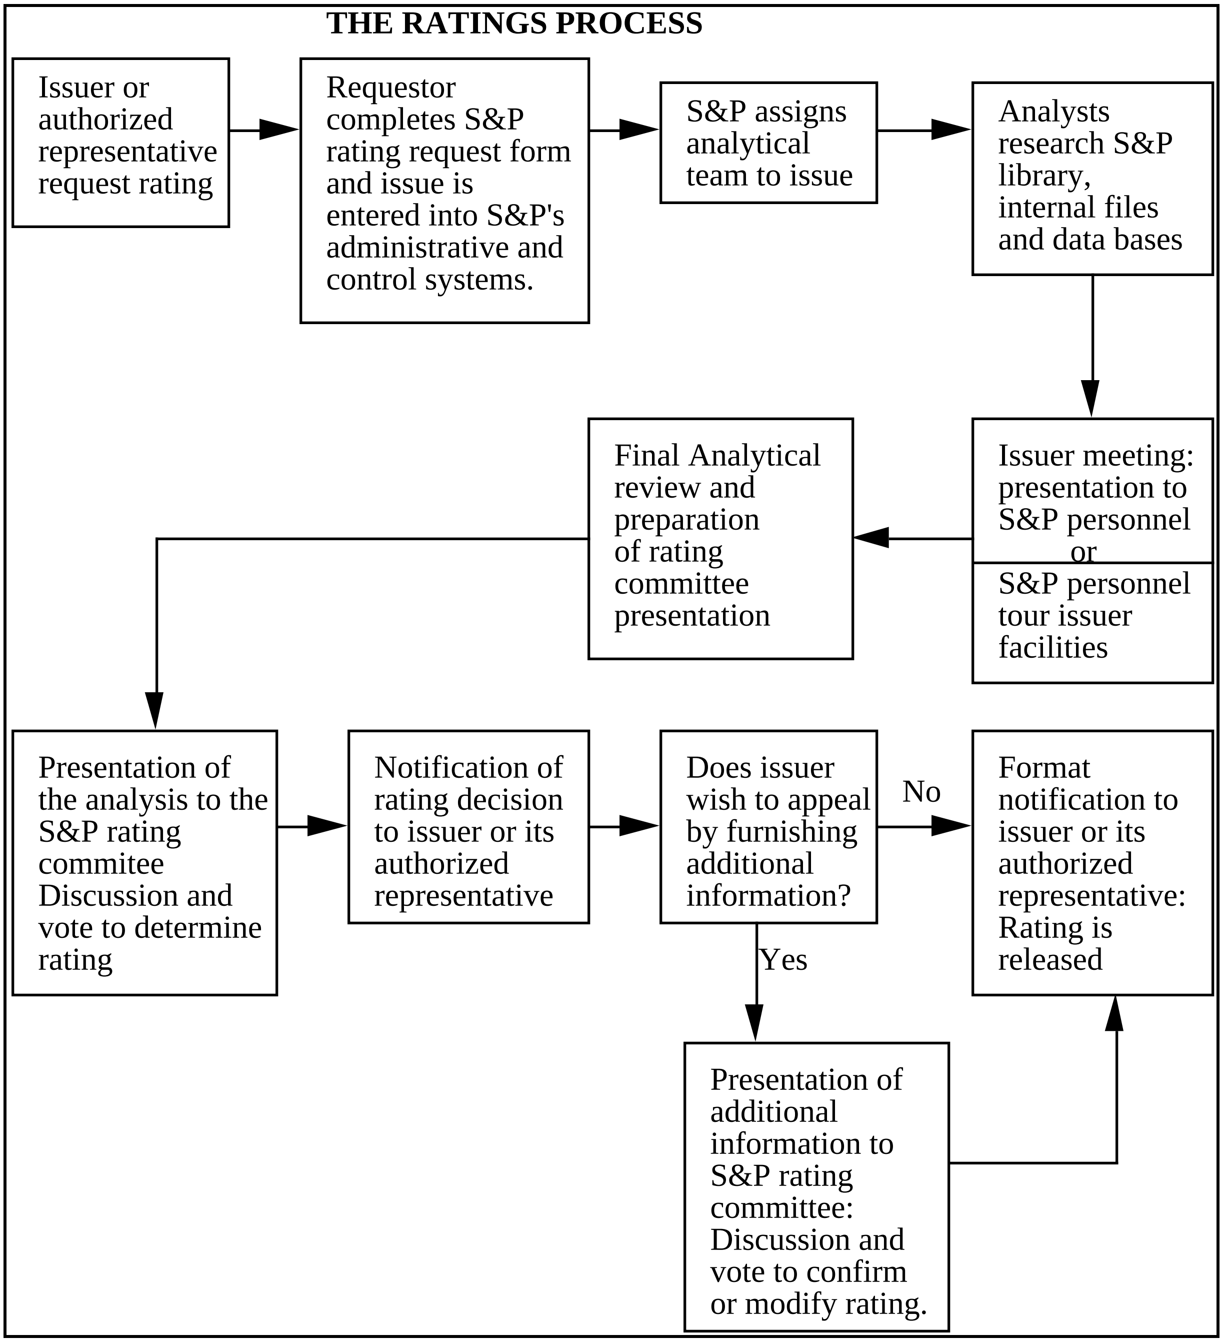

The ratings assigned by these agencies are letter ratings. A rating of AAA from Standard and Poor’s and Aaa from Moody’s represents the highest rating granted to firms that are viewed as having the lowest default risk. As the default risk increases, the ratings decrease toward D for firms in default (Standard and Poor’s). A rating at or above BBB by Standard and Poor’s is categorized as investment grade, reflecting the view of the ratings agency that there is relatively little default risk in investing in bonds issued by these firms.

Determinants of Bond Ratings

The bond ratings assigned by ratings agencies are primarily based upon publicly available information, though private information conveyed by the firm to the rating agency does play a role. The rating assigned to a company's bonds will depend in large part on financial ratios that measure the capacity of the company to meet debt payments and generate stable and predictable cash flows. While a multitude of financial ratios exist, table 2.1 summarizes some of the key ratios used to measure default risk.

Table 2.1: Financial Ratios used to measure Default Risk

Ratio	**Description**
Pretax Interest Coverage	$\frac{\text{Pretax Income from Continuing Operations} + \text{Interest Expense}}{\text{Gross Interest}}$
EBITDA Interest Coverage	$\frac{\text{EBITDA}}{\text{Gross Interest}}$
Funds from Operations / Total Debt	$\frac{\text{Net Income from Continuing Operations} + \text{Depreciation}}{\text{Total Debt}}$
Free Operating Cashflow/ Total Debt	$\frac{\left(\begin{array}{l}\text{Funds from Operations - Capital Expenditures}\\ \text{-Change in Working Capital}\end{array}\right)}{\text{Total Debt}}$
Pretax Return on Permanent Capital	$\frac{\text{Pretax Income from Continuing Operations} + \text{Interest Expense}}{\left(\begin{array}{l}\text{Average of Beginning of the year and End of the year of long and}\\ \text{short term debt, minority interest and Shareholders Equity}\end{array}\right)}$
Operating Income/Sales	$\frac{\left(\begin{array}{l}\text{Sales - COGS(before depreciation) - Selling Expenses -}\\ \text{Administrative Expenses - R \& D Expenses}\end{array}\right)}{\text{Sales}}$

Long Term Debt/ Capital	$\frac{\text{Long Term Debt}}{\text{Long Term Debt} + \text{Equity}}$
Total Debt/Capitalization	$\frac{\text{Total Debt}}{\text{Total Debt} + \text{Equity}}$

Source: Standard and Poors

There is a strong relationship between the bond rating a company receives and its performance on these financial ratios. Table 2.2 provides a summary of the median ratios[9] from 1998 to 2000 for different S&P ratings classes for manufacturing firms.

Table 2.2: Financial Ratios by Bond Rating: 1998-2000

	AAA	*AA*	*A*	*BBB*	*BB*	*B*	*CCC*
EBIT interest cov. (x)	17.5	10.8	6.8	3.9	2.3	1.0	0.2
EBITDA interest cov.	21.8	14.6	9.6	6.1	3.8	2.0	1.4
Funds flow/total debt	105.8	55.8	46.1	30.5	19.2	9.4	5.8
Free oper. cash flow/total debt (%)	55.4	24.6	15.6	6.6	1.9	–4.5	-14.0
Return on capital (%)	28.2	22.9	19.9	14.0	11.7	7.2	0.5
Oper.income/sales (%)	29.2	21.3	18.3	15.3	15.4	11.2	13.6
Long-term debt/capital (%)	15.2	26.4	32.5	41.0	55.8	70.7	80.3
Total Debt/ Capital (%)	26.9	35.6	40.1	47.4	61.3	74.6	89.4
Number of firms	10	34	150	234	276	240	23

Source: Standard and Poors

Note that the pre-tax interest coverage ratio (EBIT) and the EBITDA interest coverage ratio are stated in terms of times interest earned, whereas the rest of the ratios are stated in percentage terms.

Not surprisingly, firms that generate income and cash flows significantly higher than debt payments, that are profitable and that have low debt ratios are more likely

Companies with AAA ratings: Take a look at the companies that commanded triple AAA ratings from Standard and Poor's in the most recent period.

[9] See the Standard and Poor's online site: http://www.standardandpoors.com/ratings/criteria/index.htm

to be highly rated than are firms that do not have these characteristics. There will be individual firms whose ratings are not consistent with their financial ratios, however, because the ratings agency does add subjective judgments into the final mix. Thus, a firm that performs poorly on financial ratios but is expected to improve its performance dramatically over the next period may receive a higher rating than is justified by its current financials. For most firms, however, the financial ratios should provide a reasonable basis for guessing at the bond rating.

Synthetic Ratings and Default Risk

Not all firms that borrow money have bond ratings available on them. How do you go about estimating the cost of debt for these firms? There are two choices.

- One is to look at recent borrowing history. Many firms that are not rated still borrow money from banks and other financial institutions. By looking at the most recent borrowings made by a firm, you can get a sense of the types of default spreads being charged the firm and use these spreads to come up with a cost of debt.
- The other is to estimate a synthetic rating for the firm, i.e, use the financial ratios used by the bond ratings agencies to estimate a rating for the firm. To do this you would need to begin with the rated firms and examine the financial characteristics shared by firms within each ratings class. As an example, assume that you have an unrated firm with operating earnings of $ 100 million and interest expenses of $ 20 million. You could use the interest coverage ratio of 5.00 (100/20) to estimate a bond rating of A- for this firm.[10].

Bond Ratings and Interest Rates

The interest rate on a corporate bond should be a function of its default risk, which is measured by its rating. If the rating is a good measure of the default risk, higher rated bonds should be priced to yield lower interest rates than would lower rated bonds. The difference between the interest rate on a bond with default risk and a default-free government bond is defined to be the *default spread*. Table 2.3 summarizes default spreads for 10-year bonds in S&P's different rating classes as of December 31, 2001:

Table 2.3: Default Spreads and Bond Ratings

Rating	*Spread*

[10] This rating was based upon a table that was developed in 1999 and 2000, by listing out all rated firms, with market capitalization lower than $ 2 billion, and their interest coverage ratios, and then sorting firms based upon their bond ratings. The ranges were adjusted to eliminate outliers and to prevent overlapping ranges.

AAA	0.75%
AA	1.00%
A+	1.50%
A	1.80%
A-	2.00%
BBB	2.25%
BB	3.50%
B+	4.75%
B	6.50%
B-	8.00%
CCC	10.00%
CC	11.50%
C	12.70%
D	14.00%

Source: www.bondsonline.com

These default spreads, when added to the riskless rate, yield the interest rates for bonds with the specified ratings. For instance, a D rated bond has an interest rate about 14% higher than the riskless rate. This default spread will vary by maturity of the bond and can also change from period to period, depending on economic conditions, widening during economic slowdowns and narrowing when the economy is strong.

Summary

Risk, as we define it in finance, is measured based upon deviations of actual returns on an investment from its' expected returns. There are two types of risk. The first, which we call equity risk, arises in investments where there are no promised cash flows, but there are expected cash flows. The second, default risk, arises on investments with promised cash flows.

On investments with equity risk, the risk is best measured by looking at the variance of actual returns around the expected returns, with greater variance indicating greater risk. This risk can be broken down into risk that affects one or a few investments, which we call firm specific risk, and risk that affects many investments, which we refer to as market risk. When investors diversify, they can reduce their exposure to firm specific risk. By assuming that the investors who trade at the margin are well diversified, we conclude that the risk we should be looking at with equity investments is the market risk. The different models of equity risk introduced in this chapter share this objective of measuring market risk, but they differ in the way they do it. In the capital asset pricing model, exposure to market risk is measured by a market beta, which estimates how much risk an individual investment will add to a portfolio that includes all traded assets. The arbitrage pricing model and the multi-factor model allow for multiple sources of market risk and estimate betas for an investment relative to each source. Regression or proxy models for risk look for firm characteristics, such as size, that have been correlated with high returns in the past and use these to measure market risk. In all these models, the risk measures are used to estimate the expected return on an equity investment. This expected return can be considered the cost of equity for a company.

On investments with default risk, risk is measured by the likelihood that the promised cash flows might not be delivered. Investments with higher default risk should have higher interest rates and the premium that we demand over a riskless rate is the default premium. For most US companies, default risk is measured by rating agencies in the form of a company rating; these ratings determine, in large part, the interest rates at which these firms can borrow. Even in the absence of ratings, interest rates will include a default premium that reflects the lenders' assessments of default risk. These default-risk adjusted interest rates represent the cost of borrowing or debt for a business.

Lessons for Investors

1. Your perceptions of how risky an investment may be very different from the risk perceived by the marginal investors (the large institutional investors who set prices at the margin) in an investment. The market prices assets based upon the marginal investors' perceptions of risk.
2. Since the marginal investors are usually well diversified, the only risk that is priced is the risk that cannot be diversified away in a portfolio.
3. Individual risk and return models differ on how to measure this non-diversifable risk. The capital asset pricing model tries to measure it with one beta, whereas multi-factor models try to measure it with multiple betas.
4. The measure of risk allows us to estimate an expected return on a risky investment for the future. This expected return becomes the benchmark that the investment has to beat to be a good investment.
5. For bonds, risk is measured as default or downside risk, since there is not much potential upside. Bond with higher default risk should command higher interest rates.

CHAPTER 3

NUMBERS DON'T LIE OR DO THEY?

Financial statements provide us with the fundamental information that we use to analyze firms. While you may be able to become a successful investor without ever understanding financial statements, it does make the investment process a lot easier if you do. It is important, therefore, that we examine the principles governing these statements and how they help (or fail to help) us answer four questions:

- How valuable are the assets of a firm? The assets of a firm can come in several forms – assets with long lives such as land and buildings, assets with shorter lives such inventory, and intangible assets that still produce revenues for the firm such as patents and trademarks.
- How did the firm raise the funds to finance these assets? In acquiring these assets, firms can use the funds of the owners (equity) or borrowed money (debt), and the mix is likely to change as the assets age.
- How profitable are these assets? To evaluate whether the investments that a firm has already made are good investments, we need to estimate what returns we are making on these investments.
- How much uncertainty (or risk) is embedded in these assets? Estimating how much uncertainty there is in existing investments and the implications for a firm is clearly a first step.

We will look at the way accountants would answer these questions, and why financial statements can provide a misleading picture of a firm's health and success. Some of these differences can be traced to the differences in objectives – accountants try to measure the current standing and immediate past performance of a firm, whereas valuation is much more forward looking.

The Basic Accounting Statements

There are three basic accounting statements that summarize information about a firm. The first is the *balance sheet*, shown in Figure 3.1, which summarizes the assets owned by a firm, the value of these assets and the mix of financing, debt and equity, used to finance these assets at a point in time.

Figure 3.1: The Balance Sheet

	Assets	Liabilities	
Long Lived Real Assets	Fixed Assets	Current Liabilties	Short-term liabilities of the firm
Short-lived Assets	Current Assets	Debt	Debt obligations of firm
Investments in securities & assets of other firms	Financial Investments	Other Liabilities	Other long-term obligations
Assets which are not physical, like patents & trademarks	Intangible Assets	Equity	Equity investment in firm

The next is the *income statement*, shown in Figure 3.2, which provides information on the revenues and expenses of the firm, and the resulting income made by the firm, during a period. The period can be a quarter (if it is a quarterly income statement) or a year (if it is an annual report).

Figure 3.2: Income Statement

Gross revenues from sale of products or services	Revenues
Expenses associates with generating revenues	- Operating Expenses
Operating income for the period	= Operating Income
Expenses associated with borrowing and other financing	- Financial Expenses
Taxes due on taxable income	- Taxes
Earnings to Common & Preferred Equity for Current Period	= Net Income before extraordinary items
Profits and Losses not associated with operations	± Extraordinary Losses (Profits)
Profits or losses associated with changes in accounting rules	± Income Changes Associated with Accounting Changes
Dividends paid to preferred stockholders	- Preferred Dividends
	= Net Income to Common Stockholders

Finally, there is the *statement of cash flows*, shown in figure 3.3, which specifies the sources and uses of cash of the firm from operating, investing and financing activities, during a period.

Figure 3.3: Statement of Cash Flows

Net cash flow from operations, after taxes and interest expenses	Cash Flows From Operations
Includes divestiture and acquisition of real assets (capital expenditures) and disposal and purchase of financial assets. Also includes acquisitions of other firms.	+ Cash Flows From Investing
Net cash flow from the issue and repurchase of equity, from the issue and repayment of debt and after dividend payments	+ Cash Flows from Financing
	= Net Change in Cash Balance

The statement of cash flows can be viewed as an attempt to explain how much the cash flows during a period were, and why the cash balance changed during the period.

Asset Measurement and Valuation

When analyzing any firm, we would like to know the types of assets that it owns, the values of these assets and the degree of uncertainty about these values. Accounting statements do a reasonably good job of categorizing the assets owned by a firm, a partial job of assessing the values of these assets and a poor job of reporting uncertainty about asset values. In this section, we will begin by looking at the accounting principles underlying asset categorization and measurement, and the limitations of financial statements in providing relevant information about assets.

Accounting Principles Underlying Asset Measurement

An asset is any resource that has the potential to either generate future cash inflows or reduce future cash outflows. While that is a general definition broad enough to cover almost any kind of asset, accountants add a caveat that for a resource to be an asset. A firm has to have acquired it in a prior transaction and be able to quantify future benefits with reasonable precision. The accounting view of asset value is to a great extent grounded in the notion of **historical cost**, which is the original cost of the asset, adjusted upwards for improvements made to the asset since purchase and downwards for the loss in value associated with the aging of the asset. This historical cost is called the **book value**. While the generally accepted accounting principles for valuing an asset vary across different kinds of assets, three principles underlie the way assets are valued in accounting statements.

- *An Abiding Belief in Book Value as the Best Estimate of Value*: Accounting estimates of asset value begin with the book value. Unless a substantial reason is given to do otherwise, accountants view the historical cost as the best estimate of the value of an asset.
- *A Distrust of Market or Estimated Value:* When a current market value exists for an asset that is different from the book value, accounting convention seems to view this market value with suspicion. The market price of an asset is often viewed as both much too volatile and too easily manipulated to be used as an estimate of value for an asset. This suspicion runs even deeper when values are is estimated for an asset based upon expected future cash flows.
- *A Preference for under estimating value rather than over estimating it*: When there is more than one approach to valuing an asset, accounting convention takes the view that the more conservative (lower) estimate of value should be used rather than the less conservative (higher) estimate of value. Thus, when both market and book value are available for an asset, accounting rules often require that you use the lesser of the two numbers.

Measuring Asset Value

The financial statement in which accountants summarize and report asset value is the balance sheet. To examine how asset value is measured, let us begin with the way assets are categorized in the balance sheet. First, there are the **fixed assets**, which include the long-term assets of the firm, such as plant, equipment, land and buildings. Next, we have the short-term assets of the firm, including inventory (including raw materials, work in progress and finished goods), receivables (summarizing moneys owed to the firm) and cash; these are categorized as **current assets**. We then have investments in the assets and securities of other firms, which are generally categorized as financial investments. Finally, we have what is loosely categorized as **intangible assets**. These include assets, such as patents and trademarks that presumably will create future earnings and cash flows, and also uniquely accounting assets such as goodwill that arise because of acquisitions made by the firm.

Fixed Assets

Generally accepted accounting principles (GAAP) in the United States require the valuation of fixed assets at historical cost, adjusted for any estimated gain and loss in value from improvements and the aging, respectively, of these assets. While in theory the adjustments for aging should reflect the loss of earning power of the asset as it ages, in practice they are much more a product of accounting rules and convention, and these adjustments are called **depreciation**. Depreciation methods can very broadly be categorized into **straight line** (where the loss in asset value is assumed to be the same every year over its lifetime) and **accelerated** (where the asset loses more value in the earlier years and less in the later years). [While tax rules, at least in the United States, have restricted the freedom that firms have on their choice of asset life and depreciation methods, firms continue to have a significant amount of flexibility on these decisions for reporting purposes. Thus, the depreciation that is reported in the annual reports may not, and generally is not, the same depreciation that is used in the tax statements.

Since fixed assets are valued at book value and are adjusted for depreciation provisions, the value of a fixed asset is strongly influenced by both its depreciable life and the depreciation method used. Many firms in the United States use straight line depreciation for financial reporting while using accelerated depreciation for tax purposes, since firms can report better earnings with the former[1], at least in the years right after the asset is acquired. In contrast, Japanese and German firms often use accelerated depreciation for both tax and financial reporting purposes, leading to reported income which is understated relative to that of their U.S. counterparts.

Current Assets

Current assets include inventory, cash and accounts receivables. It is in this category that accountants are most amenable to the use of market value, especially in valuing marketable securities.

Accounts Receivable

Accounts receivable represent money owed by entities to the firm on the sale of products on credit. The accounting convention is for accounts receivable to be recorded as the amount owed to the firm, based upon the billing at the time of the credit sale. The only major valuation and accounting issue is when the firm has to recognize accounts receivable that are not collectible. Firms can set aside a portion of their income to cover expected **bad debts** from credit sales, and accounts receivable will be reduced by this reserve. Alternatively, the bad debts can be

> *Companies with largest cash balances*: Take a look at the 50 companies with the largest cash balances, as a percent of market

[1] Depreciation is treated as an accounting expense. Hence, the use of straight line depreciation (which is lower than accelerated depreciation in the first few years after an asset is acquired) will result in lower expenses and higher income.

recognized as they occur and the firm can reduce the accounts receivable accordingly. There is the danger, however, that absent a decisive declaration of a bad debt, firms may continue to show as accounts receivable amounts that they know are unlikely to be ever collected.

Cash

Cash is one of the few assets for which accountants and financial analysts should agree on value. The value of a cash balance should not be open to estimation error. Having said this, we should note that fewer and fewer companies actually hold cash in the conventional sense (as currency or as demand deposits in banks). Firms often invest the cash in interest-bearing accounts or in treasuries, so as to earn a return on their investments. In either case, market value can deviate from book value, especially if the investments are long term. While there is no real default risk in either of these investments, interest rate movements can affect their value. We will examine the valuation of marketable securities later in this section.

Inventory

Three basis approaches to valuing inventory are allowed by GAAP: FIFO, LIFO and Weighted Average.

(a) First-in, First-out (FIFO): Under **FIFO**, the cost of goods sold is based upon the cost of material bought earliest in the period, while the cost of inventory is based upon the cost of material bought latest in the year. This results in inventory being valued close to the current replacement cost. During periods of inflation, the use of FIFO will result in the lowest estimate of cost of goods sold among the three valuation approaches, and the highest net income.

(b) Last-in, First-out (LIFO): Under **LIFO**, the cost of goods sold is based upon the cost of material bought latest in the period, while the cost of inventory is based upon the cost of material bought earliest in the year. This results in finished goods being valued close to the current production cost. During periods of inflation, the use of LIFO will result in the highest estimate of cost of goods sold among the three valuation approaches, and the lowest net income.

(c) Weighted Average: Under the weighted average approach, both inventory and the cost of goods sold are based upon the average cost of all materials bought during the period. When inventory turns over rapidly, this approach will more closely resemble FIFO than LIFO.

Firms often adopt the LIFO approach for its tax benefits during periods of high inflation. The cost of goods sold is then higher because it is based upon prices paid towards to the end of the accounting period. This, in turn, will reduce the reported taxable income and net income, while increasing cash flows. Studies indicate that larger firms with rising prices for raw materials and labor, more variable inventory growth and an absence of other tax loss carry forwards are much more likely to adopt the LIFO approach.

Given the income and cash flow effects of inventory valuation methods, it is often difficult to compare the inventory values of firms that use different methods. There is, however, one way of adjusting for these differences. Firms that choose the LIFO approach to value inventories have to specify in a footnote the difference in inventory valuation between FIFO and LIFO, and this difference is termed the **LIFO reserve**. It can be used to adjust the beginning and ending inventories, and consequently the cost of goods sold, and to restate income based upon FIFO valuation.

Investments (Financial) and Marketable Securities

In the category of investments and marketable securities, accountants consider investments made by firms in the securities or assets of other firms, and other marketable securities including treasury bills or bonds. The way in which these assets are valued depends upon the way the investment is categorized and the motive behind the investment. In general, an investment in the securities of another firm can be categorized as a **minority, passive investment**; a **minority, active investment**; or a **majority, active investment**. The accounting rules vary depending upon the categorization.

Minority, Passive Investments

If the securities or assets owned in another firm represent less than 20% of the overall ownership of that firm, an investment is treated as a minority, passive investment. These investments have an acquisition value, which represents what the firm originally paid for the securities and often a market value. Accounting principles require that these assets be sub-categorized into one of three groups: investments that will be held to maturity, investments that are available for sale and trading investments. The valuation principles vary for each.

- For investments that will be held to maturity, the valuation is at historical cost or book value, and interest or dividends from this investment are shown in the income statement under net interest expenses
- For investments that are available for sale, the valuation is at market value, but the unrealized gains or losses are shown as part of the equity in the balance sheet and not in the income statement. Thus, unrealized losses reduce the book value of the equity in the firm, and unrealized gains increase the book value of equity.
- For trading investments, the valuation is at market value and the unrealized gains and losses are shown in the income statement.

Firms are allowed an element of discretion in the way they classify investments and, subsequently, in the way they value these assets. This classification ensures that firms such as investment banks, whose assets are primarily securities held in other firms for purposes of trading, revalue the bulk of these assets at market levels each period. This is called **marking-to-market** and provides one of the few instances in which market value trumps book value in accounting statements.

Minority, Active Investments

If the securities or assets owned in another firm represent between 20% and 50% of the overall ownership of that firm, an investment is treated as a **minority, active investment**. While these investments have an initial acquisition value, a proportional share (based upon ownership proportion) of the net income and losses made by the firm in which the investment was made, is used to adjust the acquisition cost. In addition, the dividends received from the investment reduce the acquisition cost. This approach to valuing investments is called the **equity approach**.

The market value of these investments is not considered until the investment is liquidated, at which point the gain or loss from the sale, relative to the adjusted acquisition cost is shown as part of the earnings under extraordinary items in that period.

Majority, Active Investments

If the securities or assets owned in another firm represent more than 50% of the overall ownership of that firm, an investment is treated as a **majority active investment**[2]. In this case, the investment is no longer shown as a financial investment but is instead replaced by the assets and liabilities of the firm in which the investment was made. This approach leads to a **consolidation** of the balance sheets of the two firms, where the assets and liabilities of the two firms are merged and presented as one balance sheet. The share of the firm that is owned by other investors is shown as a **minority interest** on the liability side of the balance sheet. A similar consolidation occurs in the financial statements of the other firm as well. The statement of cash flows reflects the cumulated cash inflows and outflows of the combined firm. This is in contrast to the equity approach, used for minority active investments, in which only the dividends received on the investment are shown as a cash inflow in the cash flow statement.

[2] Firms have evaded the requirements of consolidation by keeping their share of ownership in other firms below 50%.

Here again, the market value of this investment is not considered until the ownership stake is liquidated. At that point, the difference between the market price and the net value of the equity stake in the firm is treated as a gain or loss for the period.

Intangible Assets

Intangible assets include a wide array of assets ranging from patents and trademarks to goodwill. The accounting standards vary across intangible assets.

1. Patents and Trademarks

Patents and trademarks are valued differently depending on whether they are generated internally or acquired. When patents and trademarks are generated from internal sources, such as research, the costs incurred in developing the asset are expensed in that period even though the asset might have a life of several accounting periods. Thus, the intangible asset is not usually valued in the balance sheet of the firm. In contrast, when an intangible asset is acquired from an external party, it is treated as an asset.

Intangible assets have to be amortized over their expected lives, with a maximum amortization period of 40 years. The standard practice is to use straight-line amortization. For tax purposes, however, firms are not allowed to amortize goodwill or other intangible assets with no specific lifetime.

2. Goodwill

Intangible assets are sometimes the by-products of acquisitions. When a firm acquires another firm, the purchase price is first allocated to tangible assets and then allocated to any intangible assets such as patents or trade names. Any residual becomes **goodwill**. While accounting principles suggest that goodwill captures the value of any intangibles that are not specifically identifiable, it is really a reflection of the difference between the market value of the firm owning the assets and the book value of assets. This approach is called **purchase accounting** and it creates an intangible asset (goodwill) that is usually amortized over time.

Until recently, firms that did not want to see this charge against their earnings, often used an alternative approach called **pooling accounting**, in which the purchase price never showed up in the balance sheet. Instead, the book values of the two companies involved in the merger were aggregated to create the consolidated balance sheet of the combined firm. In 2000, pooling was banned and new rules were created for the treatment of goodwill. Instead of goodwill being written off over a fixed period of time – the old rules required a 40-year amortization period – accountants will now have to reassess the value added in an acquisition and adjust the goodwill accordingly. If the acquired company’s value has decreased since the acquisition, accountants will have to adjust the goodwill for the "impairment in value".

Measuring Financing Mix

The second set of questions that we would like to answer and accounting statements to shed some light on relates to the current value and subsequently the mixture of debt and equity used by the firm. The bulk of the information about these questions is provided on the liability side of the balance sheet and the footnotes.

Accounting Principles Underlying Liability and Equity Measurement

Just as with the measurement of asset value, the accounting categorization of liabilities and equity is governed by a set of fairly rigid principles. The first is a *strict categorization of financing into either a liability or equity* based upon the nature of the obligation. For an obligation to be recognized as a liability, it must meet three requirements:

1. It must be expected to lead to a future cash outflow or the loss of a future cash inflow at some specified or determinable date,
2. The firm cannot avoid the obligation.
3. The transaction giving rise to the obligation has happened already.

In keeping with the earlier principle of conservatism in estimating asset value, accountants recognize as liabilities only cash flow obligations that cannot be avoided.

The second principle is that the value of both liabilities and equity in a firm are *better estimated using historical costs* with accounting adjustments, rather than with expected future cash flows or market value. The process by which accountants measure the value of liabilities and equities is inextricably linked to the way they value assets. Since assets are primarily valued at historical cost or at book value, both debt and equity also get measured primarily at book value. In the section that follows, we will examine the accounting measurement of both liabilities and equity.

Measuring the Value of Liabilities and Equities

Accountants categorize liabilities into current liabilities, long term debt and long term liabilities that are neither debt nor equity. Next, we will examine the way they measure each of these.

Current Liabilities

Current liabilities include all obligations that the firm has coming due in the next accounting period. These generally include:

1. Accounts Payable – representing credit received from suppliers and other vendors to the firm. The value of accounts payable represents the amounts due to these creditors. For this item, book and market value should be similar.
2. Short term borrowing – representing short term loans (due in less than a year) taken to finance the operations or current asset needs of the business. Here again, the value shown represents the amounts due on such loans, and the book and market value should be similar, unless the default risk of the firm has changed dramatically since it borrowed the money.
3. Short term portion of long term borrowing – representing the portion of the long term debt or bonds that is coming due in the next year. Here again, the value shown is the actual amount due on these loans, and market and book value should converge as the due date approaches.
4. Other short term liabilities – which is a catch-all component for any other short term liabilities that the firm might have, including wages due to its employees and taxes due to the government.

Of all the items on the liability side of the balance sheet, absent outright fraud, current liabilities should be the one for which the accounting estimates of book value and financial estimates of market value are the closest.

Long Term Debt

Long term debt for firms can take one of two forms. It can be a long-term loan from a bank or other financial institution or it can be a long-term bond issued to financial markets, in which case the creditors are the investors in the bond. Accountants measure the value of long term debt by looking at the present value of payments due on the loan or bond at the time of the borrowing. For bank loans, this will be equal to the nominal value of the loan. With bonds, however, there are three possibilities: When bonds are issued at par value, for instance, the value of the long-term debt is generally measured in terms of the nominal obligation created, in terms of principal (face value) due on the borrowing. When bonds are issued at a premium or a discount on par value, the bonds are recorded at the issue price, but the premium or discount to the face value is amortized over the life of the bond. As an extreme example, companies that issue zero coupon debt have to record the debt at the issue price, which will be significantly below the principal (face value) due at maturity. The difference between the issue price and the face value is amortized each period and is treated as a non-cash interest expense that is tax deductible.

In all these cases, the book value of debt is unaffected by changes in interest rates during the life of the loan or bond. Note that as market interest rates rise (fall), the present value of the loan obligations should decrease (increase). This updated market value for debt is not shown on the balance sheet. If debt is retired prior to maturity, the difference between book value and the amount paid at retirement is treated as an extraordinary gain or loss in the income statement.

Finally, companies that have long-term debt denominated in non-domestic currencies have to adjust the book value of debt for changes in exchange rates. Since exchange rate changes reflect underlying changes in interest rates, it does imply that this debt is likely to be valued much nearer to market value than is debt in the home currency.

Other Long Term Liabilities

Firms often have long term obligations that are not captured in the long term debt item. These include obligations to lessors on assets that firms have leased, to employees in the form of pension fund and health care benefits yet to be paid, and to the government in the form of taxes deferred. In the last two decades, accountants have increasingly moved towards quantifying these liabilities and showing them as long-term liabilities.

1. Leases

Firms often choose to lease long-term assets rather than buy them. Lease payments create the same kind of obligation that interest payments on debt create, and they must be viewed in a similar light. If a firm is allowed to lease a significant portion of its assets and keep it off its financial statements, a perusal of the statements will give a very misleading view of the company's financial strength. Consequently, accounting rules have been devised to force firms to reveal the extent of their lease obligations on their books.

There are two ways of accounting for leases. In an **operating lease**, the lessor (or owner) transfers only the right to use the property to the lessee. At the end of the lease period, the lessee returns the property to the lessor. Since the lessee does not assume the risk of ownership, the lease expense is treated as an operating expense in the income statement and the lease does not affect the balance sheet. In a **capital lease**, the lessee assumes some of the risks of ownership and enjoys some of the benefits. Consequently, the lease, when signed, is recognized both as an asset and as a liability (for the lease payments) on the balance sheet. The firm gets to claim depreciation each year on the asset and also deducts the interest expense component of the lease payment each year. In general, capital leases recognize expenses sooner than equivalent operating leases.

Since firms prefer to keep leases off the books and sometimes to defer expenses they have a strong incentive to report all leases as operating leases. Consequently the Financial Accounting Standards Board has ruled that a lease should be treated as a capital lease if it meets any one of the following four conditions.

(a) The lease life exceeds 75% of the life of the asset.

(b) There is a transfer of ownership to the lessee at the end of the lease term.

(c) There is an option to purchase the asset at a "bargain price" at the end of the lease term.

(d) The present value of the lease payments, discounted at an appropriate discount rate, exceeds 90% of the fair market value of the asset.

The lessor uses the same criteria for determining whether the lease is a capital or operating lease and accounts for it accordingly. If it is a capital lease, the lessor records the present value of future cash flows as revenue and recognizes expenses. The lease receivable is also shown as an asset on the balance sheet and the interest revenue is recognized over the term of the lease as paid.

From a tax standpoint, the lessor can claim the tax benefits of the leased asset only if it is an operating lease, though the revenue code uses slightly different criteria[3] for determining whether the lease is an operating lease.

2. Employee Benefits

Employers provide pension and health care benefits to their employees. In many cases, the obligations created by these benefits are extensive and a failure by the firm to adequately fund these obligations needs to be revealed in financial statements.

a. Pension Plans

In a pension plan, the firm agrees to provide certain benefits to its employees, either by specifying a 'defined contribution' (wherein a fixed contribution is made to the plan each year by the employer, without any promises as to the benefits which will be delivered in the plan) or a 'defined benefit' (wherein the employer promises to pay a certain benefit to the employee). In the latter case, the employer has to put sufficient money into the plan each period to meet the defined benefits.

Under a defined contribution plan, the firm meets its obligation once it has made the pre-specified contribution to the plan. Under a defined-benefit plan, the firm's obligations are much more difficult to estimate, since they will be determined by a number of variables including the benefits that employees are entitled to, the prior contributions made by the employer, the returns the plan have earned, and the rate of return that the employer expects to make on current contributions. As these variables change, the value of the pension fund assets can be greater than, less than or equal to pension fund liabilities (which is the present

---

[3] The requirements for an operating lease in the revenue code are as follows - (a) the property can be used by someone other than the lessee at the end of the lease term, (b) the lessee cannot buy the asset using a bargain purchase option, (c) the lessor has at least 20% of its capital at risk, (d) the lessor has a positive cash flow from the lease independent of tax benefits and (e) the lessee does not have an investment in the

value of promised benefits). A pension fund whose assets exceed its liabilities is an over-funded plan, whereas one whose assets are less than its liabilities is an under-funded plan and disclosures to that effect have to be included in financial statements, generally in the footnotes.

When a pension fund is over-funded, the firm has several options. It can withdraw the excess assets from the fund, it can discontinue contributions to the plan, or it can continue to make contributions on the assumption that the over-funding is a transitory phenomenon that could well disappear by the next period. When a fund is under-funded, the firm has a liability, though accounting standards require that firms reveal only the excess of accumulated[4] pension fund liabilities over pension fund assets on the balance sheet.

b. Health Care Benefits

A firm can provide health care benefits in one of two ways: by making a fixed contribution to a health care plan, without promising specific benefits (analogous to a defined contribution plan), or by promising specific health benefits and setting aside the funds to provide these benefits (analogous to a defined benefit plan). The accounting for health care benefits is very similar to the accounting for pension obligations. The key difference between the two is that firms do not have to report[5] the excess of their health care obligations over the health care fund assets as a liability on the balance sheet, though a footnote to that effect has to be added to the financial statement.

3. Deferred Taxes

Firms often use different methods of accounting for tax and financial reporting purposes, leading to a question of how tax liabilities should be reported. Since accelerated depreciation and favorable inventory valuation methods for tax accounting purposes lead to a deferral of taxes, the taxes on the income reported in the financial statements will generally be much greater than the actual tax paid. The same principles of matching expenses to income that underlie accrual accounting suggest that the 'deferred income tax' be recognized in the financial statements. Thus a company which pays taxes of $55,000 on its taxable income based upon its tax accounting, and which would have paid taxes of $75,000 on the income reported in its financial statements, will be forced to recognize the difference

---

lease.

[4] The accumulated pension fund liability does not take into account the projected benefit obligation, where actuarial estimates of future benefits are made. Consequently, it is much smaller than the total pension liabilities.

[5] While companies might not have to report the excess of their health care obligations over assets as a liability, some firms choose to do so anyway.

($20,000) as deferred taxes in liabilities. Since the deferred taxes will be paid in later years, they will be recognized as paid.

It is worth noting that companies that actually pay more in taxes than the taxes they report in the financial statements create an asset on the balance sheet called a **deferred tax asset**. This reflects the fact that the firm's earnings in future periods will be greater as the firm is given credit for the deferred taxes.

The question of whether the deferred tax liability is really a liability is an interesting one. Firms do not owe the amount categorized as deferred taxes to any entity, and treating it as a liability makes the firm look more risky than it really is. On the other hand, the firm will eventually have to pay its deferred taxes, and treating it as a liability seems to be the conservative thing to do.

Preferred Stock

When a company issues preferred stock, it generally creates an obligation to pay a fixed dividend on the stock. Accounting rules have conventionally not viewed preferred stock as debt because the failure to meet preferred dividends does not result in bankruptcy. At the same time, the fact the preferred dividends are cumulative makes them more onerous than common equity. Thus, preferred stock is viewed in accounting as a hybrid security, sharing some characteristics with equity and some with debt.

Preferred stock is valued on the balance sheet at its original issue price, with any cumulated unpaid dividends added on. Convertible preferred stock is treated similarly, but it is treated as equity on conversion.

Equity

The accounting measure of equity is a historical cost measure. The value of equity shown on the balance sheet reflects the original proceeds received by the firm when it issued the equity, augmented by any earnings made since (or reduced by losses, if any) and reduced by any dividends paid out during the period. While these three items go into what we can call the book value of equity, a few other items also end up in this estimate.

1. When companies buy back stock for short periods, with the intent of reissuing the stock or using it to cover option exercises, they are allowed to show the repurchased stock as <u>treasury stock</u>, which reduces the book value of equity. Firms are not allowed to keep treasury stock on the books for extended periods and have to reduce their book value of equity by the value of repurchased stock in the case of actions such as stock buybacks. Since these buybacks occur at the current market price, they can result in significant reductions in the book value of equity.

2. Firms that have significant losses over extended periods or carry out massive stock buybacks can end up with negative book values of equity.
3. Relating back to our discussion of marketable securities, any unrealized gain or loss in marketable securities that are classified as available-for-sale is shown as an increase or decrease in the book value of equity in the balance sheet.

As part of their financial statements, firms provide a summary of changes in shareholders equity during the period, where all the changes that occurred to the accounting (book value) measure of equity value are summarized.

Accounting rules still do not seem to have come to grips with the effect of warrants and equity options (such as those granted by many firms to management) on the book value of equity. If warrants are issued to financial markets, the proceeds from this issue will show up as part of the book value of equity. In the far more prevalent case where options are given or granted to management, there is no effect on the book value of equity. When the options are exercised, the cash inflows from the exercise do ultimately show up in the book value of equity and there is a corresponding increase in the number of shares outstanding. The same point can be made about convertible bonds, which are treated as debt until conversion, at which point they become part of equity. In partial defense of accountants, we must note that the effect of options outstanding is often revealed when earnings and book value are computed on a per share basis. Here, the computation is made on two bases, the first on the current number of shares outstanding (primary shares outstanding) and the second on the number of shares outstanding after all options have been exercised (fully diluted shares outstanding).

As a final point on equity, accounting rules still seem to consider preferred stock, with its fixed dividend, as equity or near-equity, largely because of the fact that preferred dividends can be deferred or cumulated without the risk of default. To the extent that there can still be a loss of control in the firm (as opposed to bankruptcy), we would argue that preferred stock shares almost as many characteristics with unsecured debt as it does with equity.

Off Balance Sheet Debt

Towards the end of 2001, we witnessed the incredible collapse of Enron from a firm with more than $ 100 billion in market capitalization to a firm in bankruptcy. While there were other issues involved in the bankruptcy, one of the key ones was the failure of the firm to reveal and of analysts to find out about the billions in dollars of debt that Enron kept off its balance sheet. Enron accomplished this through the use of what are called special purpose entities – partnerships formed with the explicit objective of moving debt off the company's balance sheet. There are legitimate uses of special purpose entities, where firms carve out some of their most liquid and credit-worthy assets (accounts receivable, for instance) into separate entities and let these entities borrow at a rate much lower than what the firm could have borrowed at.[6] Enron, however, used the partnerships to remove troublesome assets of its books, claiming the earnings from these assets and not reporting the debt backing up the assets.

When analyzing a firm, you may want to pay special attention to the footnotes and the other material contained in the filings with the SEC. While this may not give you all the information you need to estimate how much a firm owes, it may give you vital clues about their existence. Firms that have multiple and complicated holding structures, with special purpose entities and partnerships, should viewed with caution. If these firms refuse to reveal fundamental information about their holdings, hiding behind accounting and legal standards, they should be avoided as investments.

Measuring Earnings and Profitability

How profitable is a firm? What did it earn on the assets that it invested in? These are the fundamental questions we would like financial statements to answer. Accountants use the income statement to provide information about a firm's operating activities over a specific time period. In terms of our description of the firm, the income statement is designed to measure the earnings from assets in place. In this section, we will examine the principles underlying earnings and return measurement in accounting, and the methods that they are put into practice.

Accounting Principles Underlying Measurement of Earnings and Profitability

Two primary principles underlie the measurement of accounting earnings and profitability. The first is the principle of **accrual accounting**. In accrual accounting, the revenue from selling a good or service is recognized in the period in which the good is sold or the service is performed (in whole or substantially). A corresponding effort is made on the expense side to match[7] expenses to revenues. This is in contrast to **cash accounting**, where revenues are recognized when payment is received and expenses are recorded when they are paid.

---

[6] If markets were rational, the firm's assets should now be much riskier and the rate at which it borrows should increase. If they are not, however, you may be able to take advantage of market frictions and end up with a much lower borrowing rate.

[7] If a cost (such as an administrative cost) cannot be easily linked with a particular revenues, it is usually recognized as an expense in the period in which it is consumed.

The second principle is the categorization of expenses into operating, financing and capital expenses. **Operating expenses** are expenses that, at least in theory, provide benefits only for the current period; the cost of labor and materials expended to create products that are sold in the current period is a good example. **Financing expenses** are expenses arising from the non-equity financing used to raise capital for the business; the most common example is interest expenses. **Capital expenses** are expenses that are expected to generate benefits over multiple periods; for instance, the cost of buying land and buildings is treated as a capital expense.

Operating expenses are subtracted from revenues in the current period to arrive at a measure of operating earnings from the firm. Financing expenses are subtracted from operating earnings to estimate earnings to equity investors or net income. Capital expenses are written off over their useful life (in terms of generating benefits) as depreciation or amortization.

Measuring Accounting Earnings and Profitability

Since income can be generated from a number of different sources, generally accepted accounting principles (GAAP) require that income statements be classified into four sections: income from continuing operations, income from discontinued operations, extraordinary gains or losses and adjustments for changes in accounting principles.

Generally accepted accounting principles require the recognition of revenues when the service for which the firm is getting paid has been performed in full or substantially and for which it has received in return either cash or a receivable that is both observable and measurable. Expenses linked directly to the production of revenues (like labor and materials) are recognized in the same period in which revenues are recognized. Any expenses that are not directly linked to the production of revenues are recognized in the period in which the firm consumes the services.

While accrual accounting is straightforward in firms that produce goods and sell them, there are special cases where accrual accounting can be complicated by the nature of the product or service being offered. For instance, firms that enter into long term contracts with their customers, for instance, are allowed to recognize revenue on the basis of the percentage of the contract that is completed. As the revenue is recognized on a percentage of completion basis, a corresponding proportion of the expense is also recognized. When there is considerable uncertainty about the capacity of the buyer of a good or service to pay for a service, the firm providing the good or service may recognize the income only when it collects portions of the selling price under the installment method.

Reverting back to our discussion of the difference between capital and operating expenses, operating expenses should reflect only those expenses that create revenues in the current period. In practice, however, a number of expenses are classified as operating expenses that do not seem to meet this test. The first is depreciation and amortization. While the notion that capital expenditures should be written off over multiple periods is reasonable, the accounting depreciation that is computed on the original historical cost often bears little resemblance to the actual economical depreciation. The second expense is research and development expenses, which accounting standards in the United States classify as operating expenses, but which clearly provide benefits over multiple periods. The rationale used for this classification is that the benefits cannot be counted on or easily quantified.

Much of financial analysis is built around the expected future earnings of a firm, and many of these forecasts start with the current earnings. It is therefore important that we know how much of these earnings come from the ongoing operations of the firm, and how much can be attributed to unusual or extraordinary events, that are unlikely to recur on a regular basis. From that standpoint, it is useful that firms categorize expenses into operating and nonrecurring expenses, since it is the earnings prior to extraordinary items that should be used in forecasting. Nonrecurring items include the following:

a. *Unusual or Infrequent items*, such as gains or losses from the divestiture of an asset or division and write-offs or restructuring costs. Companies sometimes include such items as part of operating expenses. As an example, Boeing in 1997 took a write-off of $1,400 million to adjust the value of assets it acquired in its acquisition of McDonnell Douglas, and it showed this as part of operating expenses.
b. *Extraordinary items*, which are defined as events that are unusual in nature, infrequent in occurrence and material in impact. Examples include the accounting gain associated with refinancing high coupon debt with lower coupon debt, and gains or losses from marketable securities that are held by the firm.
c. *Losses associated with discontinued operations*, which measure both the loss from the phase out period and the estimated loss on the sale of the operations. To qualify, however, the operations have to be separable from the firm.
d. *Gains or losses associated with accounting changes*, which measure earnings changes created by accounting changes made voluntarily by the firm (such as a change in inventory valuation and change in reporting period) and accounting changes mandated by new accounting standards.

Measures of Profitability

While the income statement allows us to estimate how profitable a firm is in absolute terms, it is just as important that we gauge the profitability of the firm in comparison terms or percentage returns. Two basic gauges measure profitability. One examines the profitability relative to the capital employed to get a rate of return on investment. This can be done either from the viewpoint of just the equity investors, or by looking at the entire firm. Another examines profitability relative to sales, by estimating a profit margin.

I. Return on Assets (ROA) & Return on Capital (ROC)

The *return on assets* (ROA) of a firm measures its operating efficiency in generating profits from its assets, prior to the effects of financing.

$$\text{ROA} = \frac{\text{EBIT}(1 - \text{tax rate})}{\text{Total Assets}}$$

Earnings before interest and taxes (EBIT) is the accounting measure of operating income from the income statement and total assets refers to the assets as measured using accounting rules, i.e., using book value for most assets. Alternatively, return on assets can be written as:

$$\text{ROA} = \frac{\text{Net Income} + \text{Interest Expenses}(1 - \text{tax rate})}{\text{Total Assets}}$$

Most Profitable firms: Take a look at the 50 firms with the highest returns on capital in the most recent financial year,.

By separating the financing effects from the operating effects, the return on assets provides a cleaner measure of the true return on these assets.

ROA can also be computed on a pre-tax basis with no loss of generality, by using the earnings before interest and taxes (EBIT), and not adjusting for taxes -

$$\text{Pre-tax ROA} = \frac{\text{EBIT}}{\text{Total Assets}}$$

This measure is useful if the firm or division is being evaluated for purchase by an acquirer with a different tax rate or structure.

A more useful measure of return relates the operating income to the capital invested in the firm, where capital is defined as the sum of the book value of debt and equity. This is the *return on capital* (ROC). When a substantial portion of the liabilities is either current (such as accounts payable) or non-interest bearing, this approach provides a better measure of the true return earned on capital employed in the business.

$$\text{After-Tax ROC} = \frac{\text{EBIT}(1-t)}{\text{BV of Debt} + \text{BV of Equity}}$$

$$\text{Pre-Tax ROC} = \frac{\text{EBIT}}{\text{BV of Debt} + \text{BV of Equity}}$$

II. Return on Equity

While the return on capital measures the profitability of the overall firm, the *return on equity* (ROE) examines profitability from the perspective of the equity investor by relating profits to the equity investor (net profit after taxes and interest expenses) to the book value of the equity investment.

$$\text{ROE} = \frac{\text{Net Income}}{\text{Book Value of Common Equity}}$$

Since preferred stockholders have a different type of claim on the firm than do common stockholders, the net income should be estimated after preferred dividends and the book value of common equity should not include the book value of preferred stock. This can be accomplished by using net income after preferred dividends in the numerator and the book value of common equity in the denominator.

Warning Signs in Earnings Reports

The most troubling thing about earnings reports is that we are often blindsided not by the items that get reported (such as extraordinary charges) but by the items that are hidden in other categories. We would suggest the following checklist that should be reviewed about any earnings report to gauge the possibility of such shocks.

- Is earnings growth outstripping revenue growth by a large magnitude year after year? This may well be a sign of increased efficiency, but when the differences are large and continue year after year, you should wonder about the source of these efficiencies.
- Do one-time or non-operating charges to earnings occur frequently? The charge itself might be categorized differently each year – an inventory charge one year, a restructuring charge the next and so on. While this may be just bad luck, it may also reflect a conscious effort by a company to move regular operating expenses into these non-operating items.
- Do any of the operating expenses, as a percent of revenues, swing wildly from year to year? This may suggest that the expense item (say SG&A) includes non-operating expenses that should really be stripped out and reported separately.

- Does the company manage to beat analyst estimates quarter after quarter by a cent or two? Not every company is a Microsoft. Companies that beat estimates year after year are involved in earnings management and are moving earnings across time periods. As growth levels off, this practice can catch up with them.
- Does a substantial proportion of the revenues come from subsidiaries or related holdings? While the sales may be legitimate, the prices set may allow the firm to move earnings from unit to the other and give a misleading view of true earnings at the firm.
- Are accounting rules for valuing inventory or depreciation changed frequently?
- Are acquisitions followed by miraculous increases in earnings? An acquisition strategy is difficult to make successful in the long term. A firm that claims instant success from such as strategy requires scrutiny.
- Is working capital ballooning out as revenues and earning surge? This can sometimes let us pinpoint those firms that generate revenues by lending to their own customers.

None of these factors, by themselves, suggest that we lower earnings for these firms but combinations of the factors can be viewed as a warning signal that the earnings statement needs to be held up to higher scrutiny.

Measuring Risk

How risky are the investments the firm has made over time? How much risk do equity investors in a firm face? These are two more questions that we would like to find the answer to in the course of an investment analysis. Accounting statements do not really claim to measure or quantify risk in a systematic way, other than to provide footnotes and disclosures where there might be risk embedded in the firm. In this section, we will examine some of the ways in which accountants try to assess risk.

Accounting Principles Underlying Risk Measurement

To the extent that accounting statements and ratios do attempt to measure risk, there seem to be two common themes.

a. The first is that the risk being measured is the *risk of default*, i.e. the risk that a fixed obligation, such as interest or principal due on outstanding debt, will not be met. The broader equity notion of risk, which measures the variance of actual returns around expected returns, does not seem to receive much attention. Thus, an all-equity-financed firm with positive earnings and few or no fixed obligations will generally emerge as a low-risk firm from an accounting standpoint, in spite of the fact that its earnings are unpredictable.

b. Accounting risk measures generally take *a static view of risk*, by looking at the capacity of a firm at a point in time to meet its obligations. For instance, when ratios are used to assess a firm's risk, the ratios are almost always based upon one period's income statement and balance sheet.

Accounting Measures of Risk

Accounting measures of risk can be broadly categorized into two groups. The first is disclosures about potential obligations or losses in values that show up as footnotes on balance sheets, which are designed to alert potential or current investors to the possibility of significant losses. The second is the ratios that are designed to measure both liquidity and default risk.

Disclosures in Financial Statements

In recent years, the number of disclosures that firms have to make about future obligations has proliferated. Consider, for instance, the case of **contingent liabilities**. These refer to potential liabilities that will be incurred under certain contingencies, as is the case when a firm is the defendant in a lawsuit. The general rule that has been followed is to ignore contingent liabilities which hedge against risk, since the obligations on the contingent claim will be offset[8] by benefits elsewhere. In recent periods, however, significant losses borne by firms from supposedly hedged derivatives positions (such as options and futures) have led to FASB requirements that these derivatives be disclosed as part of a financial statement. In fact, pension fund and health care obligations have moved from mere footnotes to actual liabilities for firms.

Financial Ratios

Financial statements have long been used as the basis for estimating financial ratios that measure profitability, risk and leverage. In the section on earnings, we looked at two of the profitability ratios – return on equity and return on capital . In this section, we will look at some of the financial ratios that are often used to measure the financial risk in a firm.

1. Short-Term Liquidity Risk

Short-term liquidity risk arises primarily from the need to finance current operations. To the extent that the firm has to make payments to its suppliers before it gets paid for the goods and services it provides, there is a cash shortfall that has to be met, usually through short-term borrowing. Though this financing of working capital needs is

---

[8] This assumes that the hedge is set up competently. It is entirely possible that a hedge, if sloppily set up,

done routinely in most firms, financial ratios have been devised to keep track of the extent of the firm's exposure to the risk that it will not be able to meet its short-term obligations. The two most frequently used to measure short-term liquidity risk are the current ratio and the quick ratio.

The *current ratio* is the ratio of current assets (cash, inventory, accounts receivable) to its current liabilities (obligations coming due within the next period).

$$\text{Current Ratio} = \frac{\text{Current Assets}}{\text{Current Liabilities}}$$

A current ratio below one, for instance, would indicate that the firm has more obligations coming due in the next year than assets it can expect to turn to cash. That would be an indication of liquidity risk.

While traditional analysis suggests that firms maintain a current ratio of 2 or greater, there is a trade-off here between minimizing liquidity risk and tying up more and more cash in net working capital (Net working capital = Current Assets - Current Liabilities). In fact, it can be reasonably argued that a very high current ratio is indicative of an unhealthy firm, which is having problems reducing its inventory. In recent years, firms have worked at reducing their current ratios and managing their net working capital better.

Reliance on current ratios has to be tempered by a few concerns. First, the ratio can be easily manipulated by firms around the time of financial reporting dates to give the illusion of safety; second, current assets and current liabilities can change by an equal amount, but the effect on the current ratio will depend upon its level[9] before the change.

The *quick or acid test ratio* is a variant of the current ratio. It distinguishes current assets that can be converted quickly into cash (cash, marketable securities) from those that cannot (inventory, accounts receivable).

$$\text{Quick Ratio} = \frac{\text{Cash} + \text{Marketable Securities}}{\text{Current Liabilities}}$$

The exclusion of accounts receivable and inventory is not a hard and fast rule. If there is evidence that either can be converted into cash quickly, it can, in fact, be included as part of the quick ratio.

Turnover ratios measure the efficiency of working capital management by looking at the relationship of accounts receivable and inventory to sales and to the cost of goods sold.

---

can end up costing the firm money.

[9] If the current assets and current liabilities increase by an equal amount, the current ratio will go down if it

$$\text{Accounts Receivable Turnover} = \frac{\text{Sales}}{\text{Average Accounts Receivable}}$$

$$\text{Inventory Turnover} = \frac{\text{Cost of Goods Sold}}{\text{Average Inventory}}$$

These ratios can be interpreted as measuring the speed with which the firm turns accounts receivable into cash or inventory into sales. These ratios are often expressed in terms of the number of days outstanding.

$$\text{Days Receivable Outstanding} = \frac{365}{\text{Receivable Turnover}}$$

$$\text{Days Inventory Held} = \frac{365}{\text{Inventory Turnover}}$$

A similar pair of ratios can be computed for accounts payable, relative to purchases.

$$\text{Accounts Payable Turnover} = \frac{\text{Purchases}}{\text{Average Accounts Payable}}$$

$$\text{Days Accounts Payable Outstanding} = \frac{365}{\text{Accounts Payable Turnover}}$$

Since accounts receivable and inventory are assets and accounts payable is a liability, these three ratios (standardized in terms of days outstanding) can be combined to get an estimate of how much financing the firm needs to fund working capital needs.

$$\text{Required Financing Period} = \begin{pmatrix}\text{Days Receivable} \\ \text{Outstanding}\end{pmatrix} + \begin{pmatrix}\text{Days Inventory} \\ \text{Held}\end{pmatrix} + \begin{pmatrix}\text{Days Payable} \\ \text{Outstanding}\end{pmatrix}$$

The greater the financing period for a firm, the greater is its short-term liquidity risk.

2. Long-term Solvency and Default risk

Measures of long-term solvency attempt to examine a firm's capacity to meet interest and principal payments in the long term. Clearly, the profitability ratios discussed earlier in the section are a critical component of this analysis. The ratios specifically designed to measure long term solvency try to relate profitability to the level of debt payments, to identify the degree of comfort with which the firm can meet these payments.

---

was greater than one before the increase and go up if it was less than one.

Interest Coverage Ratios

The *interest coverage ratio* measures the capacity of the firm to meet interest payments from pre-debt, pre-tax earnings.

$$\text{Interest Coverage Ratio} = \frac{\text{EBIT}}{\text{Interest Expenses}}$$

The higher the interest coverage ratio, the more secure is the firm's capacity to make interest payments from earnings. This argument however has to be tempered by the recognition that earnings before interest and taxes is volatile and can drop significantly if the economy enters a recession. Consequently, two firms can have the same interest coverage ratio but be viewed very differently in terms of risk.

The denominator in the interest coverage ratio can be easily extended to cover other fixed obligations such as lease payments. If this is done, the ratio is called a *fixed charges coverage ratio*.

$$\text{Fixed Chargeds Coverage Ratio} = \frac{\text{EBIT + Fixed Charges}}{\text{Fixed Charges}}$$

Finally, this ratio, while stated in terms of earnings, can be restated in terms of cash flows, by using earnings before interest, taxes and depreciation (EBITDA) in the numerator and cash fixed charges in the denominator.

$$\text{Cash Fixed Charges Coverage Ratio} = \frac{\text{EBITDA}}{\text{Cash Fixed Charges}}$$

Both interest coverage and fixed charge ratios are open to the criticism that they do not consider capital expenditures, a cash flow that may be discretionary in the very short term, but not in the long term if the firm wants to maintain growth. One way of capturing the extent of this cash flow, relative to operating cash flows, is to compute a ratio of the two.

$$\text{Operating Cash flow to Capital Expenditures} = \frac{\text{Cash flows from Operations}}{\text{Capital Expenditures}}$$

While there are a number of different definitions of cash flows from operations, the most reasonable way of defining it is to measure the cash flows from continuing operations, before interest but after taxes, and after meeting working capital needs.

$$\text{Cash flow from operations} = \text{EBIT (1-tax rate)} - \Delta \text{ Working Capital}$$

Debt Ratios

Interest coverage ratios measure the capacity of the firm to meet interest payments but do not examine whether it can pay back the principal on outstanding debt. *Debt ratios* attempt to do this, by relating debt to total capital or to equity. The two most widely used debt ratios are:

$$\text{Debt to Capital Ratio} = \frac{\text{Debt}}{\text{Debt + Equity}}$$

$$\text{Debt to Equity Ratio} = \frac{\text{Debt}}{\text{Equity}}$$

Most Levered firms: Take a look at the 50 firms with the highest market debt to capital and debt to equity ratios.

The first ratio measures debt as a proportion of the total capital of the firm and cannot exceed 100%. The second measures debt as a proportion of equity in the firm and can be easily derived from the first.

$$\text{Debt/Equity Ratio} = \frac{\text{Debt/Capital Ratio}}{1 - \text{Debt/Capital Ratio}}$$

While these ratios presume that capital is raised from only debt and equity, they can be easily adapted to include other sources of financing, such as preferred stock. While preferred stock is sometimes combined with common stock under the ‘equity’ label, it is better to keep it separate and to compute the ratio of preferred stock to capital (which will include debt, equity and preferred stock).

a. Variants on Debt Ratios

There are two close variants of debt ratios. In the first, only long-term debt is used rather than total debt, with the rationale that short-term debt is transitory and will not affect the long-term solvency of the firm.

$$\text{Long term Debt to Capital Ratio} = \frac{\text{Long term Debt}}{\text{Long term Debt + Equity}}$$

$$\text{Long term Debt to Equity Ratio} = \frac{\text{Long term Debt}}{\text{Equity}}$$

Given the ease with which firms can roll over short-term debt, and the willingness of many firms to use short-term financing to fund long-term projects, these variants can provide a misleading picture of the firm's financial leverage risk.

The second variant of debt ratios uses market value (MV) instead of book value, primarily to reflect the fact that some firms have a significantly greater capacity to borrow than their book values indicate.

$$\text{Market Value Debt to Capital Ratio} = \frac{\text{MV of Debt}}{\text{MV of Debt} + \text{MV of Equity}}$$

$$\text{Market Value Debt to Equity Ratio} = \frac{\text{MV of Debt}}{\text{MV of Equity}}$$

Many analysts disavow the use of market value in their calculations, contending that market values, in addition to being difficult to get for debt, are volatile and hence unreliable. These contentions are open to debate. It is true that the market value of debt is difficult to get for firms which do not have publicly traded bonds, but the market value of equity is not only easy to obtain, it is constantly updated to reflect market-wide and firm-specific changes. Furthermore, using the book value of debt as a proxy for market value in those cases where bonds are not traded does not significantly shift[10] most market-value based debt ratios.

Differences in accounting standards and practices

Differences in accounting standards across countries affect the measurement of earnings. These differences, however, are not so great as they are made out to be and they cannot explain away radical departures from fundamental principles of valuation[11]. Choi and Levich, in a survey of accounting standards across developed markets, note that most countries subscribe to basic accounting notions of consistency, realization and historical cost principles in preparing accounting statements.

The two countries that offer the strongest contrast to the United States are Germany and Japan. The key differences and their implications are as follows. First, companies in the United States generally maintain separate tax and financial reporting books, which in turn generates items like deferred taxes to cover differences between the two books. Companies in Germany and Japan do not maintain separate books. Consequently, depreciation methods in financial reports are much more likely to be accelerated and hence to reduce stated

[10] Deviations in the market value of equity from book value are likely to be much larger than deviation for debt and are likely to dominate in most debt ratio calculations.

[11] At the peak of the Japanese market, there were many investors who explained away the price-earnings multiples of 60 and greater in the market, by noting that Japanese firms were conservative in measuring earnings. Even after taking into account the general provisions and excess depreciation used by many of these firms to depress current earnings, the price-earnings multiples were greater than 50 for many firms, suggesting either extraordinary expected growth in the future or overvaluation.

income. Second, the requirement that leases be capitalized and shown as a liability is much more tightly enforced in the United States. In Japan, leases are generally treated as operating leases and do not show up as liabilities in the balance sheet. In Germany, firms can capitalize leases, but they have more leeway in classifying leases as operating and capital leases than U.S. companies. Third, goodwill, once created, can be amortized over 40 years in the United States and over much shorter time periods in Germany and Japan, again depressing stated income. Fourth, reserves in the United States can be created only for specific purposes, whereas German and Japanese companies can use general reserves to equalize earnings across periods, leading earnings to be understated during the good years, and overstated during bad years.

Most of these differences can be accounted and adjusted for when comparisons are made between companies in the U.S. and companies in other financial markets. Ratios such as price earnings, which use stated and unadjusted earnings, can be misleading when accounting standards vary widely across the companies being compared.

Summary

Financial statements remain the primary source of information for most investors and analysts. There are differences, however, in how accounting and financial analysis approach answering a number of key questions about the firm. We examine these differences in this chapter.

The first question that we examined related to the nature and the value of the assets owned by a firm. Categorizing assets into investments already made (assets in place) and investments yet to be made (growth assets), we argued that accounting statements provide a substantial amount of historical information about the former and very little about the latter. The focus on the original price of assets in place (book value) in accounting statements can lead to significant differences between the stated value of these assets and their market value. With growth assets, accounting rules result in low or no values for assets generated by internal research.

The second issue that we examined was the measurement of profitability. The two principles that seem to govern how profits are measured are accrual accounting – revenues and expenses are shown in the period where transactions occur rather than when the cash is received or paid – and the categorization of expenses into operating, financing and capital expenses. Operating and financing expenses are shown in income statements. Capital expenditures take the form of depreciation and amortization and are spread over several time periods. Accounting standards miscategorize operating leases and research and development expenses as operating expenses (when the former should be categorized as financing expenses and the latter as capital expenses).

In the last part of the chapter, we examine how financial statements deal with short-term liquidity risk and long-term default risk. While the emphasis in accounting statements is on examining the risk that firms may be unable to make payments that they have committed to make, there is very little focus on risk to equity investors.

Lessons for Investors

1. The purpose of accounting statements is to give you a measure of how a company performed in the past. Your objective in investing is to consider how a firm will perform in the future.
2. Accounting rules provide significant discretion to firms in how they measure and report earnings. Firms that adopt aggressive accounting practices, even though they might be legal, will report higher earnings than firms that adopt more conservative practices.
3. As firms age, the book value of their assets will become less and less relevant as measures of what the assets are truly worth.
4. Firms with operating leases and off-balance sheet financing owe much more than what they reveal as debt on their balance sheets.
5. The footnotes to the financial statements often carry more information than the financial statements themselves.

CHAPTER 4

SHOW ME THE MONEY: THE BASICS OF VALUATION

To invest wisely, you need to understand the principles of valuation. In this chapter, we examine those fundamental principles. In general, you can value an asset in one of three ways. You can estimate the intrinsic value of the asset by looking at its capacity to generate cashflows in the future. You can estimate a relative value, by examining how the market is pricing similar or comparable assets. Finally, you can value assets with cashflows that are contingent on the occurrence of a specific event as options.

With intrinsic valuation, we argue that the value of any asset is the present value of the expected cash flows on the asset, and it is determined by the magnitude of the cash flows, the expected growth rate in these cash flows and the uncertainty associated with receiving these cash flows. We begin by looking at assets with guaranteed cash flows over a finite period, and then we extend the discussion to cover the valuation of assets when there is uncertainty about expected cash flows. As a final step, we consider the valuation of a firm, with the potential, at least, for an infinite life and uncertainty in the cash flows.

With relative valuation, we begin by looking for similar or comparable assets. When valuing stocks, these are often defined as other companies in the same business. We then standardize convert the market values of these companies which are dollar values to multiples of some standard variable – earnings, book value and revenues are widely used. We then compare the valuations of the comparable companies to try to find misvalued companies.

There are some assets that cannot be valued using either discounted cashflow or relative valuation models because the cashflows are contingent on the occurrence of a specific event. These assets can be valued using option pricing models. We consider the basic principles that uuderlie these models in this chapter.

Intrinsic Value

We can estimate the value of an asset by taking the present value of the expected cash flows on that asset. Consequently, the value of any asset is a function of the cash flows generated by that asset, the life of the asset, the expected growth in the cash flows and the riskiness associated with the cash flows. We will begin this section by looking at valuing assets that have finite lives (at the end of which they cease to generate cash flows) and conclude by looking at the more difficult case of assets with infinite lives. We will also start the process by looking at firms whose cash flows are known with certainty and conclude by looking at how we can consider uncertainty in valuation.

The Mechanics of Present Value

Almost everything we do in intrinsic valuation rests on the concept of present value. The intuition of why a dollar today is worth more than a dollar a year from now is simple. Our preferences for current over future consumption, the effect of inflation on the buying power of a dollar and uncertainty about whether we will receive the future dollar all play a role in determining how much of a discount we apply to the future dollar. In annualized terms, this discount is measured with a discount rate. It is worth, however, reviewing the basic mechanics of present value before we consider more complicated valuation questions.

In general, there are five types of cash flows that we will encounter in valuing any asset. You can have a single cash flow in the future, a set of equal cashflows each period for a number of periods (annuity), a set of equal cashflows each period forever (perpetuity), a set of cashflows growing at a constant rate and each period for a number of periods (growing annuity) and a cash flow that grows at a constant rate forever (growing perpetuity).

The present value of a single cashflow in the future can be obtained by discounting the cashflow back at the . Thus, the value of $ 10 million in 5 years, with a discount rate of 15% can be written as:

$$\text{Present value of \$ 10 million in 5 years} = \frac{\$10}{(1.15)^5} = \text{ \$ 4.97 million}$$

You could read this present value to mean that you would be indifferent between receiving $4.97 million today or $ 10 million in 5 years.

What about the present value of an annuity? You have two choices. One is to discount each of the annual cashflows back to the present and add them all up. For instance, if you had an annuity of $ 5000 every year for the next 5 years and a discount rate of 10%, you could compute the present value of the annuity in figure 4.1:

Figure 4.1 Cash Flows on Annuiity

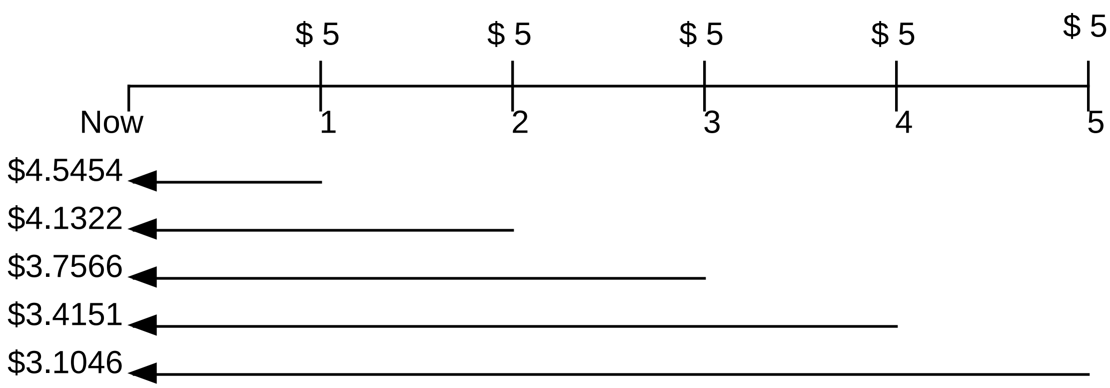

Adding up the present values yields \$18.95 million. Alternatively, you could use a short cut – an annuity formula – to arrive at the present value:

$$\text{PV of an Annuity} \;=\; A\left[\frac{1 - \dfrac{1}{(1+r)^n}}{r}\right] = 5\left[\frac{1 - \dfrac{1}{(1.1)^5}}{.10}\right] = \$18.95$$

Getting from the present value of an annuity to the present value of a perpetuity is simple. Setting n to ∞ in the above equation yields the present value of a perpetuity

$$\text{PV of an Perpetuity} \;=\; A\left[\frac{1 - \dfrac{1}{(1+r)^\infty}}{r}\right] = \frac{A}{r}$$

Thus, the present value of \$ 5 million each year forever at a discount rate of 10% is \$ 50 million (\$5 million/ .10 = \$ 50 million)

Moving from a constant cashflow to one that grows at a constant rate yields a growing annuity. For instance, if we assume that the \$ 5 million in annual cashflows will grow 20% a year for the next 5 years, we can estimate the present value in figure 4.2:

Figure 4.2 Cash Flows on Growing Annuity

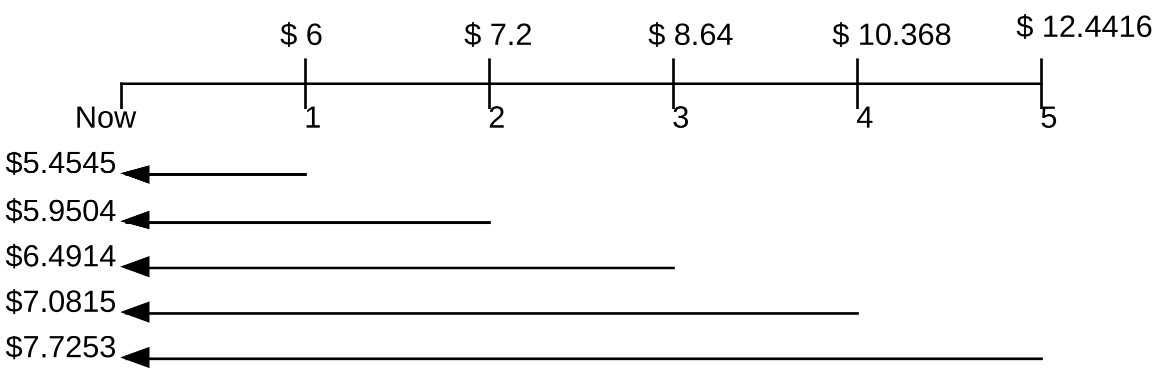

Summing up these present values yields a total value of \$32.70 million. Here again, there is a short cut available in the form of a growing annuity formula:

$$PV \text{ of a Growing Annuity} \;=\; A(1+g)\left[\frac{1 - \dfrac{(1+g)^n}{(1+r)^n}}{r - g}\right] = 5(1.20)\left[\frac{1 - \dfrac{(1.20)^5}{(1.10)^5}}{.10 - .20}\right] = \$32.70$$

Finally, consider a cashflow growing at a constant rate forever – a growing perpetuity. Substituting into the equation above, we get:

$$PV \text{ of a Growing Perpetuity } = A(1+g)\left[\frac{1 - \dfrac{(1+g)^{\infty}}{(1+r)^{\infty}}}{r - g}\right] = \frac{A(1+g)}{(r-g)}$$

Note that the fact the cashflows grow at a constant rate forever constrains this rate to be less than or equal to the growth rate of the economy in which you operate. Working with U.S. dollars, this growth rate should not exceed 5-6%.

Valuing an Asset with Guaranteed Cash Flows

The simplest assets to value have cash flows that are guaranteed, i.e, assets whose promised cash flows are always delivered. Such assets are riskless, and the interest rate earned on them is called a **riskless rate**. The value of such an asset is the present value of the cash flows, discounted back at the riskless rate. Generally speaking, riskless investments are issued by governments that have the power to print money to meet any obligations they otherwise cannot cover. Not all government obligations are not riskless, though, since some governments have defaulted on promised obligations.

Default-free Zero-coupon Bond

The simplest asset to value is a bond that pays no coupon but has a face value that is guaranteed at maturity; this bond is a *default-free zero coupon bond*. We can show the cash flow on this bond as in Figure 4.3.

Figure 4.3: Cash Flows on N-year Zero Coupon Bond

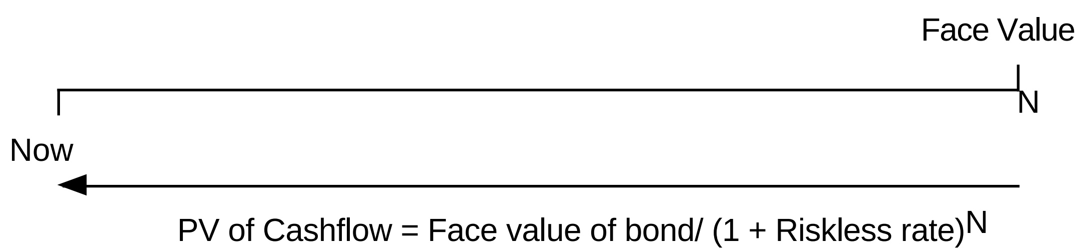

The value of this bond can be written as the present value of a single cash flow discounted back at the riskless rate where N is the maturity of the zero-coupon bond. Since the cash flow on this bond is fixed, the value of the bond will increase as the riskless rate decreases and decrease as the riskless rate increases.

To see an example of this valuation at work, assume that the ten-year interest rate on riskless investments is 4.55%, and that you are pricing a zero-coupon treasury bond, with a maturity of ten years and a face value of $ 1000. The price of the bond can be estimated as follows:

$$\text{Price of the Bond} = \frac{\$1,000}{(1.0455)^{10}} = \$\,640.85$$

Note that the face value is the only cash flow, and that this bond will be priced well below the face value of $ 1,000. Such a bond is said to be trading below par.

Conversely, we could estimate a default-free interest rate from the price of a zero-coupon treasury bond. For instance, if the 10-year zero coupon treasury were trading at $ 593.82, the default-free ten-year spot rate can be estimated as follows:

$$\text{Default-free Spot Rate} = \left(\frac{\text{Face Value of Bond}}{\text{Market Value of Bond}}\right)^{1/t} - 1 = \left(\frac{1000}{593.82}\right)^{1/10} - 1 = .0535$$

The ten-year default free rate is 5.35%.

Default-free Coupon Bond

Consider, now, a default-free coupon bond, which has fixed cash flows (coupons) that occur at regular intervals (usually semi annually) and a final cash flow (face value) at maturity. The time line for this bond is shown in Figure 4.4 (with C representing the coupon each period and N being the maturity of the bond).

Figure 4.4: Cash Flows on N-year Coupon Bond

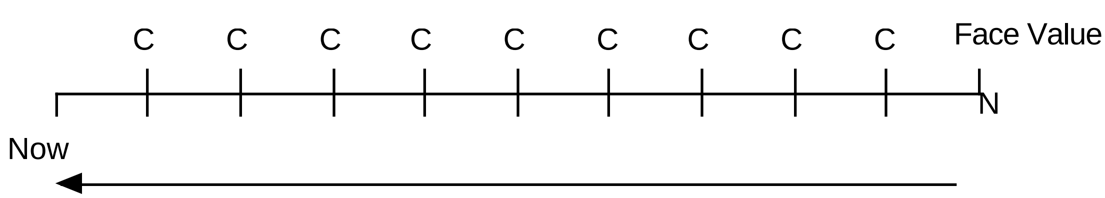

Present value of cashflows = Present value of coupons + Present value of Face Value

This bond can actually be viewed as a series of zero-coupon bonds, and each can be valued using the riskless rate that corresponds to when the cash flow comes due:

$$\text{Value of Bond} = \sum_{t=1}^{t=N} \frac{\text{Coupon}}{(1+r_1)^1} + \frac{\text{Coupon}}{(1+r_2)^2} + \frac{\text{Coupon}}{(1+r_3)^3} \ldots + \frac{\text{Coupon}}{(1+r_N)^N} + \frac{\text{Face Value of the Bond}}{(1+r_N)^N}$$

where r_t is the interest rate that corresponds to a t-period zero coupon bond and the bond has a life of N periods.

It is, of course, possible to arrive at the same value using *some weighted average* of the period-specific riskless rates used above; the weighting will depend upon how large each cash flow is and when it comes due. This weighted average rate is called the *yield to maturity*, and it can be used to value the same coupon bond:

$$\text{Value of Bond} = \sum_{t=1}^{t=N} \frac{\text{Coupon}}{(1+r)^1} + \frac{\text{Coupon}}{(1+r)^2} + \frac{\text{Coupon}}{(1+r)^3} \dots + \frac{\text{Coupon}}{(1+r)^N} + \frac{\text{Face Value of the Bond}}{(1+r_N)^N}$$

where r is the yield to maturity on the bond. Like the zero-coupon bond, the default-free coupon bond should have a value that varies inversely with the yield to maturity. As we will see shortly, since the coupon bond has cash flows that occur earlier in time (the coupons) it should be less sensitive to a given change in interest rates than a zero-coupon bond with the same maturity.

Consider now a five-year treasury bond with a coupon rate of 5.50%, with coupons paid every 6 months. We will price this bond initially using default-free spot rates for each cash flow in Table 4.1.

Table 4.1: Value of 5-year default-free bond

Time	*Coupon*	*Default-free Rate*	*Present Value*
0.5	$ 27.50	4.15%	$ 26.95
1	$ 27.50	4.30%	$ 26.37
1.5	$ 27.50	4.43%	$ 25.77
2	$ 27.50	4.55%	$ 25.16
2.5	$ 27.50	4.65%	$ 24.55
3	$ 27.50	4.74%	$ 23.93
3.5	$ 27.50	4.82%	$ 23.32
4	$ 27.50	4.90%	$ 22.71
4.5	$ 27.50	4.97%	$ 22.11
5	$ 1,027.50	5.03%	$ 803.92
			$ 1,024.78

The default-free spot interest rates reflect the market interest rates for zero coupon bonds for each maturity. The bond price can be used to estimate a weighted-average interest rate for this bond:

$$\$1,024.78 = \sum_{t=0.5}^{t=5} \frac{\$27.50}{(1+r)^t} + \frac{\$1,000}{(1+r)^5}$$

Solving for r, we obtain a rate of 4.99%, which is the yield to maturity on this bond.

Bond Value and Interest Rate Sensitivity and Duration

As market interest rates change, the market value of a bond will change. Consider, for instance, the 10-year zero coupon bond and the 5-year coupon bond described in the last two illustrations. Figure 4.5 shows the market value of each of these bonds as market interest rates vary from 3% to 10%.

Figure 4.5: Interest Rates and Bond Prices

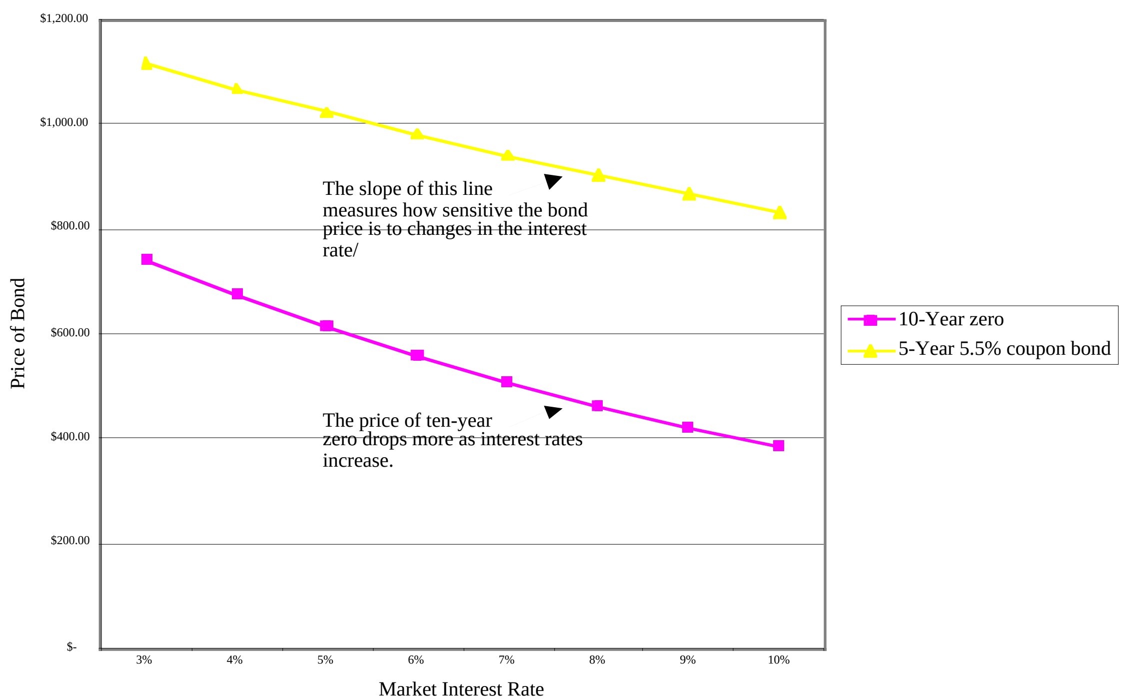

> *bondval.xls*: See the spreadsheet that includes the bond valuation examples in this chapter.

Note that the price of the 10-year zero-coupon bond is much more sensitive to interest rate changes than is the 5-year coupon bond to a given change in market interest rates. The 10-year zero coupon bond loses about half its value as interest rates increase from 3% to 10%; in contrast, the 5-year 5.5% coupon bond loses about 30% of its value. This should not be surprising since the present value effect of that interest rate increases the larger the cash flow, and the further in the future it occurs. Thus longer-term bonds will be more sensitive to interest rate changes than shorter-term bonds, with similar coupons. Furthermore, low-coupon or no-coupon bonds will be more sensitive to interest rate changes than high-coupon bonds.

The interest rate sensitivity of a bond, which is a function of both the coupon rate and the maturity of the bond, can be captured in one measure called the <u>duration</u>. The greater the duration of a bond, the more sensitive its price is to interest rate movements.. The simplest measure of duration, called Macaulay duration, can be viewed as a weighted maturity of the different cash flows on the bond.

$$\text{Duration of a Bond} = \frac{\sum_{t=1}^{t=N} t \frac{CF_t}{(1+r)^t}}{\sum_{t=1}^{t=N} \frac{CF_t}{(1+r)^t}}$$

where r is the yield to maturity on the bond.

For a zero-coupon bond, which has only one cash flow, due at maturity, the duration is equal to the maturity.

Duration of 10-year zero-coupon bond = 10 years

The duration of the 5-year coupon bond requires a few more calculations, is calculated in the Table 4.2:

Table 4.2: Value of a 5-year Coupon Bond

Time (t)	*Coupon*	*Present Value (at 4.99%)*	*t *Present Value*
0.5	$27.50	$26.84	$13.42
1	$27.50	$26.19	$26.19
1.5	$27.50	$25.56	$38.34
2	$27.50	$24.95	$49.90
2.5	$27.50	$24.35	$60.87
3	$27.50	$23.76	$71.29
3.5	$27.50	$23.19	$81.17
4	$27.50	$22.63	$90.53
4.5	$27.50	$22.09	$99.40
5	$1,027.50	$805.46	$4,027.28
Sum		$1,025.02	$4,558.39

Duration of 5-year 5.5% coupon bond = $4,558/$1,025 = 4.45

The longer the duration of a bond, the more sensitive it is to interest rate changes. In our illustrations above, the ten-year coupon bond has a higher duration and will therefore be more sensitive to interest rate changes than the five-year coupon bond.

Introducing Uncertainty into Valuation

We have to grapple with two different types of uncertainty in valuation. The first arises in the context of securities like bonds, where there is a promised cash flow to the holder of the bonds in future periods. The risk that these cash flows will not be delivered is called **default risk**; the greater the default risk in a bond, given its cash flows, the less valuable the bond will become.

The second type of risk is more complicated. When we make equity investments in assets, we are generally not promised a fixed cash flow but are entitled, instead, to whatever cash flows are left over after other claim holders (like debt) are paid; these cash flows are called *residual cash flows*. Here, the uncertainty revolves around what these residual cash flows will be, relative to expectations. In contrast to default risk, where the risk can only result in negative consequences (the cash flows delivered will be less than promised), uncertainty in the context of equity investments can cut both ways. The actual cash flows can be much lower than expected, but they can also be much higher. For the moment, we will label this risk **equity risk** and consider, at least in general terms, how best to deal with it in the context of valuing an equity investment.

Valuing an Asset with Default Risk

We will begin a section on how we assess default risk and adjust interest rates for default risk, and then consider how best to value assets with default risk.

Measuring Default Risk and Estimating Default-risk adjusted Rates

When valuing investments where the cash flows are promised, but there is a risk that they might not be delivered, it is no longer appropriate to use the riskless rate as the discount rate. The appropriate discount rate here will include the riskless rate and an appropriate premium for the default risk called a **default spread**. In chapter 3, we examined how default risk is assessed by ratings agencies and the magnitude of the default spread. It is worth noting that even in the absence of bond ratings, lenders still assess default risk and charge default spreads.

Valuing an Asset with Default Risk

The most common example of an asset with just default risk is a corporate bond, since even the largest, safest companies still have some risk of default. When valuing a corporate bond, we generally make two modifications to the bond valuation approach we developed earlier for a default-free bond. First, we will discount the coupons on the corporate bond, even though these no longer represent expected cash flows, but are instead promised cash flows[1]. Second, the discount rate used for a bond with default risk will be higher than that used for default-free bond. Furthermore, as the default risk increases, so will the discount rate used:

$$\text{Value of Corporate Coupon Bond} = \sum_{t=1}^{t=N} \frac{\text{Coupon}}{(1+\text{k}_\text{d})^t} + \frac{\text{Face Value of the Bond}}{(1+\text{k}_\text{d})^N}$$

---

[1] When you buy a corporate bond with a coupon rate of 8%, you are promised a payment of 8% of the face value of the bond each period, but the payment may be lower or non-existent, if the company defaults.

where k_d is the market interest rate given the default risk.

Consider, for instance a bond issued by Boeing with a coupon rate of 8.75%, maturing in 35 years. Based upon its default risk (measured by a bond rating assigned to Boeing by Standard and Poor's at the time of this analysis), the market interest rate on Boeing's debt is 0.5% higher than the treasury bond rate of 5.5% for default-free bonds of similar maturity. The price of the bond can be estimated as follows:

$$\text{Price of Boeing bond} = \sum_{t=0.5}^{t=35} \frac{43.875}{(1.06)^t} + \frac{1,000}{(1.06)^{35}} = \$1,404.25$$

The coupons were assumed to be semi-annual and the present value was estimated using the annuity equation. Note that the default risk on the bond is reflected in the interest rate used to discount the expected cash flows on the bond. If Boeing's default risk increases, the price of the bond will drop to reflect the higher market interest rate.

Valuing an Asset with Equity Risk

Having valued assets with guaranteed cash flows and those with only default risk, let us now consider the valuation of assets with equity risk. We will begin with the introduction to the way we estimate cash flows and consider equity risk in investments with equity risk, and then we look at how best to value these assets.

Measuring Cash Flows for an Asset with Equity Risk

Unlike the bonds that we have valued so far in this chapter, the cash flows on assets with equity risk are not promised cash flows. Instead, the valuation is based upon the *expected cash flows* on these assets over their lives. We will consider two basic questions: the first relates to how we measure these cash flows, and the second to how to come up with expectations for these cash flows.

To estimate cash flows on an asset with equity risk, let us first consider the perspective of the owner of the asset, i.e. the equity investor in the asset. Assume that the owner borrowed some of the funds needed to buy the asset. The cash flows to the owner will therefore be the cash flows generated by the asset after all expenses and taxes, and also after payments due on the debt. This cash flow, which is after debt payments, operating expenses and taxes, is called the **cash flow to equity investors**. There is also a broader definition of cash flow that we can use, where we look at not just the equity investor in the asset, but at the total cash flows generated by the asset for both the equity investor and the lender. This cash flow, which is before debt payments but after operating expenses and taxes, is called the **cash flow to the firm** (where the firm is considered to include both debt and equity investors).

Note that, since this is a risky asset, the cash flows are likely to vary across a broad range of outcomes, some good and some not so positive. To estimate the expected cash flow, we consider all possible outcomes in each period, weight them by their relative probabilities[2] and arrive at an expected cash flow for that period.

Measuring Equity Risk and Estimate Risk-Adjusted Discount Rates

When we analyzed bonds with default risk, we argued that the interest rate has to be adjusted to reflect the default risk. This default-risk adjusted interest rate can be considered the **cost of debt** to the investor or business borrowing the money. When analyzing investments with equity risk, we have to make an adjustment to the riskless rate to arrive at a discount rate, but the adjustment will be to reflect the equity risk rather than the default risk. Furthermore, since there is no longer a promised interest payment, we will term this rate a risk-adjusted discount rate rather than an interest rate. We label this adjusted discount rate the **cost of equity**.

A firm can be viewed as a collection of assets, financed partly with debt and partly with equity. The composite cost of financing, which comes from both debt and equity, is a weighted average of the costs of debt and equity, with the weights depending upon how much of each financing is used. This cost is labeled the **cost of capital**.

For instance, assume that Boeing has a cost of equity of 10.54% and a cost of debt of 3.58%. Assume also that it raised 80% of its financing from equity and 20% from debt. Its cost of capital would then be

$$\text{Cost of Capital} = 10.58\% (.80) + 3.58\% (.20) = 9.17\%$$

Thus, for Boeing, the cost of equity is 10.54% while the cost of capital is only 9.17%.

If the cash flows that we are discounting are cash flows to equity investors, as defined in the previous section, the appropriate discount rate is the cost of equity. If the cash flows are prior to debt payments and therefore to the firm, the appropriate discount rate is the cost of capital.

Valuing an Asset with Equity Risk and Finite Life

Most assets that firms acquire have finite lives. At the end of that life, the assets are assumed to lose their operating capacity, though they might still preserve some value. To illustrate, assume that you buy an apartment building and plan to rent the apartments out to earn income. The building will have a finite life, say 30 to 40 years, at the end of which it

---

[2] Note that in many cases, though we might not explicitly state probabilities and outcomes, we are implicitly doing so, when we use expected cash flows.

will have to be torn down and a new building constructed, but the land will continue to have value even if this occurs.

This building can be valued using the cash flows that it will generate, prior to any debt payments, and discounting them at the composite cost of the financing used to buy the building, i.e. , the cost of capital. At the end of the expected life of the building, we estimate what the building (and the land it sits on) will be worth and discount this value back to the present, as well. In summary, the value of a finite life asset can be written as:

$$\text{Value of Finite - Life Asset} = \sum_{t=1}^{t=N} \frac{\text{E(Cash flow on Asset}_t)}{(1+k_c)^t} + \frac{\text{Value of Asset at End of Life}}{(1+k_c)^N}$$

where k_c is the cost of capital.

This entire analysis can also be done from your perspective as the sole equity investor in this building. In this case, the cash flows will be defined more narrowly as cash flows after debt payments, and the appropriate discount rate becomes the cost of equity. At the end of the building's life, we still look at how much it will be worth but consider only the cash that will be left over after any remaining debt is paid off. Thus, the value of the equity investment in an asset with a fixed life of N years, say an office building, can be written as follows:

$$\text{Value of Equity in Finite - Life Asset} = \sum_{t=1}^{t=N} \frac{\text{E(Cash Flow to Equity}_t)}{(1+k_e)^t} + \frac{\text{Value of Equity in Asset at End of Life}}{(1+k_e)^N}$$

where k_e is the rate of return that the equity investor in this asset would demand given the riskiness of the cash flows and the value of equity at the end of the asset's life is the value of the asset net of the debt outstanding on it. Can you extend the life of the building by reinvesting more in maintaining it? Possibly. If you choose this course of action, however, the life of the building will be longer, but the cash flows to equity and to the firm each period have to be reduced[3] by the amount of the reinvestment needed for maintenance.

To illustrate these principles, assume that you are trying to value a rental building for purchase. The building is assumed to have a finite life of 12 years and is expected to have cash flows *before debt payments* of $ 1 million, growing at 5% a year for the next 12 years. The real estate is also expected to have a value of $ 2.5 million at the end of the 12th year (called the salvage value). Based upon your costs of borrowing and the cost you attach to

---

[3] By maintaining the building better, you might also be able to charge higher rents, which may provide an offsetting increase in the cash flows.

the equity you will have invested in the building, you estimate a cost of capital of 9.51%. The value of the building can be estimated in Table 4.4:

Table 4.4: Value of Rental Building

Year	*Expected Cash Flows*	*Value at End*	*PV at 9.51%*
1	$ 1,050,000		$ 958,817
2	$ 1,102,500		$ 919,329
3	$ 1,157,625		$ 881,468
4	$ 1,215,506		$ 845,166
5	$ 1,276,282		$ 810,359
6	$ 1,340,096		$ 776,986
7	$ 1,407,100		$ 744,987
8	$ 1,477,455		$ 714,306
9	$ 1,551,328		$ 684,888
10	$ 1,628,895		$ 656,682
11	$ 1,710,339		$ 629,638
12	$ 1,795,856	$ 2,500,000	$ 1,444,124
		Value of Store =	$ 10,066,749

Note that the cash flows over the next 12 years represent a growing annuity, and the present value could have been computed with a simple present value equation, as well.

$$\text{Value of Building} = \frac{1{,}000{,}000\ (1.05)(1 - \frac{(1.05)^{12}}{(1.0951)^{12}}}{(.0951 - .05)} + \frac{2{,}500{,}000}{(1.0951)^{12}} = \$10{,}066{,}749$$

This building has a value of $10.07 million to you.

Now, consider the equity investment in the rental building described above. Assume that the cash flows from the building after debt payments are expected will be $ 850,000 a year, growing at 5% a year for the next 12 years. In addition, assume that the salvage value of the building, after repaying remaining debt will be $ 1 million at the end of the 12th year. Finally, assume that your cost of equity is 9.78%. The value of equity in this building can be estimated as follows:

$$\text{Value of Equity in Building} = \frac{850{,}000\ (1.05)\left(1 - \frac{(1.05)^{12}}{(1.0978)^{12}}\right)}{(.0978 - .05)} + \frac{1{,}000{,}000}{(1.0978)^{12}} = \$8{,}053{,}999$$

Note that the value of equity in the building is also an increasing function of expected growth and the building's life, and a decreasing function of the cost of equity.

Valuing an Asset with an Infinite Life

When we value businesses and firms, as opposed to individual assets, we are often looking at entities that have no finite life. If they reinvest sufficient amounts in new assets each period, firms could keep generating cash flows forever. In this section, we value assets that have infinite lives and uncertain cash flows.

Equity and Firm Valuation

In the section on valuing assets with equity risk, we introduced the notions of cash flows to equity and cash flows to the firm. We argued that cash flows to equity are cash flows after debt payments, all expenses and reinvestment needs have been met. In the context of a business, we will use the same definition to measure the cash flows to its equity investors. These cash flows, when discounted back at the cost of equity for the business, yields the value of the equity in the business. This is illustrated in Figure 4.6:

Figure 4.6: Equity Valuation

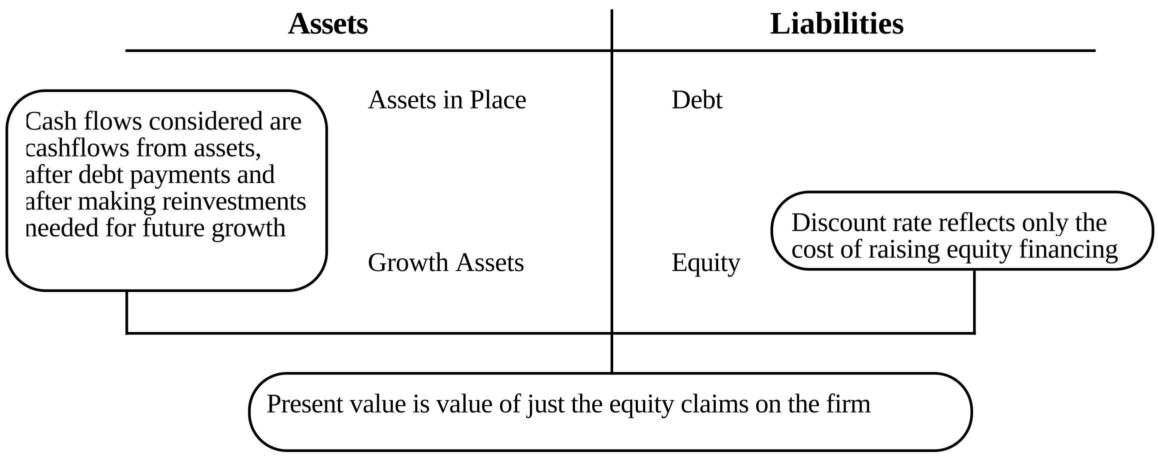

Note that our definition of both cash flows and discount rates is consistent – they are both defined in terms of the equity investor in the business.

There is an alternative approach in which, instead of valuing the equity stake in the asset or business, we look at the value of the entire business. To do this, we look at the collective cash flows not just to equity investors but also to lenders (or bondholders in the firm). The appropriate discount rate is the cost of capital, since it reflects both the cost of equity and the cost of debt. The process is illustrated in Figure 4.7.

Figure 4.7: Firm Valuation

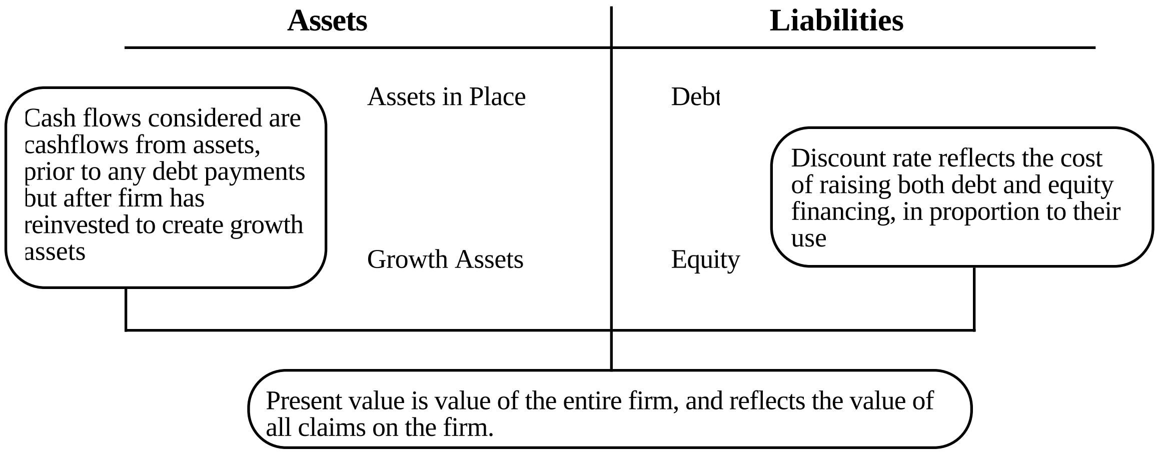

Note again that we are defining both cash flows and discount rates consistently, to reflect the fact that we are valuing not just the equity portion of the investment but the investment itself.

Dividends and Equity Valuation

When valuing equity investments in publicly traded companies, we could argue that the only cash flows investors in these investments get from the firm are dividends. Therefore, the value of the equity in these investments can be computed as the present value of expected dividend payments on the equity.

$$\text{Value of Equity (Only Dividends)} = \sum_{t=1}^{t=\infty} \frac{E(\text{Dividend}_t)}{(1+k_e)^t}$$

The mechanics are similar to those involved in pricing a bond, with dividend payments replacing coupon payments, and the cost of equity replacing the interest rate on the bond. The fact that equity in a publicly traded firm has an infinite life, however, indicates that we cannot arrive at closure on the valuation without making additional assumptions.

One way in which we might be able to estimate the value of the equity in a firm is by assuming that the dividends, starting today, will grow at a constant rate forever. If we do that, we can estimate the value of the equity using the present value formula for a perpetually growing cash flow in chapter 3. In fact, the value of equity will be

$$\text{Value of Equity (Dividends growing at a constant rate forever)} = \frac{E(\text{Dividend next period})}{(k_e - g_n)}$$

This model, which is called the **Gordon growth model**, is simple but limited, since it can value only companies that pay dividends, and only if these dividends are expected to grow at a constant rate forever. The reason this is a restrictive assumption is that no asset or firm's cash flows can grow forever at a rate higher than the growth rate of the economy. If it did, the firm would become the economy. Therefore, the constant growth rate is constrained to be less than or equal to the economy's growth rate. For valuations of firms in US dollars, this puts an upper limit on the growth rate of approximately 5-6%[4]. This constraint will also ensure that the growth rate used in the model will be less than the discount rate.

We will illustrate this model using Consolidated Edison, the utility that produces power for much of New York city, paid dividends per share of $ 2.12 in 1998. The dividends are expected to grow 5% a year in the long term, and the company has a cost of equity of 9.40%. The value per share can be estimated as follows:

Value of Equity per share = $2.12 (1.05) / (.094 - .05) = $ 50.59

The stock was trading at $ 54 per share at the time of this valuation. We could argue that based upon this valuation, the stock was mildly overvalued.

What happens if we have to value a stock whose dividends are growing at 15% a year? The solution is simple. We value the stock in two parts. In the first part, we estimate the expected dividends each period for as long as the growth rate of this firm's dividends remains higher than the growth rate of the economy, and sum up the present value of the dividends. In the second part, we assume that the growth rate in dividends will drop to a stable or constant rate forever sometime in the future. Once we make this assumption, we can apply the Gordon growth model to estimate the present value of all dividends in stable growth. This present value is called the **terminal price** and represents the expected value of the stock in the future, when the firm becomes a stable growth firm. The present value of this terminal price is added to the present value of the dividends to obtain the value of the stock today.

$$\text{Value of Equity with high-growth dividends} = \sum_{t=1}^{t=N} \frac{\text{E(Dividends}_t)}{(1+k_e)^t} + \frac{\text{Terminal Price}_N}{(1+k_e)^N}$$

where N is the number of years of high growth and the terminal price is based upon the assumption of stable growth beyond year N.

$$\text{Terminal Price} = \frac{\text{E(Dividend}_{N+1})}{(k_e - g_n)}$$

To illustrate this model, assume that you were trying to value Coca Cola. The company paid $0.69 as dividends per share during 1998, and these dividends are expected

---

[4] The nominal growth rate of the US economy through the nineties has been about 5%. The growth rate of the global economy, in nominal US dollar terms, has been about 6% over that period.

to grow 25% a year for the next 10 years. Beyond that, the expected growth rate is expected to be 6% a year forever. Assuming a cost of equity of 11% for Coca Cola, we can estimate the value of the stock in two parts and then estimate its value today.

I. Estimate the value of expected dividends during the next 10 years

The expected dividends during the high growth phase are estimated in the Table 4.5. The present values of the dividends are estimated using the cost of equity of 11% in the last column.

Table 4.5: Value of Expected Dividends during High-Growth Phase

Year	Dividends per Share	Present Value
1	$ 0.86	$ 0.78
2	$ 1.08	$ 0.88
3	$ 1.35	$ 0.99
4	$ 1.68	$ 1.11
5	$ 2.11	$ 1.25
6	$ 2.63	$ 1.41
7	$ 3.29	$ 1.58
8	$ 4.11	$ 1.78
9	$ 5.14	$ 2.01
10	$ 6.43	$ 2.26
	PV of Dividends	$ 14.05

II. Estimate the terminal value of the stock at the end of the high growth phase

To estimate the terminal price, we first estimate the dividends per share one year past the high growth phase and use the perpetual growth equation to compute present value. For Coca Cola, the estimates are as follows:

> *ddmginzu.xls*: See the spreadsheet that contains the valuation of Coca Cola.

Expected Dividends per share in year 11 = $ 6.43 *1.06 = $ 6.81

Expected Terminal Price = $ 6.81 / (.11 - .06) = $ 136.24

III. Estimate the value of the stock today

To estimate the value of the stock today, we add the present value of the terminal price estimated in the previous step to the present value of the dividends during the high growth period:

Value of Stock today = PV of Dividends in high growth + PV of Terminal Price

$$= \$\,14.05 + \$\,136.24/(1.11)^{10} = \$62.03$$

A Broader Measure of Cash Flows to Equity

There are two significant problems with the use of just dividends to value equity. The first is that it works only cash flows to the equity investors take the form of dividends.

It will not work for valuing equity in private businesses, where the owners often withdraw cash from the business but may not call it dividends, and it may not even work for publicly traded companies if they return cash to the equity investors by buying back stock, for instance. The second problem is that the use of dividends is based upon the assumption that firms pay out what they can afford to in dividends. When this is not true, the dividend discount models will mis-estimate the value of equity.

To counter this problem, we consider a broader definition of cash flow to which we call **free cash flow to equity**, defined as the cash left over after operating expenses, interest expenses, net debt payments and reinvestment needs. By **net debt payments**, we are referring to the difference between new debt issued and repayments of old debt. If the new debt issued exceeds debt repayments, the free cash flow to equity will be higher.

Free Cash Flow to Equity (FCFE) = Net Income – Reinvestment Needs – (Debt Repaid – New Debt Issued)

Think of this as potential dividends, or what the company could have paid out in dividend. To illustrate, in 1998, the Home Depot's free cash flow to equity using this definition was:

$$FCFE_{Boeing} = \text{Net Income} - \text{Reinvestment Needs} - (\text{Debt Repaid} - \text{New Debt Issued})$$
$$= \$\,1{,}614 \text{ million} - \$1{,}876 \text{ million} - (8 - 246 \text{ million}) = -\,\$\,24 \text{ million}$$

Clearly, the Home Depot did not generate positive cash flows after reinvesment needs and net debt payments. Surprisingly, the firm did pay a dividend, albeit a small one. Any dividends paid by the Home Depot during 1998 had to be financed with existing cash balances, since the free cash flow to equity is negative.

fcfeginzu.xls: See the spreadsheet that contains the valuation of the Home Depot

Once the free cash flows to equity have been estimated, the process of estimating value parallels the dividend discount model. To value equity in a firm where the free cash flows to equity are growing at a constant rate forever, we use the present value equation to estimate the value of cash flows in perpetual growth:

$$\text{Value of Equity in Infinite - Life Asset} = \frac{E(FCFE_t)}{(k_e - g_n)}$$

All the constraints relating to the magnitude of the constant growth rate used that we discussed in the context of the dividend discount model, continue to apply here.

In the more general case, where free cash flows to equity are growing at a rate higher than the growth rate of the economy, the value of the equity can be estimated again in two parts. The first part is the present value of the free cash flows to equity during the high growth phase, and the second part is the present value of the terminal value of equity, estimated based on the assumption that the firm will reach stable growth sometime in the future.

$$\text{Value of Equity with high growth FCFE} = \sum_{t=1}^{t=N} \frac{E(FCFE_t)}{(1+k_e)^t} + \frac{\text{Terminal Value of Equity}_N}{(1+k_e)^N}$$

With the FCFE approach, we have the flexibility we need to value equity in any type of business or publicly traded company.

Consider the case of the Home Depot. Assume that we expect the free cash flows to equity at the firm to become positive next period and to grow for the next 10 years at rates much higher than the growth rate for the economy. To estimate the free cash flows to equity for the next 10 years, we make the following assumptions:

- The net income of $1,614 million will grow 15% a year each year for the next 10 years.
- The firm will reinvest 75% of the net income back into new investments each year, and its net debt issued each year will be 10% of the reinvestment.

Table 4.6 summarizes the free cash flows to equity at the firm for this period and computes the present value of these cash flows at the Home Depot's cost of equity of 9.78%.

Table 4.6: Value of FCFE

Year	*Net Income*	*Reinvestment Needs*	*Net Debt Issued*	*FCFE*	*PV of FCFE*
1	$ 1,856	$ 1,392	$ (139)	$ 603	$ 549
2	$ 2,135	$ 1,601	$ (160)	$ 694	$ 576
3	$ 2,455	$ 1,841	$ (184)	$ 798	$ 603
4	$ 2,823	$ 2,117	$ (212)	$ 917	$ 632
5	$ 3,246	$ 2,435	$ (243)	$ 1,055	$ 662
6	$ 3,733	$ 2,800	$ (280)	$ 1,213	$ 693
7	$ 4,293	$ 3,220	$ (322)	$ 1,395	$ 726
8	$ 4,937	$ 3,703	$ (370)	$ 1,605	$ 761
9	$ 5,678	$ 4,258	$ (426)	$ 1,845	$ 797
10	$ 6,530	$ 4,897	$ (490)	$ 2,122	$ 835
Sum of PV of FCFE =					$6,833

Note that since more debt is issued than paid, net debt issued increases the free cash flows to equity each year. To estimate the terminal price, we assume that net income will grow 6% a year forever after year 10. Since lower growth will require less reinvestment, we will assume that the reinvestment rate after year 10 will be 40% of net income; net debt issued will remain 10% of reinvestment.

$FCFE_{11}$ = Net Income_{11} – Reinvestment_{11} – Net Debt Paid $(\text{Issued})_{11}$

= $6,530 (1.06) – $6,530 (1.06) (0.40) – (-277) = $ 4,430 million

Terminal Price_{10} = $FCFE_{11}/(k_e - g)$ = $ 4,430 / (.0978 - .06) = $117,186 million

The value per share today can be computed as the sum of the present values of the free cash flows to equity during the next 10 years and the present value of the terminal value at the end of the 10^{th} year.

Value of the Stock today = \$ 6,833 million + \$ 117,186/$(1.0978)^{10}$ = \$52,927 million

On a free cash flow to equity basis, we would value the equity at the Home Depot at \$ 52.93 billion.

From Valuing Equity to Valuing the Firm

A firm is more than just its equity investors. It has other claim holders, including bondholders and banks. When we value the firm, therefore, we consider cash flows to all of these claim holders. We define the **free cash flow to the firm** as being the cash flow left over after operating expenses, taxes and reinvestment needs, but before any debt payments (interest or principal payments).

Free Cash Flow to Firm (FCFF) = After-tax Operating Income – Reinvestment Needs

The two differences between FCFE and FCFF become clearer when we compare their definitions. The free cash flow to equity begins with net income, which is after interest expenses and taxes, whereas the free cash flow to the firm begins with after-tax operating income, which is before interest expenses. Another difference is that the FCFE is after net debt payments, whereas the FCFF is before net debt.

What exactly does the free cash flow to the firm measure? On the one hand, it measures the cash flows generated by the assets before any financing costs are considered and thus is a measure of operating cash flow. On the other, the free cash flow to the firm is the cash flow used to service all claim holders' needs for cash – interest and principal to debt holders and dividends and stock buybacks to equity investors.

To illustrate the estimation of free cash flow to the firm, consider Boeing in 1998. In that year, Boeing had adjusted operating income of \$ 2,736 million, a tax rate of 35% and reinvested \$1,719 million in new investments. The free cash flow to the firm for Boeing in 1998 is then:

$$FCFF_{Boeing} = \text{Operating Income (1- Tax Rate)} - \text{Reinvestment Needs}$$

$$= \$\ 2{,}736\ (1-.35) - \$\ 1{,}719 \text{ million} = \$\ 59 \text{ million}$$

Once the free cash flows to the firm have been estimated, the process of computing value follows a familiar path. If valuing a firm or business with free cash flows growing at a constant rate forever, we can use the perpetual growth equation:

$$\text{Value of Firm with FCFF growing at constant rate} = \frac{E(FCFF_1)}{(k_c - g_n)}$$

There are two key distinctions between this model and the constant-growth FCFE model used earlier. The first is that we consider cash flows before debt payments in this model, whereas we used cash flows after debt payments when valuing equity. The second is that we then discount these cash flows back at a composite cost of financing, i.e., the cost of capital to arrive at the value of the firm, while we used the cost of equity as the discount rate when valuing equity.

To value firms where free cash flows to the firm are growing at a rate higher than that of the economy, we can modify this equation to consider the present value of the cash flows until the firm is in stable growth. To this present value, we add the present value of the terminal value, which captures all cash flows in stable growth.

$$\text{Value of high - growth business} = \sum_{t=1}^{t=N} \frac{E(\text{FCFF}_t)}{(1+k_c)^t} + \frac{\text{Terminal Value of Business}_N}{(1+k_c)^N}$$

Assume now that Boeing is interested in selling its information, space and defense systems division. The division reported cash flows before debt payments but after reinvestment needs of $ 393 million in 1998, and the cash flows are expected to grow 5% a year in the long term. The cost of capital for the division is 9%. The division can be valued as follows:

Value of Division = $ 393 (1.05) / (.09 - .05) = $ 10,318 million

You can extend this model to value Boeing as a firm. To do this valuation, assume that Boeing has cash flows before debt payments but after reinvestment needs and taxes of $ 850 million in the current year. Further, assume that these cash flows will grow at 15% a year for the next 5 years and at 5% thereafter. Boeing has a cost of capital of 9.17%. The value of Boeing as a firm can then be estimated in Table 4.7:

fcffginzu.xls: See the spreadsheet that contains the valuation of Boeing as a firm.

Table 4.7: Value of Boeing

Year	Cash Flow	Terminal Value	Present Value
1	$978		$895
2	$1,124		$943
3	$1,293		$994
4	$1,487		$1,047
5	$1,710	$43,049	$28,864
Value of Boeing as a firm =			$32,743

The terminal value is estimated using the free cash flow to the firm in year 6, the cost of capital of 9.17% and the expected constant growth rate of 5% as follows:

Terminal Value = $ 1710 (1.05)/(.0917-.05) = $ 43,049 million

It is then discounted back to the present to get the value of the firm today shown above as $32,743 million.

Note that this is not the value of the equity of the firm. To get to the value of the equity, we would need to subtract out debt from $32,743 million the value of all non-equity claims in the firm.

II. Relative Valuation

In intrinsic valuation the objective is to find assets that are priced below what they should be, given their cash flow, growth and risk characteristics. In relative valuation, the philosophical focus is on finding assets that are cheap or expensive relative to how "similar" assets are being priced by the market right now. It is therefore entirely possible that an asset that is expensive on an intrinsic value basis may be cheap on a relative basis.

A. Standardized Values and Multiples

To compare the valuations of "similar" assets in the market, we need to standardize the values in some way. They can be standardized relative to the earnings that they generate, the book value or replacement value of the assets themselves or relative to the revenues that they generate. Each approach is used widely and has strong adherents.

1. Earnings Multiples

One of the more intuitive ways to think of the value of any asset is as a multiple of the earnings generated by it. When buying a stock, it is common to look at the price paid as a multiple of the earnings per share generated by the company. This *price/earnings ratio* can be estimated using current earnings per share (which is called a trailing PE) or a expected earnings per share in the next year (called a forward PE). When buying a business (as opposed to just the equity in the business) it is common to examine the value of the business as a multiple of the operating income (or EBIT) or the operating cash flow (EBITDA). While a lower multiple is better than a higher one, these multiples will be affected by the growth potential and risk of the business being acquired.

2. Book Value or Replacement Value Multiples

While markets provide one estimate of the value of a business, accountants often provide a very different estimate of the same business in their books. This latter estimate, which is the *book value*, is driven by accounting rules and are heavily influenced by what was paid originally for the asset and any accounting adjustments (such as depreciation) made since. Investors often look at the relationship between the price they pay for a stock and the book value of equity (or net worth) as a measure of how over or undervalued a stock it; the price/book value ratio that emerges can vary widely across sectors, depending again upon the growth potential and the quality of the investments in each. When valuing businesses, this ratio is estimated using the value of the firm and the book value of all assets (rather than just the equity). For those who believe that book value is not a good measure of the true value of the assets, an alternative is to use the replacement cost of the assets; the ratio of the value of the firm to replacement cost is called *Tobin's Q*.

3. Revenue Multiples

Both earnings and book value are accounting measures and are affected by accounting rules and principles. An alternative approach, which is far less affected by these factors, is to look at the relationship between value of an asset and the revenues it generates. For equity investors, this ratio is the *price/sales ratio*, where the market value per share is divided by the revenues generated per share. For firm value, this ratio can be modified as the *value/sales ratio*, where the numerator becomes the total value of the firm. This ratio, again, varies widely across sectors, largely as a function of the profit margins in each. The advantage of these multiples, however, is that it becomes far easier to compare firms in different markets, with different accounting systems at work.

B. The Fundamentals Behind Multiples

One reason commonly given for relative valuation is that it requires far fewer assumptions than does discounted cash flow valuation. In my view, this is a misconception. The difference between discounted cash flow valuation and relative valuation is that the assumptions that an analyst makes have to be made explicit in the former and they can remain implicit in the latter. It is important that we know what the variables are that drive multiples, since these are the variables we have to control for when comparing these multiples across firms.

To look under the hood, so to speak, of equity and firm value multiples, we will go back to fairly simple discounted cash flow models for equity and firm value and use them to derive our multiples. Thus, the simplest discounted cash flow model for equity which is a stable growth dividend discount model would suggest that the value of equity is:

$$\text{Value of Equity} = P_0 = \frac{DPS_1}{k_e - g_n}$$

where DPS_1 is the expected dividend in the next year, k_e is the cost of equity and g_n is the expected stable growth rate. Dividing both sides by the earnings, we obtain the discounted cash flow model for the PE ratio for a stable growth firm:

$$\frac{P_0}{EPS_0} = PE = \frac{\text{Payout Ratio} * (1+g_n)}{k_e - g_n}$$

Dividing both sides by the book value of equity, we can estimate the Price/Book Value ratio for a stable growth firm:

$$\frac{P_0}{BV_0} = PBV = \frac{\text{ROE} * \text{Payout Ratio} * (1+g_n)}{k_e - g_n}$$

where ROE is the return on equity. Dividing by the Sales per share, the price/sales ratio for a stable growth firm can be estimated as a function of its profit margin, payout ratio, profit margin and expected growth.

$$\frac{P_0}{Sales_0} = PS = \frac{\text{Profit Margin} * \text{Payout Ratio} * (1+g_n)}{k_e - g_n}$$

We can do a similar analysis from the perspective of firm valuation. The value of a firm in stable growth can be written as:

$$\text{Value of Firm} = V_0 = \frac{FCFF_1}{k_c - g_n}$$

Dividing both sides by the expected free cash flow to the firm yields the Value/FCFF multiple for a stable growth firm:

$$\frac{V_0}{FCFF_1} = \frac{1}{k_c - g_n}$$

Since the free cash flow the firm is the after-tax operating income netted against the net capital expenditures and working capital needs of the firm, the multiples of EBIT, after-tax EBIT and EBITDA can also be similarly estimated. The value/EBITDA multiple, for instance, can be written as follows:

$$\frac{\text{Value}}{\text{EBITDA}} = \frac{(1-t)}{k_c - g} + \frac{\text{Depr (t)/EBITDA}}{k_c - g} - \frac{\text{CEx/EBITDA}}{k_c - g} - \frac{\Delta \text{ Working Capital/EBITDA}}{k_c - g}$$

The point of this analysis is not to suggest that we go back to using discounted cash flow valuation but to get a sense of the variables that may cause these multiples to vary across firms in the same sector. An analyst who is blind to these variables might conclude that a stock with a PE of 8 is cheaper than one with a PE of 12, when the true reason may by that the latter has higher expected growth, or that a stock with a P/BV ratio of 0.7 is cheaper than one with a P/BV ratio of 1.5, when the true reason may be that the latter has a much higher return on equity. The following table lists out the multiples that are widely used and the variables driving each; the variable, which in my view, is the most significant is highlighted for each multiple. This is what I would call the *companion variable* for this multiple, i.e., the one variable I would need to know in order to use this multiple to find under or over valued assets.

Table 4.8: Multiples and Companion Variables

Companion variables are in bold type

Multiple	*Determining Variables*
Price/Earnings Ratio	*Growth*, Payout, Risk
Price/Book Value Ratio	Growth, Payout, Risk, ***ROE***
Price/Sales Ratio	Growth, Payout, Risk, *Net Margin*
Value/EBIT Value/EBIT (1-t) Value/EBITDA	Growth, ***Reinvestment Needs***, Leverage, Risk
Value/Sales	Growth, Net Capital Expenditure needs, Leverage, Risk, ***Operating Margin***
Value/Book Capital	Growth, Leverage, Risk and ***ROC***

C. The Use of Comparables

Most analysts who use multiples use them in conjunction with “comparable” firms to form conclusions about whether firms are fairly valued or not. At the risk of being simplistic, the analysis begins with two decisions - the multiple that will be used in the analysis and the group of firms that will comprise the comparable firms. The multiple is computed for each of the comparable firms, and the average is computed. To evaluate an individual firm, the analyst then compares its multiple to the average computed; if it is significantly different, the analyst makes a subjective judgment on whether the firm’s individual characteristics (growth, risk ..) may explain the difference. Thus, a firm may have a PE ratio of 22 in a sector where the average PE is only 15, but the analyst may conclude that this difference can be justified by the fact that the firm has higher growth potential than the average firm in the sector. If, in the analysts’ judgment, the difference on the multiple cannot be explained by the fundamentals, the firm will be viewed as over valued (if its multiple is higher than the average) or undervalued (if its multiple is lower than the average).

1. Choosing Comparables

The heart of this process is the selection of the firms that comprise comparable firms. From a valuation perspective, a comparable firm is one with similar cash flows, growth potential and risk. If life were simple, the value of a firm would be analyzed by looking at how an exactly identical firm - in terms of risk, growth and cash flows - is priced. In most analyses, however, a comparable firm is defined to be one in the same business as the firm being analyzed. If there are enough firms in the sector to allow for it, this list will be pruned further using other criteria; for instance, only firms of similar size may be considered. Implicitly, the assumption being made here is that firms in the same sector have similar risk, growth and cash flow profiles and therefore can be compared with much more legitimacy. This approach becomes more difficult to apply under two conditions:

1. There are relatively few firms in a sector. In most markets outside the United States, the number of publicly traded firms in a particular sector, especially if it is defined narrowly, is small.
2. The differences on risk, growth and cash flow profiles across firms within a sector is large. Thus, there may be hundreds of computer software companies listed in the United States, but the differences across these firms are also large.

The tradeoff is therefore a simple one. Defining a sector more broadly increases the number of firms that enter the comparable firm list, but it also results in a more diverse group.

2. Controlling for Differences across Firms

Since it is impossible to find identical firms to the one being valued, we have to find ways of controlling for differences across firms on the relevant ways. The advantage of the discounted cash flow models introduced in the prior section is that we have a clear idea of what the fundamental determinants of each multiple are, and therefore what we should be controlling for; table 1 provides a summary of the variables. The process of controlling for the variables can range from very simple approaches, which modify the multiples to take into account differences on one key variable, to more complex approaches that allow for differences on more than one variable.

Let us start with the simple approaches. Here, the basic multiple is modified to take into account the most important variable determining that multiple. Thus, the PE ratio is divided by the expected growth rate in EPS for a company to come up with a growth-adjusted PE ratio. Similarly, the PBV ratio is divided by the ROE to come up with a value ratio, and the price sales ratio by the net margin. These modified ratios are then compared across companies in a sector. Implicitly, the assumption made is that these firms are comparable on all the other dimensions of value, besides the one being controlled for.

Illustration 4: Comparing PE ratios and growth rates across firms: Software companies

In the following table, we have listed the PE ratios and expected analyst consensus growth rates over 5 years for a selected list of software companies:

Company	*PE*	*Expected Growth Rate*	*PE/Expected Growth (PEG)*
Acclaim Entertainment	13.70	23.60%	0.58
Activision	75.20	40.00%	1.88
Broderbund	32.30	26.00%	1.24
Davidson Associates	44.30	33.80%	1.31
Edmark	88.70	37.50%	2.37
Electronic Arts	33.50	22.00%	1.52
The Learning Co.	33.50	28.80%	1.16
Maxis	73.20	30.00%	2.44
Minnesota Educational	69.20	28.30%	2.45
Sierra On-Line	43.80	32.00%	1.37

While comparisons on the PE ratio alone do not factor in the differences in expected growth, the PEG ratio in the last column can be viewed as growth adjusted PE ratio and that would suggest that Acclaim is the cheapest company in this group and Minnesota Educational is the most expensive. This conclusion holds only if these firms are of equivalent risk, however.

Controlling for more than one variable

When firms vary on more than one dimension, it becomes difficult to modify the multiples to take into account the differences across firms. It is, however, feasible to run regressions of the multiples against the variables and then use these regressions to get predicted values for each firm. This approach works reasonably well when the number of comparable firms is large and the relationship between the multiple and variable is strong. When these conditions do not hold, a few outliers can cause the coefficients to change dramatically and make the predictions much less reliable.

oilcos.xls: See the spreadsheet that contains the relative valuation of oil companies used in this example.

Illustration 5: PBV Ratios and ROE: The Oil Sector

The following table summarizes Price/Book Value ratios of oil companies and reports on their returns on equity and expected growth rates:

Company Name	*P/BV*	*ROE*	*Expected Growth*

Total ADR B	0.90	4.10	9.50%
Giant Industries	1.10	7.20	7.81%
Royal Dutch Petroleum ADR	1.10	12.30	5.50%
Tesoro Petroleum	1.10	5.20	8.00%
Petrobras	1.15	3.37	15%
YPF ADR	1.60	13.40	12.50%
Ashland	1.70	10.60	7%
Quaker State	1.70	4.40	17%
Coastal	1.80	9.40	12%
Elf Aquitaine ADR	1.90	6.20	12%
Holly	2.00	20.00	4%
Ultramar Diamond Shamrock	2.00	9.90	8%
Witco	2.00	10.40	14%
World Fuel Services	2.00	17.20	10%
Elcor	2.10	10.10	15%
Imperial Oil	2.20	8.60	16%
Repsol ADR	2.20	17.40	14%
Shell Transport & Trading ADR	2.40	10.50	10%
Amoco	2.60	17.30	6%
Phillips Petroleum	2.60	14.70	7.50%
ENI SpA ADR	2.80	18.30	10%
Mapco	2.80	16.20	12%
Texaco	2.90	15.70	12.50%
British Petroleum ADR	3.20	19.60	8%
Tosco	3.50	13.70	14%

Since these firms differ on both growth and return on equity, we ran a regression of PBV ratios on both variables:

PBV = -0.11 + 11.22 (ROE) + 7.87 (Expected Growth) R^2 = 60.88%

(5.79) (2.83)

The numbers in brackets are t-statistics and suggest that the relationship between PBV ratios and both variables in the regression are statistically significant. The R-squared indicates the percentage of the differences in PBV ratios that is explained by the independent variables. Finally, the regression itself can be used to get predicted PBV ratios for the companies in the list. Thus, the predicted PBV ratio for Repsol would be:

Predicted PBV_{Repsol} = -0.11 + 11.22 (.1740) + 7.87 (.14) = 2.94
Since the actual PBV ratio for Repsol was 2.20, this would suggest that the stock was undervalued by roughly 25%.

Both approaches described above assume that the relationship between a multiple and the variables driving value are linear. Since this is not necessarily true, it is possible to run non-linear versions of these regressions.

3. Expanding the Comparable Firm Universe

Searching for comparable firms within the sector in which a firm operates is fairly restrictive, especially when there are relatively few firms in the sector or when a firm operates in more than one sector. Since the definition of a comparable firm is not one that is in the same business but one that has the same growth, risk and cash flow characteristics as the firm being analyzed, it is also unclear why we have to stay sector-specific. A software firm should be comparable to an automobile firm, if we can control for differences in the fundamentals.

The regression approach that we introduced in the previous section allows us to control for differences on those variables that we believe cause differences in multiples across firms. Using the minimalist version of the regression equations here, we should be able to regress PE, PBV and PS ratios against the variables that should affect them:

PE = a + b (Growth) + c (Payout ratios) + d (Risk)

PBV = a + b (Growth) + c (Payout ratios) + d (Risk) + e (ROE)

PS = a + b (Growth) + c (Payout ratios) + d (Risk) + e (Margin)

It is, however, possible that the proxies that we use for risk (beta) , growth (expected growth rate) and cash flow (payout) may be imperfect and that the relationship may not be linear. To deal with these limitations, we can add more variables to the regression - e.g., the size of the firm may operate as a good proxy for risk - and use transformations of the variables to allow for non-linear relationships.

The first advantage of this approach over the “subjective” comparison across firms in the same sector described in the previous section is that it does quantify, based upon actual market data, the degree to which higher growth or risk should affect the multiples. It is true that these estimates can be noisy, but this noise is a reflection of the reality that many analysts choose not to face when they make subjective judgments. Second, by looking at all firms in the universe, it allows analysts operating in sectors with relatively few firms in them to make more powerful comparisons. Finally, it gets analysts past the tunnel vision induced by comparing firms within a sector, when the entire sector may be under or over valued.

Valuing an Asset with Contingent Cash Flows (Options)

In general, the value of any asset is the present value of the expected cash flows on that asset. In this section, we will consider an exception to that rule when we will look at assets with two specific characteristics:

- They derive their value from the values of other assets.
- The cash flows on the assets are contingent on the occurrence of specific events.

These assets are called options, and the present value of the expected cash flows on these assets will understate their true value. In this section, we will describe the cash flow characteristics of options, consider the factors that determine their value and examine how best to value them.

Cash Flows on Options

There are two types of options. A call option gives the buyer of the option the right to buy the underlying asset at a fixed price, whereas a put option gives the buyer the right to sell the underlying asset at a fixed price. In both cases, the fixed price at which the underlying asset can be bought or sold is called the **strike or exercise price**.

To look at the payoffs on an option, consider first the case of a call option. When you buy the right to sell an asset at a fixed price, you want the price of the asset to increase above that fixed price. If it does, you make a profit, since you can buy at the fixed price and then sell at the much higher price; this profit has to be netted against the cost initially paid for the option. However, if the price of the asset decreases below the strike price, it does not make sense to exercise your right to buy the asset at a higher price. In this scenario, you lose what you originally paid for the option. Figure 4.8 summarizes the cash payoff at expiration to the buyer of a call option.

Figure 4.8: Payoff on Call Option

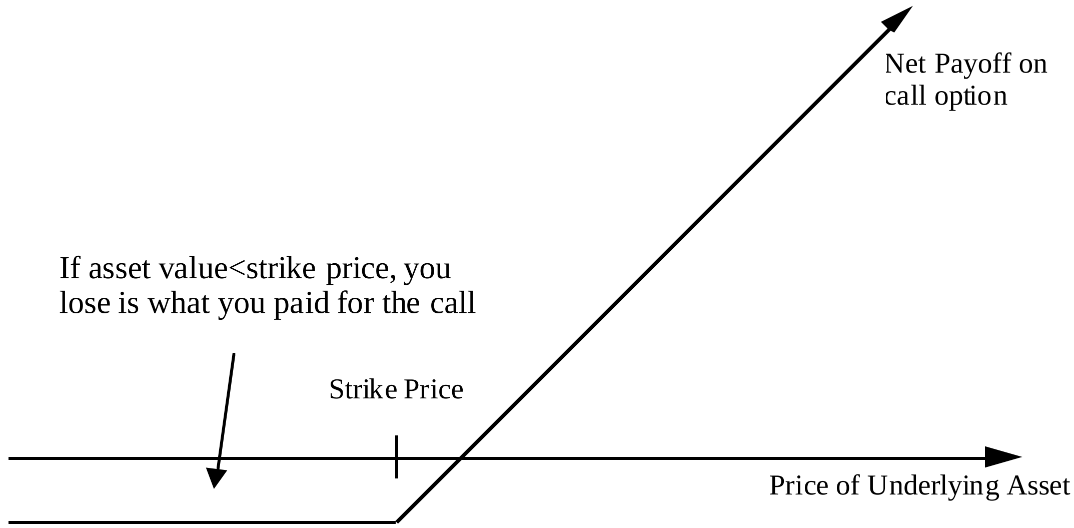

With a put option, you get the right to sell at a fixed price, and you want the price of the asset to decrease below the exercise price. If it does, you buy the asset at the exercise price and then sell it back at the current price, claiming the difference as a gross profit. When the initial cost of buying the option is netted against the gross profit, you arrive at an estimate of the net profit. If the value of the asset rises above the exercise price, you will not exercise the right to sell at a lower price. Instead, the option will be allowed to expire without being exercised, resulting in a net loss of the original price paid for the put option. Figure 4.9 summarizes the net payoff on buying a put option.

Figure 4.9: Payoff on Put Option

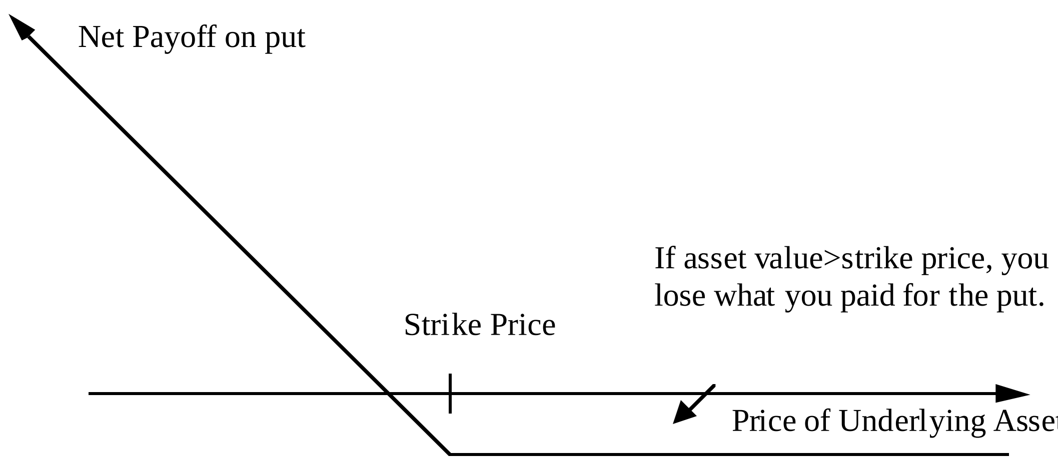

With both call and put options, the potential for profit to the buyer is significant, but the potential for loss is limited to the price paid for the option.

Determinants of Option Value

What is it that determines the value of an option? At one level, options have expected cash flows just like all other assets, and that may seem like good candidates for discounted cash flow valuation. The two key characteristics of options -- that they derive their value from some other traded asset, and the fact that their cash flows are contingent on the occurrence of a specific event -- does suggest an easier alternative. We can create a portfolio that has the same cash flows as the option being valued, by combining a position in the underlying asset with borrowing or lending. This portfolio is called a **replicating portfolio** and should cost the same amount as the option. The principle that two assets (the option and the replicating portfolio) with identical cash flows cannot sell at different prices is called the **arbitrage principle**.

Options are assets that derive value from an underlying asset; increases in the value of the underlying asset will increase the value of the right to buy at a fixed price and reduce the value to sell that asset at a fixed price. On the other hand, increasing the strike price will reduce the value of calls and increase the value of puts.

While calls and puts move in opposite directions when stock prices and strike prices are varied, they both increase in value as the life of the option and the variance in the underlying asset's value increases. The reason for this is the fact that options have limited losses. Unlike traditional assets that tend to get less valuable as risk is increased, options become more valuable as the underlying asset becomes more volatile. This is so because the added variance cannot worsen the downside risk (you still cannot lose more than what you paid for the option) while making potential profits much higher. In addition, a longer life for the options just allows more time for both call and put options to appreciate in value. Since calls provide the right to buy the underlying asset at a fixed price, an increase in the value of the asset will increase the value of the calls. Puts, on the other hand, become less valuable as the value of the asset increase.

The final two inputs that affect the value of the call and put options are the riskless interest rate and the expected dividends on the underlying asset. The buyers of call and put options usually pay the price of the option up front, and wait for the expiration day to exercise. There is a present value effect associated with the fact that the promise to buy an asset for $ 1 million in 10 years is less onerous than paying it now. Thus, higher interest rates will generally increase the value of call options (by reducing the present value of the price on exercise) and decrease the value of put options (by decreasing the present value ofthe price received on exercise). The expected dividends paid by assets make them less valuable; thus, the call option on a stock that does not pay a dividend should be worth more than a call option on a stock that does pay a dividend. The reverse should be true for put options.

Conclusion

In this chapter, we lay the foundations for the models that we will be using to value both assets and firms in the coming chapters. There are three classes of valuation models. The more general of these models, discounted cash flow valuation, can be used to value any asset with expected cash flows over its life. The value is the present value of the expected cash flows at a discount rate that reflects the riskiness of the cash flows, and this principle applies whether one is looking at a zero-coupon government bond or equity in high risk firms. The second set of models are relative valuation models, where we value assets based upon how similar assets are priced by the market. There are some assets that generate cash flows only in the event of a specified contingency, and these assets will not be valued accurately using discounted cash flow models. Instead, they should be viewed as options and valued using option pricing models.

Lessons for Investors

1. All assets that generate or are expected to generate cashflows can be valued by discounting the expected cash flows back at a rate that reflects the riskiness of the cashflows – more risky cash flows should be discounted at higher rates.
2. The value of an on-going business is a function of four variables – how much the business generates in cashflows from existing investments, how long these cashflows can be expected to grow at a rate higher than the growth rate of the economy (high growth period), the level of the growth rate during this period and the riskiness of the cashflows. Companies with higher cashflows, higher growth rates, longer high-growth periods and lower risk will have higher values.
3. Alternatively, assets can be valued by looking at how similar assets are priced in the market. This approach is called relative valuation and is built on the presumption that the market is correct, on average.
4. Assets whose cashflows are contingent on the occurrence of specific events are called options and can be valued using option pricing models.

CHAPTER 5

MANY A SLIP: TRADING, EXECUTION AND TAXES

As investors consider different investment strategies, they have to keep two important issues in mind – trading costs and taxes. It costs to trade, and some strategies create larger trading costs than others. The cost of trading clearly imposes a drag on the performance of all investors and turns otherwise winning portfolios into losing portfolios. As we debate the extent of these costs, we need to get a measure of what the costs are, how they vary across investment strategies and how investors can minimize these costs. In this chapter, we will take an expansive view of trading costs and argue that the brokerage costs (which is what many investors consider as the only trading cost) is only one component (and often the smallest) of trading costs. We will also look at the trading costs associated with holding real assets (such as real estate) and non-traded investments (like equity in a private business) We will also discuss the trade off between trading costs and trading speed and how to devise ways of keeping the trading costs low.

There is a second equally important element in investment success. Investors get to take home after-tax returns and not before tax returns. Thus, strategies that perform exceptionally well before taxes may be money losers if considered after tax. Taxes are particularly difficult to deal with, partly because they are investor and investment specific (different investors have different tax rates) and partly because the tax code itself changes over time, often in unpredictable ways. We will consider the evidence that has accumulated that mutual funds have done their investors a disservice by not considering taxes and that the after-tax returns lag pre-tax returns considerably. We will also look at ways in which we can adjust our investment strategies to keep our tax liabilities low.

The Trading Cost Drag

While we debate what constitutes trading costs and how to measure them, there is a fairly simple way in which we can estimate, at the minimum, how much trading costs affect the returns of the average portfolio manager. Active money managers trade because they believe that there is profit in trading, and the return to any active money manager has three ingredients to it:

Return on active money manager = Expected Return$_{Risk}$ + Return from active trading - Trading costs

Looking across all active money managers, we can reasonably assume that the average expected return has to be equal to the return on the market index. Thus, subtracting the average return made by active money managers from the return on the index should give us a measure of the payoff to active money management:

Average Return$_{\text{Active Money Managers}}$ - Return on Index = Return from Active Trading - Trading Costs

Here the evidence becomes quite depressing. The average active money manager has underperformed the index in the last decade by about 1%. If we take the view that active trading adds no excess return, on average, the trading costs, at the minimum, should be 1% of the portfolio on an annual basis. If we take the view that active trading does add to the returns, the trading costs will be greater than 1% of the portfolio on an annual basis.

There are also fairly specific examples of real portfolios that have been constructed to replicate hypothetical portfolios, where the magnitude of the trading costs are illustrated starkly. For decades, Value Line has offered advice to individual investors on what stocks to buy and which ones to avoid, and ranked stocks from 1 to 5 based upon their desirability as investments. Studies by academics and practitioners found that Value Line rankings seemed to correlate with actual returns. In 1979, Value Line decided to create a mutual fund that would invest in the stocks that it was recommending to its readers. In figure 5.1, we consider the difference in returns between 1979 and 1991 between the fund that Value Line ran and the paper portfolio that Value Line has used to compute the returns that its stock picks would have had.

Figure 5.1: Value Line - Paper Portfolio versus Real Fund

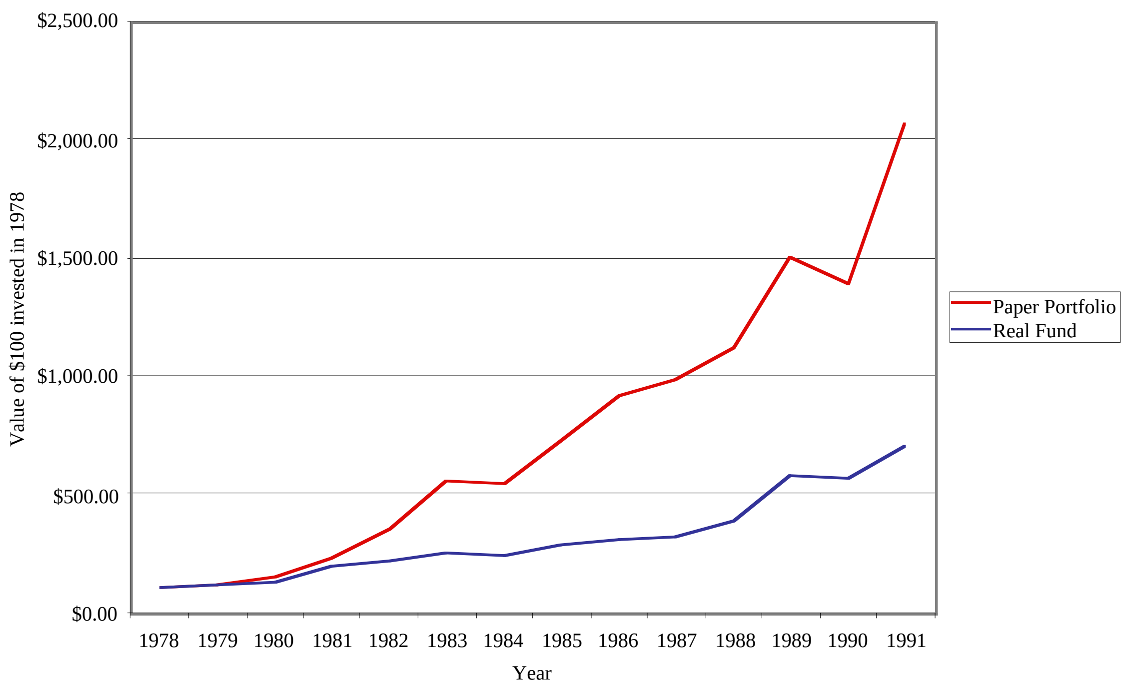

The paper portfolio had an annual return of 26.2%, whereas the Value Line fund had a return of 16.1%. While part of the difference can be attributed to Value Line waiting until its subscribers had a chance to trade, a significant portion of the difference can be explained by the costs of trading.

Looking at the evidence, there are a couple of conclusions that we would draw. The first is that money managers either underestimate trading costs, over estimate the returns to active trading or both. The second is that trading costs are a critical ingredient to any investment strategy, and can make the difference between a successful strategy and an unsuccessful one.

The Components of Trading Costs: Traded Financial Assets

There are some investors who undoubtedly operate under the misconception that the only cost of trading stocks is the brokerage commission that they pay when they buy or sell assets. While this might be the only cost that they pay explicitly, there are other costs that they incur in the course of trading that generally dwarf the commission cost. When trading any asset, they are three other ingredients that go into the trading costs. The first is the spread between the price at which you can buy an asset (the ask price) and the price at which you can sell the same asset at the same point in time (the bid price). The second is the price impact that an investor can create by trading on an asset, pushing the price up when buying the asset and pushing it down while selling. The third cost, which was first proposed by Jack Treynor in his article[1] on transactions costs, is the opportunity cost associated with waiting to trade. While being a patient trader may reduce the first two components of trading cost, the waiting can cost profits both on trades that are made and in terms of trades that would have been profitable if made instantaneously but which became unprofitable as a result of the waiting. It is the sum of these costs that makes up the trading cost on an investment strategy.

The Bid-Ask Spread

There is a difference between that a buyer will pay and the seller will receive, at the same point in time for the same asset, in almost every traded asset market. The bid-ask spread refers to this difference. In the section that follows, we will examine why this difference exists, how large it is as a cost, the determinants of its magnitude and its effects on returns in different investment strategies.

---

[1] This was proposed in his article titled "What does it take to win the trading game?", published in the Financial Analysts Journal, January-February 1981.

Why is there a bid-ask spread?

In most markets, there is a dealer or market maker who sets the bid-ask spread, and there are three types of costs that the dealer faces that the spread is designed to cover. The first is the risk and the cost of holding inventory; the second is the cost of processing orders and the final cost is the cost of trading with more informed investors. The spread has to be large enough to cover these costs and yield a reasonable profit to the market maker on his or her investment in the profession.

1. The Inventory Rationale

Consider a market maker or a specialist on the floor of the exchange who has to quote bid prices and ask prices, at which he is obligated to execute buy and sell orders from investors[2]. These investors, themselves, could be trading because of information they have received (informed traders), for liquidity (liquidity traders) or based upon their belief that an asset is under or over valued (value traders). In such a market, if the market makers set the bid price too high, they will accumulate an inventory of the stock. If market makers set the ask price too low, they will find themselves with a large short position in the stock. In either case, there is a cost to the market makers that they will attempt to recover by increasing the spread between the bid and ask prices.

Market makers also operate with inventory constraints, some of which are externally imposed (by the exchanges or regulatory agencies) and some of which are internally imposed (due to capital limitations and risk). As the market makers' inventory positions deviate from their optimal positions, they bear a cost and will try to adjust the bid and ask prices to get back to their preferred position.

2. The Processing Cost Argument

Since market makers incur a processing cost with the paperwork and fees associated with orders, the bid-ask spread has to cover, at the minimum, these costs. While these costs are likely to be very small for large orders of stocks traded on the exchanges, they become larger for small orders of stocks that might be traded only through a dealership market. Furthermore, since a large proportion of this cost is fixed, these costs as a percentage of the price will generally be higher for low-priced stocks than for high-priced stocks.

Technology clearly has reduced the processing cost associated with trades as computerized systems take over from traditional record keepers. These cost reductions

---

[2] This model was set up by Amihud and Mendelson to explain why bid-ask spreads are different for different firms.

should be greatest for stocks where the bulk of the trades are small trades - small stocks held by individual rather than institutional investors.

3. The Adverse Selection Problem

The adverse selection problem arises from the different motives investors have for trading on an asset - liquidity, information and views on valuation. Since investors do not announce their reasons for trading at the time of the trade, the market maker always runs the risk of trading against more informed investors. Since market makers can expect to lose on such trades, they have to charge an average spread that is large enough to compensate for such losses. This theory would suggest that spreads will increase with the proportion of informed traders in an asset market, the "differential" information possessed, on average, by these traders and uncertainty about future information on the asset.

The Magnitude of the Bid-Ask Spread

The New York Stock Exchange reported[3] that the average bid-ask spread across all NYSE stocks in 1996 was $0.23, which seems trivial especially when one considers the fact that the average price of a NYSE stock is between $ 40 and $ 50. This average, however, obscures the large differences in the cost as a percentage of the price across stocks, based upon capitalization, stock price level and trading volume. A study[4] by Thomas Loeb in 1983, for instance, reported the spread as a percentage of the stock price for companies as a function of their marker capitalization for small orders. These results are summarized in Figure 5.2:

---

[3] See 1996 NYSE Fact Book for a listing of the average spread across all NYSE stocks, by month.

[4] See "Trading Costs: The Critical Link Between Investment Information and Results" in the Financial Analysts Journal, May/June 1983.

Figure 5.2: Prices and Spreads by Market Cap

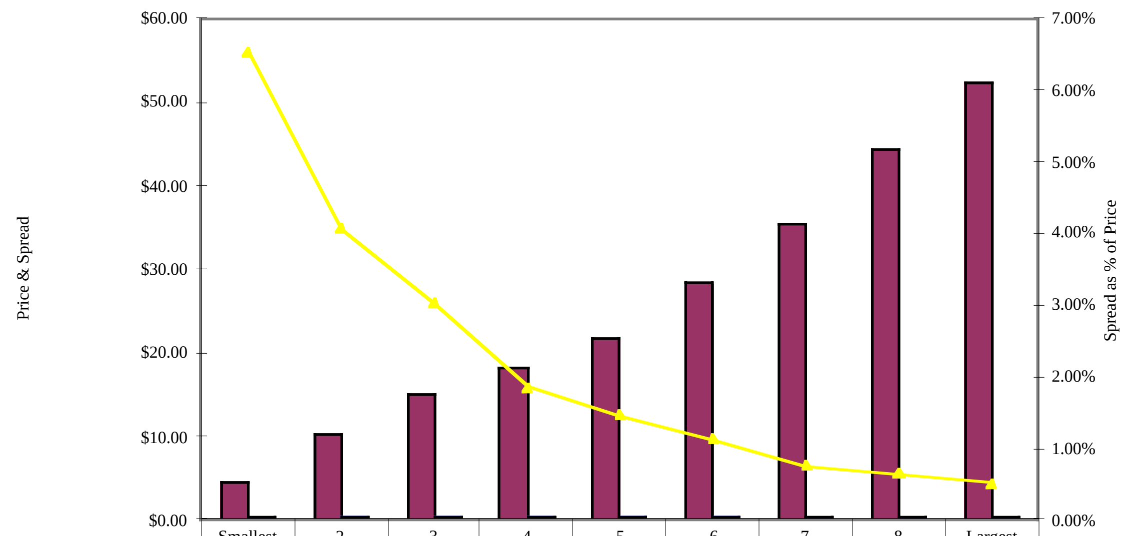

	Smallest	2	3	4	5	6	7	8	Largest
Average Price	$4.58	$10.30	$15.16	$18.27	$21.85	$28.31	$35.43	$44.34	$52.40
Average Spread	$0.30	$0.42	$0.46	$0.34	$0.32	$0.32	$0.27	$0.29	$0.27
Spread/Price	6.55%	4.07%	3.03%	1.86%	1.46%	1.13%	0.76%	0.65%	0.52%

Market Capitalization Class

Average Price   Average Spread   Spread/Price

While the dollar spread is not that different across market capitalization classes, the smallest companies also tend to have lower priced stocks. Consequently, the spread is as high as 6.55% of the price for small capitalization stocks and as low as 0.52% of the price for large capitalization companies. Another study by Huang and Stoll found that the stocks in the top 20% in terms of trading volume had an average spread of only 0.62% of the price while the stocks in the bottom 20% had a spread of 2.06%. There are also large differences in bid-ask spreads across different exchanges in the United States. Looking at only NASDAQ stocks, researchers found[5] that the average was almost 6% of the price in 1992, and much higher for low-priced stocks on the exchange. Some of the difference can be attributed to the fact that NASDAQ stocks are generally much smaller and riskier than stocks listed on the NYSE or AMEX.

> *Stocks with highest Bid-ask Spreads*: Take a look at the 50 stocks with the highest bid-ask spreads.

While these studies looked at traded U.S. equities, there are bid-ask spreads in other markets as well. While no single comprehensive study of all these spreads exists, the following conclusions seem warranted:

---

[5] See "Trading Costs and the Trading Systems for NASDAQ stocks" by M. Kothare and P.A. Laux in Financial Analysts Journal (March/April 1995)

1. The spreads in U.S. government securities are much lower than the spreads on traded stocks in the United States. For instance, the typical bid-ask spread on a Treasury bill is less than 0.1% of the price.
2. The spreads on corporate bonds tend to be larger than the spreads on government bonds, with safer (higher rated) and more liquid corporate bonds having lower spreads than riskier (lower rated) and less liquid corporate bonds.
3. The spreads in non-U.S. equity markets are generally much higher than the spreads on U.S. markets, reflecting the lower liquidity in those markets and the smaller market capitalization of the traded firms.
4. While the spreads in the traded commodity markets are similar to those in the financial asset markets, the spreads in other real asset markets tend to be much larger.

The Determinants of the Bid-Ask Spread

A number of studies have looked at the variables that determine (or, at the very least, correlate with) the bid-ask spread. Studies[6] find that spreads as a percentage of the price are greater for low-priced stocks with higher volatility and lower trading volume; spreads also seem to increase as the number of market makers or dealers in the stock decreases. Each of these findings is consistent with the theory on the bid-ask spread. The negative correlation with price level can be explained by the higher processing cost as a percentage of the price. Higher volume reduces the need for market makers to maintain inventory and also allows them to turn over their inventory rapidly, resulting in lower inventory costs. The higher volatility leads to higher bid-ask spreads partly because the adverse selection problem is greater for more volatile stocks; there will generally be more informed traders, a greater "information differential" and greater uncertainty about future information on these stocks. It is also worth noting that variables such as price level, volatility and trading volume are not only correlated with each other, but are also correlated with other variables such as firm size.

The study quoted in the previous section, by Kothare and Laux, that looked at average spreads on the NASDAQ also looked at differences in bid-ask spreads across stocks on the NASDAQ. In addition to noting similar correlations between the bid-ask spreads, price level and trading volume, they uncovered an interesting new variable. They

---

[6] See "Competition and the Pricing of Dealer Service in the Over-the-Counter Market" by S.Tinic and R. West in Journal of Financial and Quantitative Analysis (June 1972), "The Pricing of Security Dealer Services: An Empirical Analysis of NASDAQ stocks" by H. Stoll in Journal of Finance (November 1978) and "Liquidity Effects of the Introduction of the S&P 500 Futures Contract on the Underlying Stocks" in Journal of Business (April 1993).

found that stocks where institutional activity increased significantly had the biggest increase in bid-ask spreads. While some of this can be attributed to the concurrent increase in volatility in these stocks, it might also reflect the perception on the part of market makers that institutional investors tend to be informed investors with more or better information. Note, though, that institutional investors also increase liquidity which should reduce the order processing cost component of the bid-ask spread, and in some cases the net effect can lead to a lower spread.[7]

Can firms affect the bid-ask spreads on their stocks? There is some evidence that they can by improving the quality of information that they disclose the financial markets, thus reducing the advantages that informed traders may have relative to the rest of the market. In 2001, Heflin. Shaw and Wild looked at 221 firms and examine the relationship between information disclosure quality – they measure this using disclosure quality scores assigned by the Corporate Information Committee of the Financial Analysts Federation – and the bid-ask spread. They find that bid-ask spreads decrease as information quality increases.

While most of the studies quoted above have looked at differences in spreads across stocks, Hasbrouck investigated why spreads change for the same stock at different points in time. He notes that large trades cause spreads to widen, relative to small trades, and hypothesizes that this is because large trades are more likely to contain information.

Market Microstructure and Bid-Ask Spreads

Does the market in which a stock trades matter, when it comes to how big the bid-ask spread should be? Studies indicate that bid-ask spreads have historically been much higher on the NASDAQ than on the New York Stock Exchange, even after controlling for differences in the variables mentioned above – trading volume and price level. In fact, the bid-ask spreads of stocks drop when they switch from the NASDAQ to the NYSE.[8]

A 1994 study by Christie and Schultz provided one explanation for the phenomenon. They found that there were a disproportionately large number of 1/4 quotes and far too few 1/8 quotes.[9] They argued that dealers on the NASDAQ were colluding to set quotes too high and that investors were therefore paying the price with larger bid-ask

---

[7] Dey and Radhakrishna (2001) provide some evidence of this for stocks listed on the NYSE.

[8] See Barclay, M. "Bid-Ask Spreads and the Avoidance of Odd-Eighth Quotes on Nasdaq: An Examination of Exchange Listings." *Journal of Financial Economics*, 45 (1997), 35-60.

[9] If 1/8 and 1/4 quotes are equally likely to show up, roughly half of all quotes should end with an eighth (1/8, 3/8, 5/8 or 7/8) and half should end with a quarter (1/4, 1/2, 3/4).

spreads. This triggered an investigation by the Securities and Exchange Commission (SEC) which agreed that dealers were indeed engaged in anti-competitive behavior. Eventually, the exchange settled the lawsuit for more than a billion dollars. An alternative explanation is that the higher spreads on the NASDAQ, relative to the NYSE, can be explained by structural differences across the markets. Consider, for example, how limit orders are handled on the two exchanges. The specialists on the floor of the New York Stock Exchange are required to reflect in their bid-ask spread the limit prices, if they are better than their own quotes, and this has the effect of reducing the bid-ask spread. On the NASDAQ, limit orders do not affect the bid-ask quotes, and are executed only if prices move against the limit. You would expect larger bid-ask spreads as a consequence.[10]

In 2000, the New York Stock Exchange abandoned its historical practice of quoting prices in fractions (1/8, 1/4… etc) and shifted to decimal prices. Since you can get finer gradations of prices in decimals, it was hypothesized that this should lead to lower bid-ask spreads. Studies since the shift indicate that there has been a decline in spreads on the smaller, less liquid stocks but no discernible impact on the more liquid listings.

Role in Investment Strategies

Looking at the evidence, it is clear that bid-ask spreads will affect the returns from investment strategies, but that the effect will vary, depending upon the strategy. While a strategy of buying under valued companies in the S&P 500 and holding for the long term should not be affected very much by the bid-ask spread, a strategy of buying small over-the-counter stocks or emerging market stocks after information releases, and trading frequently, might lose its allure, when bid-ask spreads are factored into the returns.

To show the effect of the bid-ask spread on returns, consider the strategy of buying "losers". Researchers present evidence[11] that a strategy of buying the stocks that have the most negative returns over the previous year and holding for a five-year period earns significant positive returns. A follow-up study, however, noted that many of these "losers" were low-priced stocks, and that putting in a constraint that the prices be greater than $10 on this strategy resulted in a significant drop in the excess returns. Since bid-ask spreads tend to be largest for low-priced stocks, it is an open question as to whether an investment

---

[10] Chung, Van Ness and Van Ness (2001) tested both explanations. While they find that the treatment of limit orders does lower the bid-ask spread on the NYSE, they conclude that collusion among dealers still leads to higher spreads on the NASDAQ.

[11] See "Does the Stock Market Overreact?" by F.M DeBondt and R. Thaler in Journal of Finance (July 1985)

strategy of buying losers will yield excess returns in practice. In fact, similar concerns should exist about any strategy that recommends investing in low-priced, illiquid and small-cap stocks, or in asset classes that have high volatility and low liquidity.

The Price Impact

Most investors assume that trading costs become smaller as portfolios become larger. While this is true for brokerage commissions, it is not always the case for the other components of trading costs. There is one component where larger investors bear a more substantial cost than do smaller investors and that is in the impact that their trading has on prices. If the basic idea behind successful investing is to buy low and sell high, pushing the price up as you buy and then down as you sell reduces the profits from investing.

Why is there a price impact?

There are two reasons for the price impact, when investors trade. The first is that markets are not completely liquid. A large trade can create an imbalance between buy and sell orders, and the only way in which this imbalance can be resolved is with a price change. This price change that arises from lack of liquidity, will generally be temporary and will be reversed as liquidity returns to the market.

The second reason for the price impact is informational. A large trade attracts the attention of other investors in that market because if might be motivated by new information that the trader possesses. Notwithstanding claims to the contrary, investors usually assume, with good reason, that an investor buying a large block is buying in advance of good news and that an investor selling a large block has come into possession of some bad news about the company. This price effect will generally not be temporary, especially when we look at a large number of stocks where such large trades are made. While investors are likely to be wrong a fair proportion of the time on the informational value of large block trades, there is reason to believe that they will be right almost as often.

How large is the price impact?

There is conflicting evidence on how much of an impact large trades have on stock prices. On the one hand, studies of block trades on the exchange floor seem to suggest that markets are liquid and that the price impact of trading is small and is reversed quickly. These studies, however, have generally looked at heavily traded stocks at the New York Stock exchange. On the other hand, there are others who argue that the price impact is likely to be large, especially for smaller and less liquid stocks.

Studies of the price reaction to large block trades on the floor of the exchange conclude that prices adjust within a few minutes to such trades. An early study examined the speed of the price reaction by looking at the returns an investor could make by buying stock right around the block trade and selling later[12]. They estimated the returns after transactions as a function of how many minutes the acquisition took place after the block trade, and found that only trades made within a few minutes of the block trade had a chance of making excess returns. (See Figure 5.3) Put another way, prices adjusted to the liquidity effects of the block trade within five minutes of the block. While this may be understated because of the fact that these were block trades on large stocks on the NYSE, it is still fairly strong evidence of the capacity of markets to adjust quickly to imbalances between demand and supply.

Figure 5.3: Annualized Returns from buying after block trades

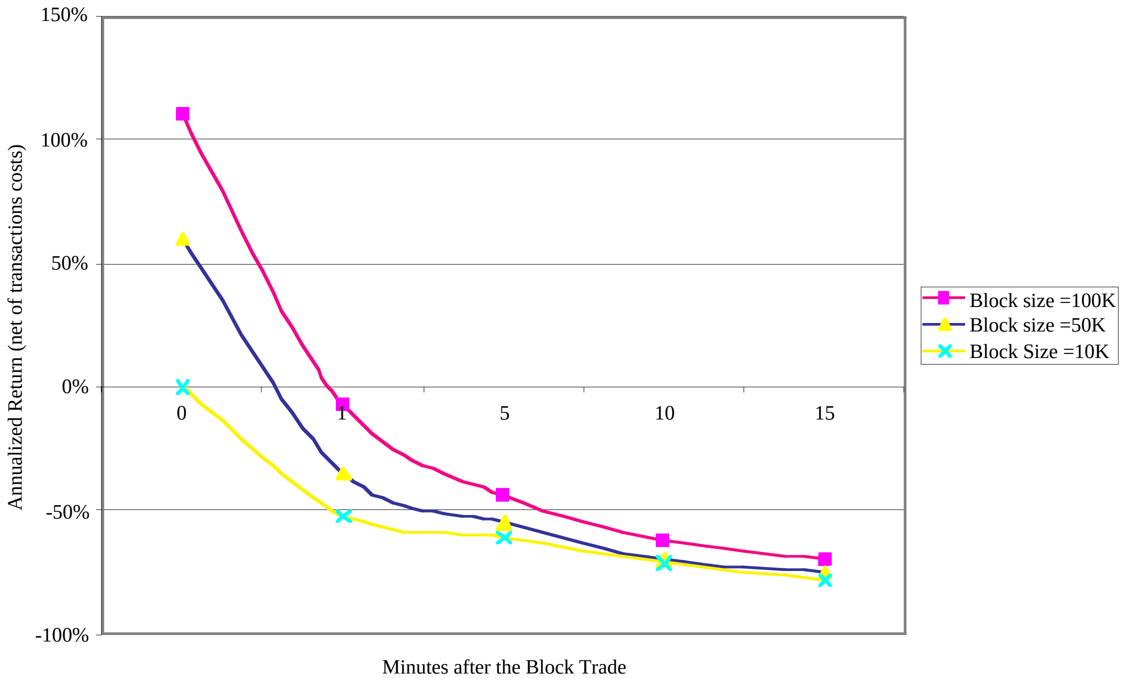

This study suffers from a sampling bias - it looks at large block trades in liquid stocks on the exchange floor. Studies that look at smaller, less liquid stocks find that the price impact tends to be larger and the adjustment back to the correct price is slower than it is for the more liquid stocks.[13] There are other interesting facts about block trades that have emerged from other studies. First, while stock prices go up on block buys and go down on block sell, they are far more likely to bounce back after sell trades. In other words, when

[12] See Dann, Mayers and Rabb.(1978)

[13] Joel Haasbrouck looked at a detailed data set that contained information on quotes, trades and spreads of stocks listed on the NYSE and came to this conclusion.

prices go up after a block buy, they are more likely to stay up.[14] A recent study[15] looks at both liquid and illiquid stocks on the NYSE also finds a tendency on the part of markets to overshoot. When a block buy is made, the price seems to go up too much and it can take several days for it to revert back to a normal level for illiquid stocks.

These studies, while they establish a price impact, also suffer from another selection bias, insofar as they look only at actual executions. The true cost of market impact arises from those trades that would have been done in the absence of a market impact but were not because of the perception that it would be large. In one of few studies of how large this cost could be, Thomas Loeb collected bid and ask prices from specialists and market makers, at a point in time, for a variety of block sizes. Thus, the differences in the spreads as the block size increases can be viewed as an expected price impact from these trades. Table 5.2 summarizes his findings across stocks, classified by market capitalization:

Table 5.2: Round-Trip Transactions Costs as a Function of Market Capitalization and Block Size

	Dollar Value of Block ($ thousands)								
Sector	*5*	*25*	*250*	*500*	*1000*	*2500*	*5000*	*10000*	*20000*
Smallest	17.30%	27.30%	43.80%						
2	8.90%	12.00%	23.80%	33.40%					
3	5.00%	7.60%	18.80%	25.90%	30.00%				
4	4.30%	5.80%	9.60%	16.90%	25.40%	31.50%			
5	2.80%	3.90%	5.90%	8.10%	11.50%	15.70%	25.70%		
6	1.80%	2.10%	3.20%	4.40%	5.60%	7.90%	11.00%	16.20%	
7	1.90%	2.00%	3.10%	4.00%	5.60%	7.70%	10.40%	14.30%	20.00%
8	1.90%	1.90%	2.70%	3.30%	4.60%	6.20%	8.90%	13.60%	18.10%
Largest	1.10%	1.20%	1.30%	1.71%	2.10%	2.80%	4.10%	5.90%	8.00%

The sectors refer to market capitalization, and show the negative relationship between size and price impact. Note, however the effect of increasing block sizes on expected price impact, within each sector; larger trades elicit much larger price impact than do smaller trades.

---

[14] See Holthausen, R. W., R. W. Leftwich, and D. Mayers, 1990, Large-Block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects," Journal of Financial Economics, 26, 71-95. and Keim, D. B., and A. Madhavan, 1995, Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Trades," Journal of Financial Economics, 37, 371-398.

[15] See Spierdijk, Nijman, and van Soest (2002)

While the Loeb studies suggest that price impact can create very large costs, studies of actual equity transactions suggest that institutional investors have learned how to reduce, if not eliminate, these costs by modifying their trading behavior. A study by Leinweber, who looked[16] at 13,651 equity transactions, totaling about $ 2 billion, by a large corporate pension plan in 1991, found a very weak relationship between trade size and trading cost. Figure 5.4 presents his findings on the percent trading cost and the size of the trade as a percent of the three-day average trading volume:

Figure 5.4: Percent Trading Costs and Trade Size

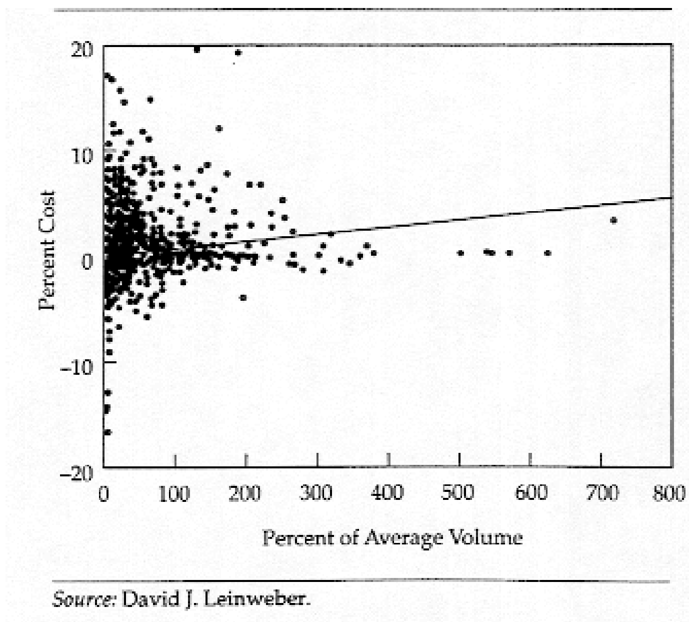

Note the bulge around the smallest trades, which seem to have both the lowest cost and the highest cost trades. Figure 5.5 presents the net trading loss by order size:

---

[16] See "Using Information from Trading in Trading and Portfolio Management" by D.J. Leinweber in Execution Techniques, True Trading Costs and Microstructure of Markets, AIMR.

Figure 5.5: Net Trading Loss By Order Size

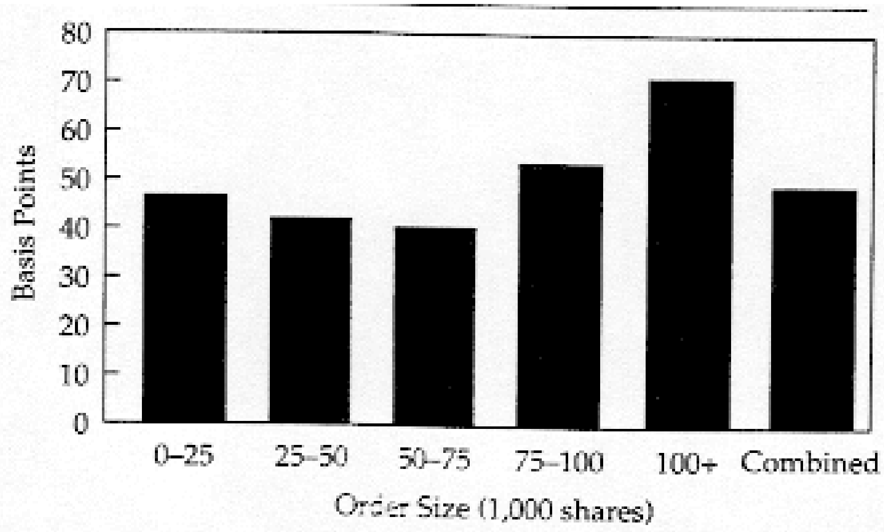

Source: David J. Leinweber

In fact, while smaller trades (<25,000 shares), on average, had lower trading costs than larger trades, they cumulatively accounted for almost 30% of the total trading costs for the fund. Thus, it may be just as important to worry about trading costs on small trades as on large trades, especially given the sheer number of small trades made by many portfolio managers and investors.

Determinants of the Price Impact

Looking at the evidence, the variables that determine that price impact of trading seem to be the same variables that drive the bid-ask spread. That should not be surprising. The price impact and the bid-ask spread are both a function of the liquidity of the market. The inventory costs and adverse selection problems are likely to be largest for stocks where small trades can move the market significantly.

Since you can reduce the price impact of trades by breaking them up into smaller trades, the price impact cost is likely to be greatest for investment strategies that require instantaneous trading. Thus, a portfolio manager who buys small, illiquid stocks because they are under valued is likely to face a smaller price impact cost than an investor who buys the same stocks after positive earnings announcements. The former can afford to spread his trades over time whereas the latter has to trade right after the announcement. We will consider this issue in more detail in the next section.

The Opportunity Cost of Waiting

The final component of trading costs is the opportunity cost of waiting. An investor can reduce the bid-ask spread and price impact costs of trading by trading patiently. If, in fact, there was no cost to waiting, even a large investor could break up trades into small lots and buy or sell large quantities without affecting the price or the spread significantly. There is, however, a cost to waiting. In particular, the price of an asset that an investor wants to buy because he or she believes that it is undervalued may rise while the investor waits to trade, and this, in turn, can lead to one of two consequences. One is that the investor does eventually buy, but at a much higher price, reducing expected profits from the investment. The other is that the price rises so much that the asset is no longer under valued and the investor does not trade at all. A similar calculus applies when an investor wants to sell an asset that he or she thinks is overvalued.

The cost of waiting will depend in great part on the probability that the investor assigns that the price will rise (fall) while he or she waits to buy (sell). We would argue that this probability will be a function of why the investor thinks the asset is under or over valued. In particular, the following factors should affect this probability:

1. *Is the valuation assessment based upon private information or is based upon public information?* Private information tends to have a short shelf life in financial markets, and the risks of sitting on private information are much greater than the risks of waiting when the valuation assessment is based upon public information. Thus, the cost of waiting is much larger when the strategy is to buy on the rumors (or information) of a possible takeover than it would be in a strategy of buying low PE ratio stocks.
2. *How active is the market for information?* Building on the first point, the risks of waiting, when one has valuable information, is much greater in markets where there are other investors actively searching for the same information. Again, in practical terms, the costs of waiting might be greater when there are dozens of analysts following the target stock than when there are few other investors paying attention to the stock.
3. *How long term or short term is the strategy?* While this generalization does not always hold, short-term strategies are much likely to be affected by the cost of waiting than longer term strategies. Some of this can be attributed to the fact that short term strategies are more likely to be motivated by private information, whereas long term strategies are more likely to be motivated by views on value.
4. *Is the investment strategy a "contrarian" or "momentum" strategy?* In a contrarian strategy, where investors are investing against the prevailing tide (buying when others are selling or selling when others are buying), the cost of waiting is likely to be smaller precisely because of this behavior. In contrast, the cost of waiting in a "momentum"

strategy are likely to be higher since the investor is buying when other investors are buying and selling when others are selling.

In summary, the cost of waiting is likely to be greatest for short term investment strategies, based upon private information or momentum, in markets with active information gathering. It will be less of an issue for long term investment strategies based upon public information and for contrarian strategies.

Investment Strategy and Total Trading Costs

The fact that assets which have high bid-ask spreads also tend to be assets where trading can have a significant price impact makes it even more critical that we examine investment strategies that focus disproportionately in these assets with skepticism. With the price impact, the effect of the size of the portfolio becomes much more critical, since large portfolios beget large trading blocks, which, in turn, have the biggest price impact. Thus, a strategy of investing in low-priced stocks which are not followed by analysts may yield excess returns, even after the bid-ask spread is considered, for a portfolio of $ 25 million but cease to be profitable if that same portfolio becomes $ 500 million.

Keim and Madhavan illustrate the interrelationship between total trading costs – implicit (including price impact and opportunity costs) as well as explicit (commissions and spreads) – and investment strategies.[17] Not surprisingly, they find that strategies that require large block trades have much higher total trading costs than strategies with smaller trades. They also find that the total trading costs are much greater for investors who buy small stocks as opposed to large ones. Table 5.3 provides a summary of their estimates of total trading costs for small cap and large cap companies listed on the NYSE and NASDAQ from 1991 to 1993.

Table 5.3: Total Round Trip Trading Costs and Market Capitalization

Market Capitalization	*Implicit Cost*	*Explicit Cost*	*Total Trading Costs (NYSE)*	*Total Trading Costs (NASDAQ)*
Smallest	2.71%	1.09%	3.80%	5.76%
2	1.62%	0.71%	2.33%	3.25%
3	1.13%	0.54%	1.67%	2.10%
4	0.69%	0.40%	1.09%	1.36%
Largest	0.28%	0.28%	0.31%	0.40%

[17] See "The Cost of Institutional Equity Trades" by Keim, D.B. and A. Madhavan in the Financial Analysts Journal, July/August 1998.

Note that the smallest companies have total round-trip trading costs that are significantly higher than the largest companies. They also find significant differences in costs between managers with different trading styles, with technical traders having the highest costs (presumably because of their need for immediate execution) and value traders the lowest costs.

Trading Costs with Non-traded Assets

If the cost of trading stocks can be substantial, it should be even more significant if your investment strategy requires you to hold assets that are not traded regularly such as collectibles, real estate or equity positions in private companies. In this section, we will consider these costs.

Trading Costs on Real Assets

If your investment strategy requires you to hold real assets, you may be exposed to very large trading costs. Real assets can range from gold to real estate to fine art and the transactions costs associated with trading these assets can also vary substantially. The smallest transactions costs are associated with precious commodities – gold, silver or diamonds – since they tend to come in standardized units. With residential real estate, the commission that you have to pay a real estate broker or salesperson can be 5-6% of the value of the asset. With commercial real estate, commissions may be smaller for larger transactions, but they will be well in excess of commissions on financial assets. With fine art or collectibles, the commissions become even higher. If you sell a Picasso through one of the auction houses, you may have to pay15-20% of the value of the painting as a commission. Why are the costs so high? The first reason is that there are far fewer intermediaries in real asset businesses than there are in the stock or bond markets; this reduces competition. The second is that real estate and fine art are not standardized products. In other words, one Picasso can be very different from another, and you often need the help of experts to judge value. This adds to the cost in the process.

Trading Costs on Private Equity/ Businesses

If your strategy requires you to take positions in private businesses – private equity as it is called – you have to allow for the fact that lucrative though the returns from these investments may be, they are illiquid. It is common, in fact, for investors in private businesses to assess an illiquidity discount on value to reflect their expectation that the cost of getting out of the position will be high. In this section, we will consider some of the factors that will determine this cost and empirical assessments of how big the cost may be.

Determinants of Illiquidity Cost

The cost of illiquidity is likely to vary across both firms and buyers, which renders rules of thumb useless. Let us consider first some of the factors that may cause the cost to vary across firms.

1. *Liquidity of assets owned by the firm*: The fact that a private firm is difficult to sell may be rendered moot if its assets are liquid and can be sold with no significant loss in value. A private firm with significant holdings of cash and marketable securities should have a lower illiquidity costs than one with factories or other assets for which there are relatively few buyers.
2. *Financial Health and cashflows of the firm*: A private firm that is financially healthy should be easier to sell than one that is not healthy. In particular, a firm with strong income and positive cash flows should be subject to a smaller illiquidity cost than one with negative income and cash flows.
3. *Possibility of going public in the future:* The greater the likelihood that a private firm can go public in the future, the lower should be the illiquidity cost. In effect, the probability of going public is built into the valuation of the private firm.
4. *Size of the Firm*: If we state the illiquidity cost as a percent of the value of the firm, it should become smaller as the size of the firm increases. In other words, the illiquidity discount should be smaller as a percent of firm value for firms like Cargill and Koch Industries, which are worth billions of dollars, than it should be for a small firm worth $15 million.

The illiquidity cost is also likely to vary across potential buyers because the desire for liquidity varies among individuals. It is likely that those buyers who have deep pockets and see little or no need to cash out their equity positions will face lower illiquidity costs, for similar firms, than buyers that have less of a safety margin.

Empirical Evidence on Illiquidity Cost

How large is the cost of being illiquid? This is a very difficult question to answer empirically because the discount attached to an asset's value itself cannot be observed. Even if we were able to obtain the terms of all private firm transactions, note that what is reported is the price at which private firms are bought and sold. The value of these firms is not reported and the illiquidity discount is the difference between the value and the price.

In fact, much of the evidence on illiquidity discounts comes from examining "restricted stock" at publicly traded firms. Restricted securities are securities issued by a publicly traded company, but not registered with the SEC, that can be sold through private placements to investors, but cannot be resold in the open market for a two-year holding period, and limited amounts can be sold after that. When this stock is issued, the issue price is set much lower than the prevailing market price, which is observable, and the difference is viewed as a discount for illiquidity. There have been several studies of restricted stock, and while they vary on the degree of the discount at which restricted stock are place, they all report significant discounts.[18]

In summary, then, there seems to be a substantial discount attached, at least on average, when an investment is not liquid. Much of the practice of estimating illiquidity discounts seems to build on these averages. For instance, rules of thumb often set the illiquidity discount at 20-30% of estimated value and there seems to be little or no variation across firms.

The Management of Trading Costs

The preceding discussion make clear out not only how large the "trading cost" problem is for active money managers, but also how difficult it is to develop a strategy to minimize the collective cost. Actions taken to reduce one type of trading cost (say, the brokerage commission or bid-ask spread) may increase another (for instance, the price impact). Strategies designed to minimize the collective impact of the bid-ask spread and the price impact (such as breaking up trades and using alternative trading routes) may increase the opportunity cost of waiting. In this section, we will examine ways in which trading costs can be managed within the broader construct of maximizing portfolio returns, given an investment philosophy.

Step 1: Develop a coherent investment philosophy and a consistent investment strategy

The first step in managing trading costs is developing and staying with a coherent investment philosophy and strategy. The portfolio managers who pride themselves on style switching and moving from one investment philosophy to another are the ones who bear the biggest burden in terms of transactions costs, partly because style switching increases turnover and partly because it is difficult to develop a trading strategy without a consistent investment strategy.

Step 2: Estimate the cost of waiting given the investment strategy

Mutual funds with highest trading costs: Take a look at the 50 mutual funds with the highest trading costs.

[18] Maher reports a discount of 35.43% on restricted stock issues made between 1969 and 1973; Moroney, using data from 1970, reports an average discount of 35%. Silber looked at restricted stock issues from 1949 to 1989 and estimates a discount of 33.75%.

The second step in the process is determining the cost of waiting for the investment strategy that is being followed. As noted in the previous section, the cost of waiting is likely to small for long-term, contrarian strategies and greater for short-term, information-based and momentum strategies. If the cost of waiting is very high, then the objective has to be minimize this cost, which essentially translates into trading as quickly as one can, even if the other costs of trading increase as a consequence.

Step 3: Look at the alternatives available to minimize transactions costs, given the cost of waiting

Once the cost of waiting has been identified, the investor can consider the third step which is to minimize the effect of the bid-ask spread and the price impact on portfolio returns. While we have talked about trading primarily in terms of trading on the floor of the exchange, there are a number of options that an investor can use to reduce the trading costs. Rose and Cushing[19] make some of the following suggestions to reduce trading costs on a portfolio for an institutional investor.

1. Take advantage of the alternatives to trading on the floor of the exchange. Among these alternatives are using the upstairs block market (where large buyers and sellers trade with each other), the dealer market (where trades are made with a dealer) and crossing networks, where trades are executed over a network. The trade off is straightforward - the approaches that yield the most liquidity (the exchange floor and the dealer market) are also the ones that have the highest trading costs.
2. Trade portfolios rather than individual stocks, when multiple orders have to be placed. Portfolio trades generally result in lower trading costs and allow for better risk-management and hedging capabilities.
3. Use technology to reduce the paperwork associated with trading and to keep track of trades which have already been made. By allowing traders to have information on whether their trades have been executed, and on trades that have already been made, technology can help control costs.
4. Be prepared prior to trading on ways to control liquidity and splits between manual and electronic trading. This “pre-trade” analysis will allow traders to identify the least costly and most efficient way to make a trade.
5. After the trade has been executed, do a post-trade analysis, where the details of the trade are provided in addition to a market impact analysis, which lists among other

---

[19] See “Making the Best Use of Trading Alternatives” by J.D. Rose and D.C. Cushing in Execution Techniques, True Trading Costs and the Microstructure of Markets, AIMR>

information, the benchmarks that can be used to estimate the price impact, including the mid-point of the bid-ask spread before the trade and the previous day's close. These post-trade analyses can then be aggregated across types of trades., securities and markets to give portfolio managers a measure of where their costs are greatest and how to control them.

Step 4: Stay within a portfolio size that is consistent with the investment philosophy and trading strategy that has been chosen

While it is tempting to most portfolio managers to view portfolio growth as the fruit of past success, there is a danger that arises from allowing portfolios to become too big. How big is too big? It depends upon both the portfolio strategy that has been chosen, and the trading costs associated with that strategy. While a long-term value investor who focuses well-known, large-capitalization stocks might be able to allow his or her portfolio to increase to almost any size, an investor in small-cap, high growth stocks or emerging market stocks may not have the same luxury, because of the trading costs we have enumerated in the earlier sections.

Step 5: Consider whether your investment strategy is yielding returns that exceed the costs

The ultimate test of an investment strategy lies in whether it earns excess returns after transactions costs. Once an investor has gone through the first four steps, the moment of truth always arrives when the performance of the portfolio is evaluated. If a strategy consistently delivers returns that are lower than the costs associated with implementing the strategy, the investor has one of two choices - he or she can switch to a passive investing approach (such as an index fund) or to a different active investing strategy, with higher expected returns or lower trading costs or both.

Taxes

As has often been said, the only two things that are certain in life are taxes and death. While investors may get a chance to pause and admire the pre-tax returns they make on their investment portfolios, they can spend only the returns that they have left after taxes. Strategies that yield attractive pre-tax returns can generate sub-standard after tax returns. There are two reasons why taxes are ignored by both researchers looking at investment strategies and portfolio managers who put these strategies into practice. The first is that taxes affect different investors differently, ranging from no impact on tax-exempt investors such as pension funds to very large effects on older and wealthier individual investors. The second is the complexity of the tax laws is such that the same investor may face different tax rates on different parts of his or her income (dividends versus capital gains) and different portions of his or her portfolio (pension fund versus savings).

Investment Returns and Taxes

How big of a drag are taxes on investment returns? Studies that look at returns on the U.S. stock market and government bonds show that stocks have generated much higher returns and ending portfolio managers for investors than treasury bills or bonds. Figure 5.6 presents the ending value of $ 100 invested in stocks, treasury bonds and treasury bills in 1928 and held through the end of 2001.

Figure 5.6: Portfolio Value from 1928 to 2001 - Stocks, T.Bonds and T.Bills

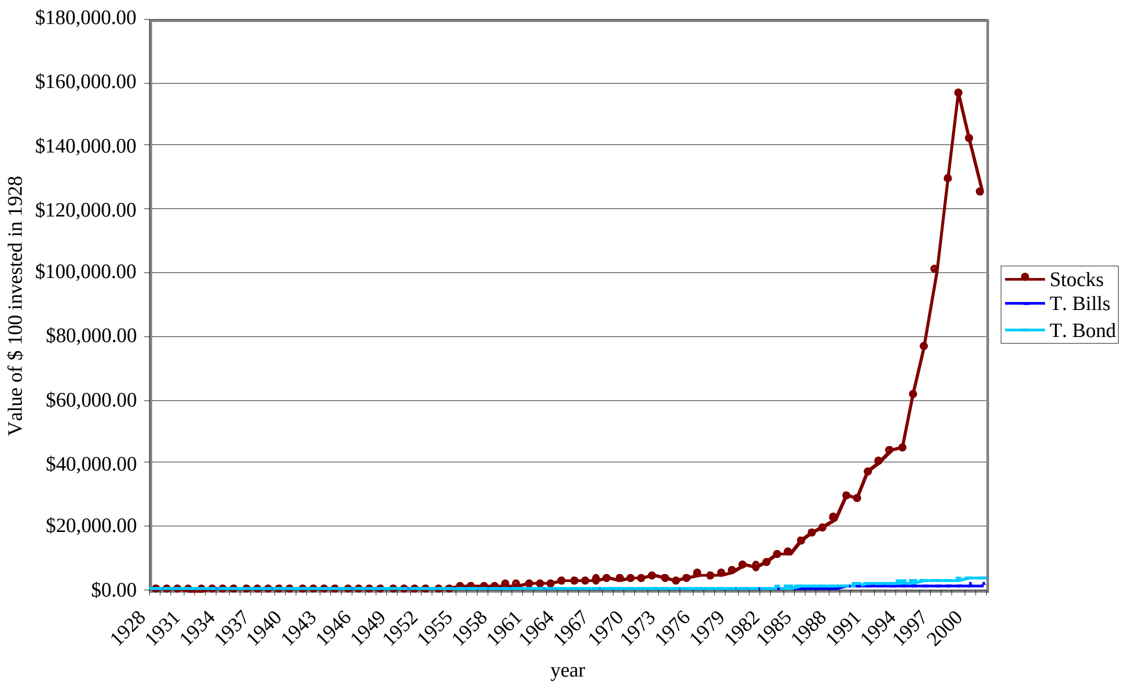

Thus, $ 100 invested in stocks would have grown to $ 125,599, significant higher than what your portfolio would have been worth if invested in T.Bills ($1,713) or T.Bonds ($3,587). This is impressive but it is also before taxes and transactions costs. Let us for the moment consider the effects of taxes on these returns. Assume that the investor buying these stocks faced a tax rate of 35% on dividends and 20% on capital gains over this period. To compute the effect of taxes on returns, we do have to consider how often this investor trades. If we assume that he turns over his entire portfolio at the end of each year, he would have to pay taxes on both dividends and the price appreciation each year. Figure 5.7 shows the effect on the portfolio value over the period and the effect of taxes on the ending portfolio:

Figure 5.7: Value of $ 100 invested in Stocks: Before and After Taxes

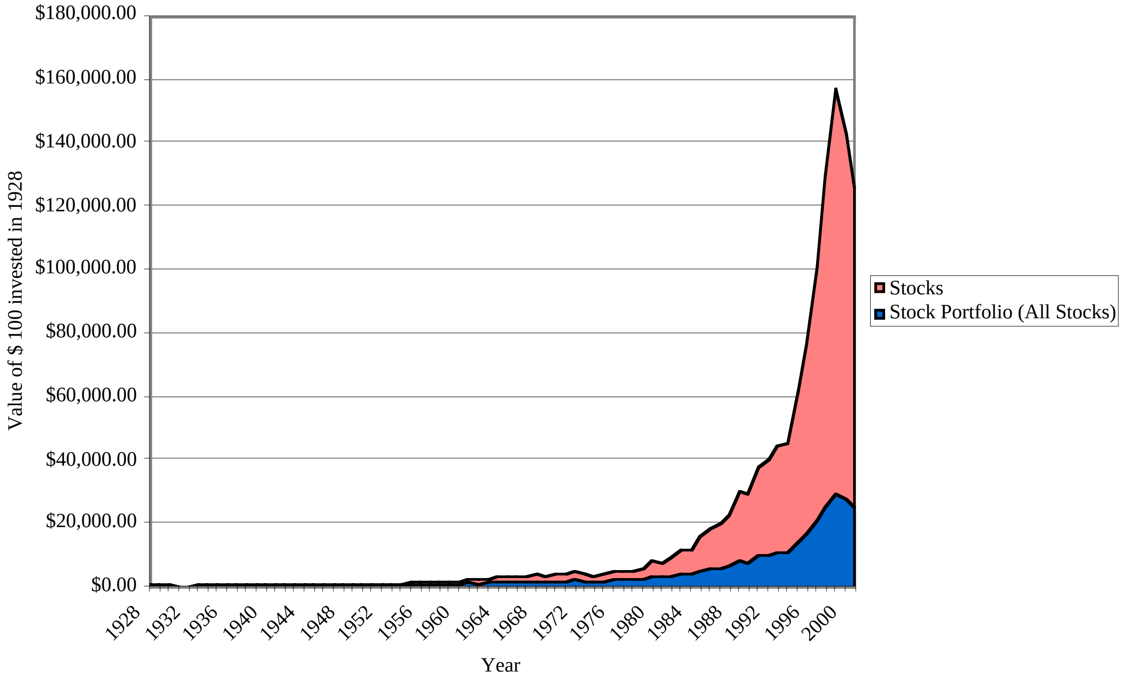

Note that introducing taxes into returns reducing the ending value of the portfolio by more than two thirds from $125,598 to $39,623.

But what if this investor, instead of turning over his entire portfolio once every year, had turned it over once every 2 years (or 3 or 5). Trading less often does not reduce the tax bite from dividends but it does allow investors to delay paying capital gains taxes, thus increasing the ending portfolio value. This insight about the relationship between taxes and trading frequency is a key one. Since much of the return when investing in stocks comes from price appreciation, the more frequently you trade, the higher your tax bill is likely to be for any given pre-tax return. In fact, the effect is likely to be exacerbated by the higher tax rates on short-term capital gains (which have been generally similar to ordinary tax rates) than long-term capital gains.

There is one final point to be made about the tax effect. While the taxes on capital gains can be deferred by not trading on your winners, the taxes on dividends have to be paid each period that you receive dividends. Thus, a strategy of investing in stocks that have higher dividend yields than average will result in less flexibility when it comes to tax timing and more taxes, at least relative to investing in low dividend yield stocks for the long term. We illustrate this in figure 5.8 for an investor by contrasting the performance of a portfolio with a dividend yield half that of the market each year to one with twice the dividend yield, keeping the total returns constant.[20]

Figure 5.8: Value of $ 100 invested in stocks in 1928 & Dividend Yields

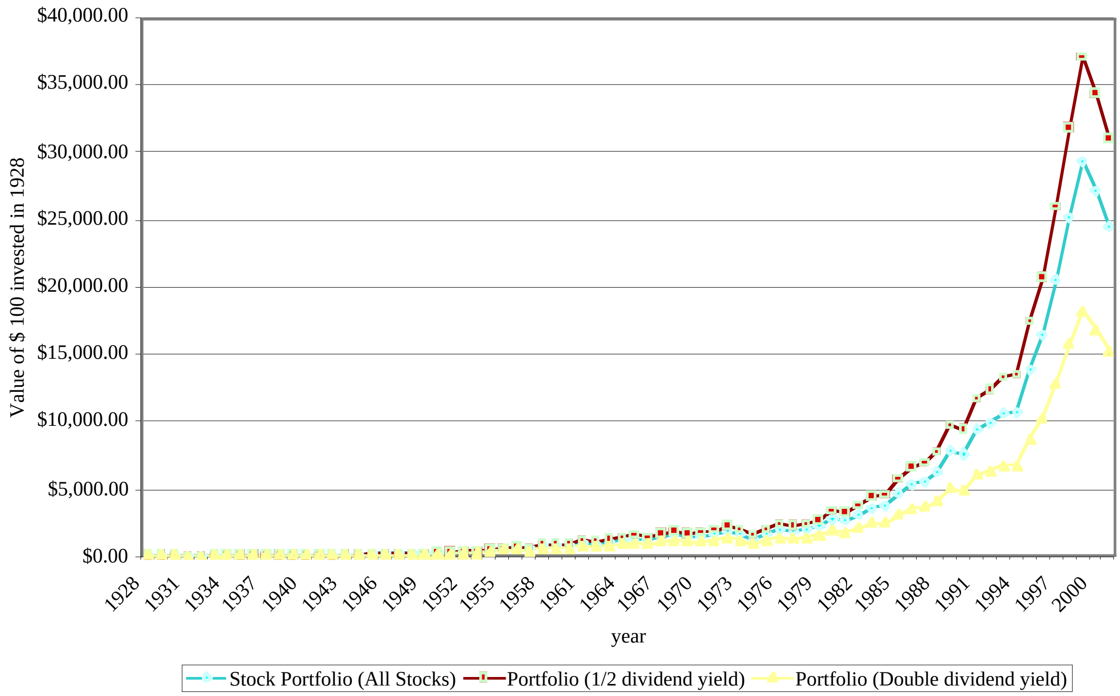

Note that the portfolio of stocks with half the dividend yield of the market has an ending value of just over $ 30,000 in 2001, whereas one with a dividend yield twice that of the market has an ending value of roughly half that amount.

The Tax Drag on Returns

How well do investors manage their tax liabilities? All too often, investment performance has been measured in terms of pre-tax returns. The rankings of mutual funds done by services such as Morningstar and Forbes have been based upon pre-tax returns. Until recently, the promotional material for most funds presented the pre-tax returns of these funds, contrasted with the S&P 500. This focus on pre-tax returns may be explained by the fact that investors have very different tax profiles and that it is difficult to find a typical investor, but it has also had the undesirable side effect. Money managers often adopt strategies that expose their investors to substantial tax bills because they feel that they will

---

[20] To provide an example, the average dividend yield across all stocks in 1996 was 3.20% and the total return was 23.82%. The half dividend yield portfolio was estimated to have a dividend yield of 1.60% and a price appreciation of 22.22% for a total return of 23.82%. The double dividend yield portfolio had a dividend yield of 6.40% and a price appreciation of 17.42% for a total return of 23.82%.

not be penalized for this tax exposure. Figure 5.9 presents the pre-tax and after-tax returns between 1997 and 2001 at the 14 largest domestic mutual funds in the United States in March 2002.

Figure 5.9: Pre-tax and After-tax Returns: Largest U.S. Mutual Funds

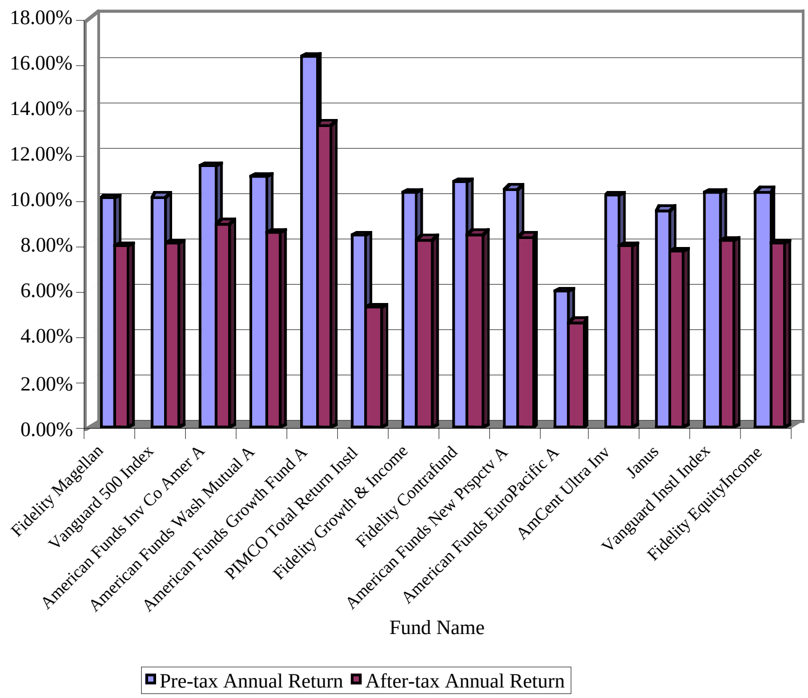

The after-tax returns are significantly lower than the pre-tax returns for each of the funds.

There are encouraging signs for investors concerned about taxes. The first is that the SEC has started requiring mutual funds to report their after-tax returns in conjunction with pre-tax returns in their promotional material. The second is that the mutual fund families have begun offering tax-efficient funds, where the objective is to maximize after-tax rather than pre-tax returns. The third is that the performance evaluators, such as Morningstar, have woken up to the tax costs being imposed on investors by mutual funds. In fact, the latest Morningstar reports on mutual funds report not only the after-tax returns over the last few years on these funds but also a measure of tax efficiency for each fund obtained by dividing the after-tax return by the pre-tax return. A fund that generates a pre-tax return on 9% and an after-tax return on 6% will therefore have a tax efficiency ratio of 67% (6/9). Figure 5.10 reports pre-tax and after-tax returns between 1999 and 2001 for equity mutual funds in the United States, categorized by style:

Figure 5.10: Pre-tax and After-tax Returns at U.S. equity mutual funds- 1999-2001

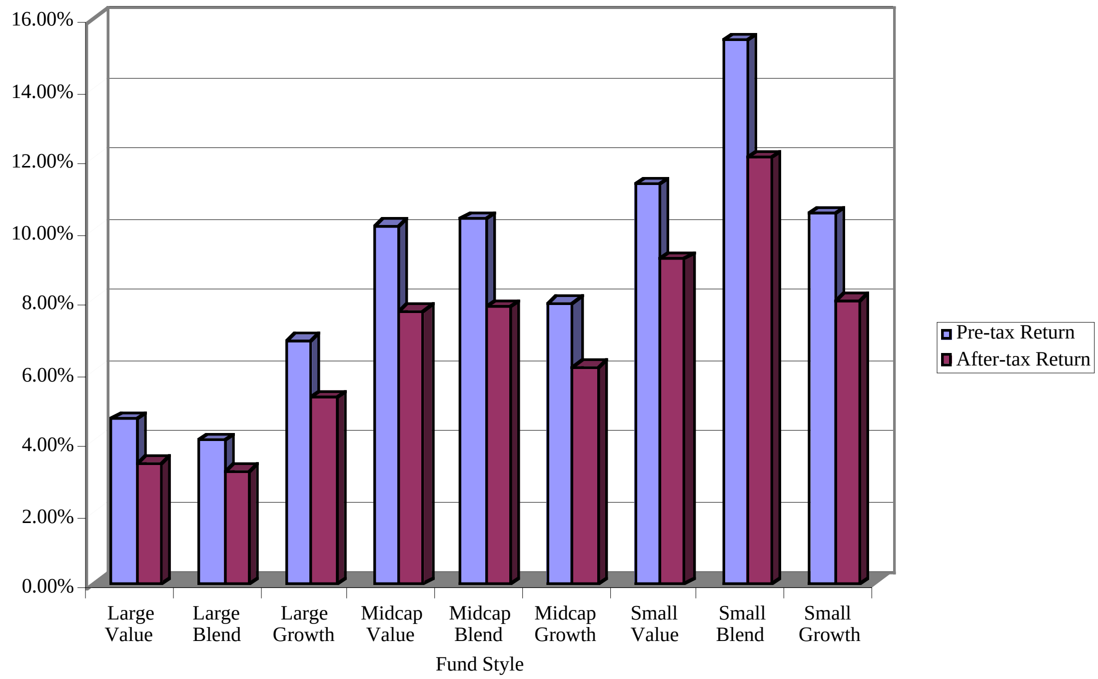

As you can see from the graph, there are significant differences between pre-tax and after-tax returns at many funds and the tax drag on returns cuts across all fund styles. While there are some funds that are tax-efficient, there are others where the after-tax return is less than 60% of the pre-tax return. What are the factors that cause it to be different? It is as function of a number of variables:

 *Mutual funds with highest tax drag*: Take a look at the 50 mutual funds with the largest difference between pre-tax and after-tax returns

- Higher turnover ratios (and more active trading) seem to give rise to higher tax costs for investors. After all, capital gains taxes are assessed only when you sell stocks. In fact, figure 5.11 categorizes mutual funds (both bond and equity) into six classes based upon turnover ratios and reports on the difference between pre-tax and after-tax returns in each class:

Figure 5.11: Tax Effect and Turnover Ratio: U.S. Mutual funds- 1999-2001

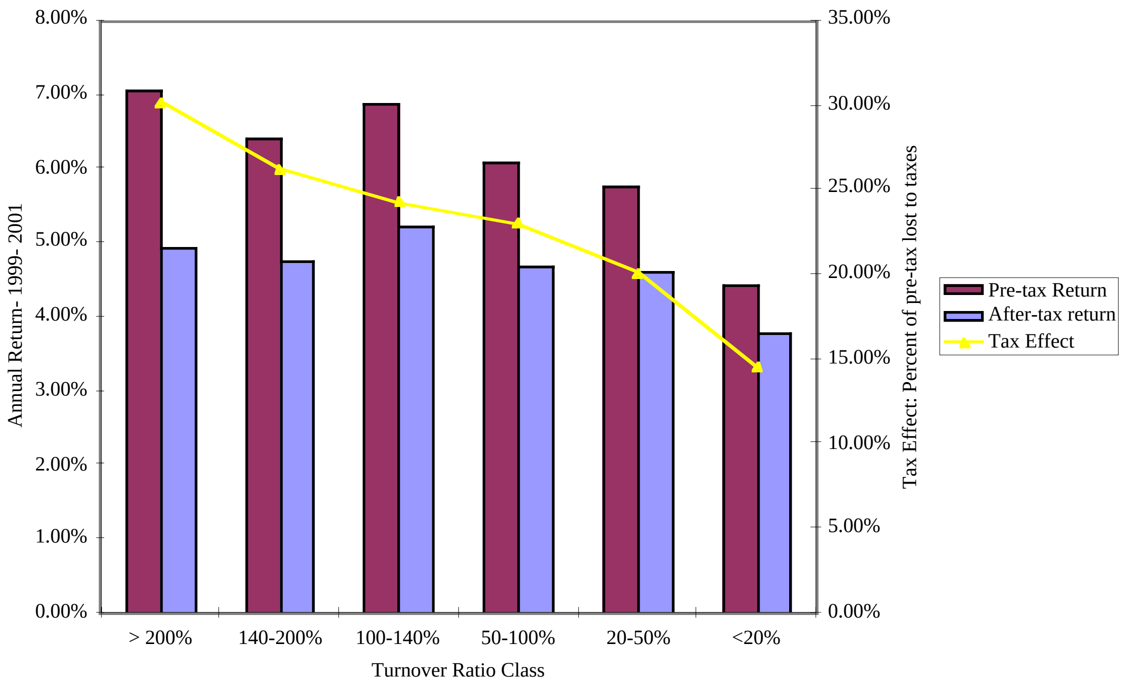

We measure the tax effect in each category by looking at the ratio of after-tax to pre-tax returns. Thus, the firms with turnover ratios that exceed 200% are the most tax inefficient, since the after-tax returns are more than 30% lower than the pre-tax returns. In fact, the pattern is consistent with the tax effect becoming smaller as the turnover ratios decrease.

- The after-tax return is also affected by money flowing in (inflows) and out of the fund (redemptions). Why might that be? Redemptions may require a fund to sell holdings to raise cash and, in the process, convert paper gains into taxable capital gains.
- The tax cost will also vary depending upon whether the mutual fund is actively trying to manage the cost. For instance, a fund that wants to minimize tax costs will sell a few losers when it sells winners and offset the capital losses on the former against the capital gains on the latter.

Tax Management Strategies

While tax considerations alone should not determine any investor's portfolio, they clearly have to be a factor in how portfolios are put together and the type of investment strategies adopted. In this section, we consider ways in which investors may be able to reduce how much they lose because of taxes.

Minimal Turnover

The simplest and most effective way to reduce your taxes is to trade less often. As noted in the prior sections, portfolio turnover is a key determinant of tax costs. While you may not be able to reduce your turnover ratios to those of an index fund (which often has a turnover ratio of 5% or less) if you want to actively beat the market, you can still try to minimize trading, given any strategy that you adopt. In addition, your choice of strategy should be influenced by the amount of trading that goes with it. Does this mean that you should avoid strategies that require a lot of trading? Not necessarily. If you can earn a high enough return to cover the additional taxes you have to pay, you may still choose to go with this strategy.

Tax-Based Trading

The other way to minimize taxes is to consider trades specifically for the purpose of reducing your tax bill. There are several forms of tax-based trading:

- In its simplest form, you may sell stocks in your portfolio that have gone down just before a tax year ends and use the capital losses from the sale to offset capital gains on other stocks that you may have sold during the year.
- If your tax status varies over time, you may choose to claim your capital gains in a year in which your tax rate is low and your capital losses in a year in which your tax rate is high.
- In its most dangerous form, you may make investments specifically because they offer a chance to reduce your tax bill. While such tax shelters have long been utilized by investors, you should recognize that the tax authorities usually require these entities to have an economic purpose that goes beyond tax reduction. All too often, investors are beguiled by promised tax savings in tax shelters that are never delivered.

In theory, while tax-based trading offers promise, investors have to keep it in perspective. In particular, trading purely for the purposes of reducing your tax bills strikes us as more likely to reduce than increase returns. However, augmenting a sensible trading strategy with tax considerations makes sense. Thus, if you were planning to prune and rebalance your portfolio, you should probably do it before the end of the tax year and consider the effect of these trades on your tax bill. If you are trying to decide between selling two stocks in your portfolio, the fact that one of them could reduce taxes paid by $ 100,000 should tip the balance.

Tax Arbitrage

The fact that all investors do not face the same tax rates and that each investor may, in fact, be taxed differently on different parts of her portfolio does raise an interesting implication. If your tax rate as an investor is much lower than the tax rates of other investors in the market, you may be able to exploit the difference to earn excess returns. To see why, consider the following scenario. Assume that you are tax exempt and that every other investor in the market faces a 40% tax rate on both dividends and capital gains. Let us also assume that these other investors price stocks to earn an after-tax return of 9%. On a pre-tax basis, stocks will have to earn 15%. As a tax-exempt investor, you will be able to earn the 15% pre-tax return as well, but your after-tax returns will also be 15%.

Is tax arbitrage feasible if you are the only taxable investor in a tax-exempt universe? You will either have to settle for lower after-tax return than other investors in the market or not buy stocks at all. In reality, there are both tax-exempt and taxable investors in every market, and tax rates vary widely across taxable investors. The market prices of assets will reflect the relative magnitudes of each group, and there will always be groups of assets that yield more favorable returns for each group. Thus, tax exempt investors may find their best bargains in stocks that generate the greatest tax liabilities for other investors – high dividend paying stocks, for instance. High tax rate investors migrate towards stocks where they are penalized the least by their tax status – non-dividend paying, high growth stocks would be an example.

Conclusion

Trading costs are an integral part of any investment portfolio and can make the difference between a portfolio that beats the market and one that does not. The overall evidence suggests that trading costs impose a significant drag on portfolio returns, and may explain why active money managers under perform the market. The reason trading costs are large is that they include not just brokerage costs, but also the costs associated with the bid-ask spread, the price impact created by trading and the cost of waiting. The reason they are difficult to control is that actions taken to reduce one component of the trading cost tend to increase the other components. Trading costs do not impose a uniform burden on all investment strategies. They punish short term, information-based strategies far more than they do long term value-based strategies; they affect strategies that focus on less-liquid assets far more than they do strategies that are built around liquid assets. No matter what the strategy, though, it is the portfolio manager’s job to manage trading costs, given the constraints of the strategy, and earn an excess return that covers these costs.

Taxes can also have significant implications for investment strategy. Since you get to keep only after-tax profits, the tax exposure generated by different investment strategies has to be an important factor in which strategy you choose. With any given investment strategy, you can try to reduce your tax bill by trading less often, tax-based trading or the use of tax shelters.

Lessons for Investors

1. The brokerage costs, which are often the most explicit costs of trading, represent only a small portion of the total trading cost. There are at least three other costs associated with trading, the bid-ask spread that you bear at the time of trading, the price impact that you have as a result of trading and the opportunity cost of waiting to trade.
2. While large investors may have an advantage over small investors when it comes to brokerage costs and even the bid-ask spread – they often face a narrower spread – they face much large price impact and opportunity costs than small investors.
3. The drag imposed by trading costs on returns will depend upon what type of stocks you invest in – it will be higher for smaller, less liquid stocks – and how much you trade – higher turnover will create higher trading costs.
4. As an investor, you get to spend after-tax income, not pre-tax income. The portion of your returns that will be devoured by taxes will depend, like trading costs, on what you hold – dividend paying stocks will create larger tax bills – and how much you trade – more trading will generate more taxes

CHAPTER 6

TOO GOOD TO BE TRUE? TESTING INVESTMENT STRATEGIES

As investors, we are constantly bombarded with sales pitches from experts claiming to have found the secret formula or the magic model that guarantees investment success. Buy stocks using this strategy, they say, and you will get a portfolio that has low risk and high returns. While you do not want to rule out the possibility that such strategies exist, it pays to be skeptical. In this chapter, we consider how to test investment strategies. In the process, we will also examine what we mean when we say that markets are efficient and cannot be beaten or that markets are inefficient and can be beaten

What is an efficient market? What does it imply for investment and valuation models? Clearly, market efficiency is a concept that is controversial and attracts strong views, pro and con, partly because of differences between individuals about what it really means, and partly because it is a core belief that in large part determines how an investor approaches investing. This chapter provides a simple definition of market efficiency, considers the implications of an efficient market for investors and summarizes some of the basic approaches that are used to test investment schemes, thereby proving or disproving market efficiency.

Market Efficiency and Investment Philosophies

The question of whether markets are efficient, and if not, where the inefficiencies lie, is central to choosing your investment philosophy. If markets are, in fact, efficient, the market price provides the best estimate of value, and the process of valuation becomes one of justifying the market price. As an investor, you would not then try to pick under or over valued stocks or time the market. Instead, you would diversify across a broad band of stocks and not trade very often.

If markets are not efficient, the market price may in fact be wrong, and which investment philosophy you pick will depend upon why you believe markets make mistakes and how they correct them. Those investors who can pinpoint these misvalued stocks, then, will then be able to make 'higher' returns than other investors, thus accomplishing the very difficult task of beating the market. But what if markets are really efficient but you mistakenly pick stocks, thinking they are inefficient. You will bear both the cost of the resources you spend (in terms of time and money) in picking stocks and the additional taxes and transactions costs of this strategy. Consequently, you will end up with a far lower return than that earned by your neighbor who invested her wealth in an index fund.

Examining where and when there are market inefficiencies can also help us in the far more prosaic task of picking investment strategies. A value investor, for example, may have to decide between low price to book value stocks and low price earnings ratio companies. The evidence may yield a clue as to which strategy is more effective at highlighting undervalued stocks, In addition, market 'inefficiencies' can provide the basis for screening the universe of stocks to come up with a sub-sample that is more likely to have under valued stocks. Given the number of stocks that you get to pick from, this not only saves time for you as an investor but increases the odds significantly of finding under and over valued stocks. For instance, some efficiency studies suggest that stocks that are 'neglected' by institutional investors are more likely to be undervalued and earn excess returns. A strategy that screens firms for low institutional investment (as a percentage of the outstanding stock) may yield a sub-sample of neglected firms, which can then be analyzed to arrive at a portfolio of undervalued firms. If the research is correct, the odds of finding undervalued firms should increase in this sub-sample.

Market Efficiency: Definition and Implications

Market efficiency means different things to different people. To those who are true believers in efficient markets, it is an article of faith that defines how they look at or explain market phenomena. To critics, it indicates an academic notion of infallible and supremely rational investors who always know what the true value of an asset is. In this section, we would like to first define what we believe an efficient market is and then follow it up with implications for investment strategies.

What is an efficient market?

In its most general sense, an efficient market is one where the market price is an unbiased estimate of the true value of the investment. Implicit in this derivation are several key concepts -

(a) Contrary to popular view, market efficiency does not require that the market price be equal to true value at every point in time. All it requires is that errors in the market price be unbiased, i.e., that prices can be greater than or less than true value for individual stocks, as long as these deviations are random[1].

(b) The fact that the deviations from true value are random implies, in a rough sense, that there is an equal chance that any given stock is under or over valued at any point in time, and that the deviation is uncorrelated with any observable variable. For instance, in an efficient

[1] Randomness implies that there is an equal chance that stocks are under or over valued at any point in time.

market, stocks with lower PE ratios should be no more or less likely to under valued than stocks with high PE ratios.

(c) If the deviations of market price from true value are random, it follows that no group of investors should be able to consistently find under or over valued stocks using any investment strategy.

Market Efficiency, Investors and Information

Definitions of market efficiency have to be specific not only about the market that is being considered but also the investor group that is covered. It is extremely unlikely that all markets are efficient to all investors at all times, but it is entirely possible that a particular market (for instance, the New York Stock Exchange) is efficient with respect to the average investor. It is also possible that some markets are efficient while others are not, and that a market is efficient with respect to some investors and not to others. This is a direct consequence of differential tax rates and transactions costs, which confer advantages on some investors relative to others.

Definitions of market efficiency are also linked up with assumptions about what information is available to investors and reflected in the price. For instance, a strict definition of market efficiency that assumes that all information, public as well as private, is reflected in market prices would imply that even investors with precise inside information will be unable to beat the market. One of the earliest classifications of levels of market efficiency was provided by Fama (1971), who argued that markets could be efficient at three levels, based upon what information was reflected in prices. Under weak form efficiency, the current price reflects the information contained in all past prices, suggesting that charts and technical analyses that use past prices alone would not be useful in finding under valued stocks. Under semi-strong form efficiency, the current price reflects the information contained not only in past prices but all public information (including financial statements and news reports) and no approach that was predicated on using and massaging this information would be useful in finding under valued stocks. Under strong form efficiency, the current price reflects all information, public as well as private, and no investors will be able to consistently find under valued stocks.

Implications of market efficiency

An immediate and direct implication of an efficient market is that no group of investors should be able to consistently beat the market using a common investment strategy. An efficient market would also carry very negative implications for many investment strategies -

(a) In an efficient market, equity research and valuation would be a costly task that provided no benefits. The odds of finding an undervalued stock would always be 50:50, reflecting the randomness of pricing errors. At best, the benefits from information collection and equity research would cover the costs of doing the research.

(b) In an efficient market, a strategy of randomly diversifying across stocks or indexing to the market, carrying little or no information cost and minimal execution costs, would be superior to any other strategy that created larger information and execution costs. There would be no value added by portfolio managers and investment strategists.

(c) In an efficient market, a strategy of minimizing trading, i.e., creating a portfolio and not trading unless cash is needed would be superior to a strategy that required frequent trading.

It is therefore no wonder that the concept of market efficiency evokes such strong reactions on the part of portfolio managers and analysts, who view it, quite rightly, as a challenge to their existence.

It is also important that there be clarity about what market efficiency does not imply. An efficient market does not imply that -

(a) Stock prices cannot deviate from true value; in fact, there can be large deviations from true value. The only requirement is that the deviations be random.

(b) No investor will 'beat' the market in any time period. To the contrary, approximately half[2] of all investors, prior to transactions costs, should beat the market in any period.

(c) No group of investors will beat the market in the long term. Given the number of investors in financial markets, the laws of probability would suggest that a fairly large number are going to beat the market consistently over long periods, not because of their investment strategies but because they are lucky. It would not, however, be consistent if a disproportionately large number[3] of these investors used the same investment strategy.

In an efficient market, the expected returns from any investment will be consistent with the risk of that investment over the long term, though there may be deviations from these expected returns in the short term.

---

[2] Since returns are positively skewed, i.e., large positive returns are more likely than large negative returns (since this is bounded at -100%), less than half of all investors will probably beat the market.

[3] One of the enduring pieces of evidence against market efficiency lies in the performance records posted by many of the investors who learnt their lessons from Ben Graham in the fifties. No probability statistics could ever explain the consistency and superiority of their records.

Necessary conditions for market efficiency

Markets do not become efficient automatically. It is the actions of investors, sensing bargains and putting into effect schemes to beat the market, that make markets efficient. The necessary conditions for a market inefficiency to be eliminated are as follows -

(1) The market inefficiency should provide the basis for a trading scheme to beat the market and earn excess returns. For this to hold true -

(a) The asset (or assets) which is the source of the inefficiency has to be traded.

(b) The transactions costs of executing the scheme have to be smaller than the expected profits from the scheme.

(2) There should be profit maximizing investors who

(a) recognize the 'potential for excess return'

(b) can replicate the beat the market scheme that earns the excess return

(c) have the resources to trade on the asset until the inefficiency disappears

The internal contradiction of claiming that there is no possibility of beating the market in an efficient market and requiring profit-maximizing investors to constantly seek out ways of beating the market and thus making it efficient has been explored by many. If markets were, in fact, efficient, investors would stop looking for inefficiencies, which would lead to markets becoming inefficient again. It makes sense to think about an efficient market as a self-correcting mechanism, where inefficiencies appear at regular intervals but disappear almost instantaneously as investors find them and trade on them.

Propositions about market efficiency

A reading of the conditions under which markets become efficient leads to general propositions about where investors are most likely to find inefficiencies in financial markets-

Proposition 1: *The probability of finding inefficiencies in an asset market decreases as the ease of trading on the asset increases. To the extent that investors have difficulty trading on an asset, either because open markets do not exist or there are significant barriers to trading, inefficiencies in pricing can continue for long periods.*

This proposition can be used to shed light on the differences between different asset markets. For instance, it is far easier to trade on stocks that it is on real estate, since markets are much more open, prices are in smaller units (reducing the barriers to entry for new traders) and the asset itself does not vary from transaction to transaction (one share of IBM is identical to another share, whereas one piece of real estate can be very different from another piece, a stone's throw away. Based upon these differences, there should be a greater likelihood of finding inefficiencies (both under and over valuation) in the real estate market.

Proposition 2: *The probability of finding mispricing in an asset market increases as the transactions and information cost of exploiting the inefficiency increases. The cost of collecting information and trading varies widely across markets and even across investments in the same markets. As these costs increase, it pays less and less to try to exploit these inefficiencies.*

Consider, for instance, the perceived wisdom that smaller companies that are not followed by analysts or held by institutions are more likely to be under valued. This may be true in terms of raw returns, but transactions costs are likely to be much higher for these stocks since-

(a) They are often unlisted or listed on the NASDA!, leading to higher brokerage commissions and expenses

(b) Due to their illiquidity and the fact that they are low priced stocks, the bid-ask spread becomes a much higher fraction of the total price paid.

(c) Trading is often thin on these stocks, and small trades can cause prices to change resulting in a higher 'buy' price and a lower 'sell' price.

Once you consider the transactions costs, you may very well find that the excess returns you perceived in these stocks may be gone.

Corollary 1: *Investors who can establish a cost advantage (either in information collection or transactions costs) will be more able to exploit small inefficiencies than other investors who do not possess this advantage.*

There are a number of studies that look at the effect of block trades on prices, and conclude that while they affect prices, that investors will not be exploit these inefficiencies because of the number of times they will have to trade and their transactions costs. These concerns are unlikely to hold for a specialist on the floor of the exchange, who can trade quickly, often and at no or very low costs. It should be pointed out, however, that if the market for specialists is efficient, the value of a seat on the exchange should reflect the present value of potential benefits from being a specialist.

This corollary also suggests that investors who work at establishing a cost advantage, especially in relation to information, may be able to generate excess returns on the basis of these advantages. Thus a John Templeton, who started investing in Japanese and other Asian markets well before other portfolio managers, might have been able to exploit the informational advantages he had over his peers to make excess returns on his portfolio.

Proposition 3: *The speed with which an inefficiency is resolved will be directly related to how easily the scheme to exploit the inefficiency can be replicated by other investors. The ease with which a scheme can be replicated itself is inversely related to the time, resources* *and information needed to execute it. Since very few investors single-handedly possess the resources to eliminate an inefficiency through trading, it is much more likely that an inefficiency will disappear quickly if the scheme used to exploit the inefficiency is transparent and can be copied by other investors.*

To illustrate this point, assume that stocks are consistently found to earn excess returns in the month following a stock split. Since firms announce stock splits publicly, and any investor can buy stocks right after these splits, it would be surprising if this inefficiency persisted over time. This can be contrasted with the excess returns made by some 'arbitrage funds' in index arbitrage, where index futures are bought (sold), and stocks in the index are sold short (bought). This strategy requires that investors be able to obtain information on index and spot prices instantaneously, have the capacity (in terms of margin requirements and resources) to buy and sell index futures and to sell short on stocks, and to have the resources to take and hold very large positions until the arbitrage unwinds. Consequently, inefficiencies in 'index futures pricing' are likely to persist at least for the most efficient arbitrageurs, with the lowest execution costs and the speediest execution times.

Inefficiency or Anomaly: A Simple Test

When you do find an investment strategy that seems to beat the market, it is always an open question as to whether you have found a market mistake that can be exploited for excess returns or just a phenomenon that occurs in financial markets that you are unable to explain because the models you use are incorrect or the data you use is incomplete or erroneous. You would categorize the first as an inefficiency and the second as an anomaly. The pragmatic difference is that you would try to make money off the first but not of the second.

One way to tell the difference is to observe what happens to the excess returns once a strategy has been uncovered and publicized. If it has uncovered an inefficiency, you should see the excess returns rapidly disappear after the strategy is made public. If it is an anomaly, you will see the excess returns continue unabated even after it is publicized.

Testing market efficiency

Tests of market efficiency look at the whether specific investment strategies or portfolio managers beat the market. But what does beating the market involve? Does it just imply that someone earns a return greater than what the market (say, the S&P 500) earns in a specific year? We will begin by looking at what beating the market involves and define what we mean be excess returns. We will then follow up by looking at three standard ways of testing market efficiency and when and why we may choose one over the other.

Beating the Market

The fundamental question that we often attempt to answer when we test an investment strategy is whether the return we earn from the strategy is above or below a benchmark return on an alterative strategy of equivalent risk. But what should that benchmark return be? As we shall see, it is almost impossible to measure the success or failure of an investment strategy without taking a point of view on how risk should be measured.

Performance Benchmarks

If you can estimate the returns that you could have made by adopting an investment strategy in the past or observed the returns made by a portfolio manager or investor over a period, you can evaluate those returns. To make the evaluation, you have to choose an appropriate benchmark. In this section, we will consider the alternatives that are available to us in making this choice.

1. Comparison to Indices

When you have estimated the returns on a strategy, the simplest comparison you can make is to the returns you would have made by investing in an index. Many portfolio managers and investors still compare the returns they make on their portfolios to the returns on the S& P 500. While this comparison may be simple, it can also be dangerous when you have a strategy that does not have the same risk as investing in the index and the bias can cut both ways. If you have a strategy that is riskier than investing in the index – say investing in small, high growth stocks - you are biasing yourself towards concluding that the strategy works (i.e., it beats the market). If you have a strategy that is much safer than investing in the index, such as buying high dividend paying, mature companies, you are biasing yourself towards concluding that the strategy does not work.

There are slightly more sophisticated versions of this approach that are less susceptible to this problem. For instance, some services that judge mutual funds do so by comparing them to an index of funds that have the same style as the fund being judged. Thus, a fund that invests in large market cap companies with low price to book ratios will be compared to other large-cap, value funds. The peril remains, though, since categorizing investors into neat boxes is easier said than done. A fund manager may begin the year calling herself a large-cap, value investor and during the course of the year shift to being an investor in high growth, risky companies.

2. Risk and Return Models

In chapter 3, we considered the basics of risk and put forth several risk and return models. All of these models tried to measure the risk in an investment, though they differed on how best to measure it, and related the expected return on the investment to the risk measure. You could use these models to measure the risk in an investment strategy, and then examine the returns relative to this risk measure. We will consider some of these risk-adjusted measures of performance in this section.

a. Mean Variance Measures

The simplest measures of risk-adjusted performance have their roots in the mean-variance framework developed by Harry Markowitz in the early 1950s. In the mean-variance world, the standard deviation of an investment measures its risk and the return earned is the reward. If you compare two investments with the same standard deviation in returns, the investment with the higher average return would be considered to be the better one.

I. Sharpe Ratio

Extending this concept to investment strategies, you could look at the payoff to each unit of risk taken by dividing the return earned using the strategy by the standard deviation of return, in a measure called the Sharpe ratio.[4]

Sharpe Ratio = Average Return on Strategy/ Standard deviation of Returns from Strategy

To compute the standard deviation, you would need to track the returns each period for several periods. For instance, the average monthly returns over the last 5 years at mutual funds can be divided by the standard deviations in monthly returns at mutual funds over the 5 years to come up with Sharpe ratios for mutual funds. Once you have the Sharpe ratios for individual funds, you can compare them across funds to find the funds that earn the highest reward per unit of risk (standard deviation) or you can compare the fund's ratios to the Sharpe ratio for the entire market to make a judgment on whether active investing pays.

> *Funds with highest Sharpe Ratios*: Take a look at the 50 mutual funds with the highest Sharpe ratios.

The Sharpe ratio is a versatile measure that has endured the test of time. Its focus on the standard deviation as the measure of risk does bias it against portfolios that are not diversified widely across the market. A sector-specific mutual fund (such as a bio-tech or health care fund) will tend to do poorly on a Sharpe ratio basis because its standard deviation will be higher because of the presence of sector-specific

---

[4] This measure was developed by Bill Sharpe in the mid 1960s as a measure of mutual fund performance.

risk. Since investors in these funds can diversity away that risk by holding multiple funds, it does seem unfair to penalize these funds for them.

Illustration 6.1: Computing Sharpe Ratios for largest U.S. Mutual funds – 1997-2001

In table 6.1, we compute the Sharpe ratios for the 15 largest mutual funds in the United States in March 2002, using data from 1997 to 2001.

Table 6.1: Sharpe Ratios for Large U.S. Mutual Funds: 1997-2001

Fund Name	*Average Return*	*Standard deviation*	*Sharpe Ratio*
Fidelity Magellan	10.72%	20.82%	0.51
Vanguard 500 Index	10.14%	19.63%	0.52
American Funds Inv Co Amer A	12.84%	15.69%	0.82
American Funds Wash Mutual A	12.36%	16.64%	0.74
American Funds Growth Fund A	17.66%	25.00%	0.71
Fidelity Growth & Income	10.32%	16.10%	0.64
Fidelity Contrafund	11.47%	18.33%	0.63
American Funds New Prspctv A	11.82%	18.33%	0.64
American Funds EuroPacific A	7.20%	18.75%	0.38
AmCent Ultra Inv	10.17%	26.15%	0.39
Janus	9.57%	27.31%	0.35
Vanguard Instl Index	10.27%	19.65%	0.52
Fidelity EquityIncome	10.39%	17.89%	0.58
S & P 500	10.35%	19.55%	0.53

Note that the American funds delivered Sharpe ratios that were much higher than the S&P 500. In contrast, the Janus fund delivered a lower Sharpe ratio. The two biggest funds, Fidelity Magellan and the Vanguard 500 (not surprisingly), delivered Sharpe ratios that were very similar to the S&P 500.

II. Information Ratio

A close relative of the Sharpe ratio is the information ratio. It is the ratio of the excess return earned by a fund over an index to the excess volatility of this fund to the volatility of the index. To measure the latter, we estimate what is commonly called tracking error, which measures the deviations of the fund returns from the index returns each period over several periods. In its most common form, the excess return over the S&P 500 for a fund is divided by the tracking error of the fund relative to the S&P 500.

Information Ratio = (Return on Strategy – Return on Index)/ Tracking Error versus the Index

Information ratios differ from Sharpe ratios because of their fidelity to an index. In other words, you can have a portfolio with low standard deviation but it can have high tracking error, if it contains stocks that are not in the index. For instance, a portfolio of low risk stocks with low market capitalization may have low standard error but it will still have a high tracking error versus the S&P 500.

III. M Squared

In the 1990s, these measures were refined slightly to come up with a measure called M squared.[5] Instead of dividing the total return of a strategy or fund by its standard deviation, you compute the expected return you would have had on the fund, if you had to adjust its standard deviation down to that of the index and compare this expected return to the return on the index. For instance, assume that you have a fund with a return of 30% and a standard deviation of 50% and that the return on the S&P 500 is 15% and the standard deviation of the S&P 500 is 20%. To make the fund's standard deviation comparable to that of the S&P 500, you would have had to invest 60% of your money in T.Bills (earning 3%) and 40% in the fund:

Adjusted Standard Deviation of Portfolio = .4 (Std deviation of fund) + .6(0)

= .4 (50%) = 20%

The return on this portfolio can then be calculated:

Expected return on portfolio = .4 (30%) + .6 (3%) = 13.8%

Since this return is lower than the S&P's return of 15%, you would categorize this fund as an underperformer.

This measure of performance is closely related to the Sharpe ratio and is susceptible to the same biases. Since the expected return is adjusted to make the risk of the mutual fund similar to that of the index, funds that are not diversified widely across the market will score poorly.

b. Capital Asset Pricing Model

The capital asset pricing model emerged from the mean-variance framework to become the first model for risk and return in finance. In the CAPM, as described in chapter 3, the expected return on an investment can be written as function of its beta:

---

[5] The developed of this measure was Leah Modigliani, a strategist at Morgan Stanley. Without taking anything away from her accomplishments, it is worth noting that she is the daughter of Franco Modigliani, Nobel price winner in Economics.

Expected return = Riskfree rate + Beta (Expected return on market – Riskfree rate)

In chapter 2, we used the model to estimate the expected returns for the next period, using the current riskfree rate, the beta and the average premium earned by stocks over the riskfree rate as inputs. In this section, we will consider how the capital asset pricing model can be adapted to judge past performance.

i. Excess Return (Alpha or Jensen's Alpha)

The simplest way to use to the capital asset pricing model to evaluate performance is to compare the actual return to the return your investment or strategy should have made over the evaluation period, given its beta and given what the market did over the period. As an example, assume that you are analyzing a strategy that generated 12% in returns over the last year. Assume that your calculations indicate that the strategy has a beta of 1.2, and that the riskfree rate at the beginning of the last year was 4% and that the return on the market over the last year was 11%. You can compute the excess return as follows:

Expected return over last year = 4% + 1.2 (11% - 4%) = 12.4%

Excess Return = Actual Return – Expected Return = 12% - 12.4% = -0.4%

This strategy under performed the market, after adjusting for risk, by 0.4%. This excess return is also called an abnormal return.

What are the differences between what we are doing here and what we did in chapter 2 to forecast expected returns in the future? The first is that we use the riskfree rate at the beginning of the evaluation period when we do evaluation, whereas we use the current riskfree rate when making forecasts. The second is that we use the actual return on the market over the period, even if it is negative, when we do evaluation rather than the historical or an implied equity premium that we use when computing forecasted returns. Finally, the beta we use in evaluation should measure the risk you were exposed to during your evaluation period, while a forward-looking beta should be used for forecasts.

The excess returns on a strategy can be computed for any return period you want – daily, weekly, monthly or annual. You would need to adjust your riskfree rate and market return appropriately, using, for instance, the weekly riskfree rate and a weekly market return if you want to compute weekly excess returns from a strategy. An alternative approach to estimating the excess return, which should yield the same results, is to run a regression of the returns on your strategy, in excess of the riskfree rate, against return on a market index, in excess of the riskfree rate.

(Return on strategy – Riskfree rate) = a + b (Returns on market index – Riskfree rate)

The slope of this regression gives you the historical beta, but the intercept of this regression yields the excess return by period for your strategy.[6] Using the statistical term for the intercept, the excess return is often called an alpha. In some quarters, it is called Jensen's alpha, reflecting the fact that it was first used in a study of mutual funds in the 1960s by Michael Jensen, one of the pioneers in empirical finance from the University of Chicago.

There is one final point that should be made about excess returns. When you compute excess returns by day or week over a longer period (say six months or a year), you may also want to compute how the strategy performed over the entire period. To do this, you usually look at the compounded return over the period. This compounded return is called a cumulative excess return or a cumulative abnormal return (CAR). Defining the excess return in each interval as ER_t, you can write the excess return over a period as follows:

Cumulative Abnormal return over n intervals= $(1+ ER_1)\ (1+ ER_2)\ (1+ ER_3)\ldots\ (1+ ER_n)$

A cumulative abnormal return that is greater than zero indicates that the strategy beat the market, at least over the period of your test.

Unlike the variance-based measures in the last section, Jensen's alpha does not penalize sector-specific funds that are not diversified because it looks at the beta of a portfolio and not its standard deviation. The measure's fidelity to the capital asset pricing model, however, exposes it to all of the model's limitations. Since the model has historically under estimated the expected returns of small cap stocks, with low PE and low price to book ratios, you will tend to find that strategies that focus on stocks with these characteristics earn positive excess returns.

There are variations that have appeared on Jensen's alpha. An early variation replaced the capital asset pricing model with what is commonly called the market model, where the expected return on an investment is based upon a past regression alpha.[7] In the

---

[6] To see why, lets work through the algebra. The expected return in the CAPM can be written as:

Expected Return on Strategy = Riskfree Rate + Beta (Return on market – Riskfree rate)

Expected Return on Strategy – Riskfree Rate = Beta (Return on market – Riskfree rate)

In other words, if your stock did exactly as predicted by the CAPM, the intercept should be zero. If the intercept is different from zero, that must indicate under performance (if it is negative) or out performance (if it is positive)

[7] In the market model, the excess return is written as

Excess return = Actual return – (a + b Return on Market)

Where a is the intercept and b is the slope of a regression of returns on the stock against returns on the market index.

last decade, for instance, researchers have developed a version of the measure that allows the beta to change from period to period for a strategy – these are called time varying betas. This is clearly more realistic than assuming one beta for the entire testing period.

ii. Treynor Index

The excess return is a percentage measure. But is earning a 1% excess return over an expected return of 15% equivalent to earning a 1% excess return over an expected return of 7%? There are many who would argue that the latter strategy is a more impressive one. The Treynor Index attempts to correct for this by converting the excess return into a ratio, relative to the beta.[8] It is computed by dividing the difference between the returns on a strategy and the riskfree rate by the beta of the investment. This value is then compared to the difference between the returns on the market and the riskfree rate.

Treynor Index = (Return on Strategy – Riskfree Rate)/ Beta

To illustrate, assume that you are considering the strategy that we described in the last section with a beta of 1.2 that earned a return of 12% in the most recent year. In that example, the return on the market over the same year was 11% and the riskfree rate was 4%. The Treynor Index for this strategy would be

Treynor Index for Strategy= (12% - 4%)/ 1.2 = 6.67%

Treynor Index for Market = (11% - 4%)/ 1 = 7.00%

This strategy underperformed the market.

The Treynor Index is closely related to the alpha measure described in the last section. The measures will always agree on whether a strategy under or out performs the market, but will disagree on rankings. The Treynor Index will rank lower beta strategies higher than the alpha measure, because it looks at excess returns earned per unit beta.

Illustration 6.2: Estimating Jensen's Alpha and Treynor's Index: Large U.S. Mutual Funds

In table 6.2 below, we estimate Jensen's alpha and Treynor's Index for the 15 largest mutual funds in the United States. For simplicity, we assumed that the average annual risk free rate during the period was 5%.

Table 6.2: Jensen's Alpha and Treynor Index: Large Mutual Funds – 1997 to 2001

Fund Name	*Return on Fund*	*Beta*	*Expected Return*	*Jensen's Alpha*	*Treynor's Index*
Fidelity Magellan	10.72%	1.02	10.46%	0.26%	5.61%

[8] This measure is named after Jack Treynor, who used it to evaluate mutual fund performance.

Vanguard 500 Index	10.14%	1.00	10.35%	-0.21%	5.14%
American Funds Inv Co Amer A	12.84%	0.71	8.80%	4.04%	11.04%
American Funds Wash Mutual A	12.36%	0.58	8.10%	4.26%	12.69%
American Funds Growth Fund A	17.66%	1.07	10.72%	6.94%	11.83%
Fidelity Growth & Income	10.32%	0.70	8.75%	1.58%	7.60%
Fidelity Contrafund	11.47%	0.67	8.58%	2.89%	9.66%
American Funds New Prspctv A	11.82%	0.82	9.39%	2.43%	8.32%
American Funds EuroPacific A	7.20%	0.78	9.17%	-1.97%	2.82%
AmCent Ultra Inv	10.17%	1.26	11.74%	-1.57%	4.10%
Janus	9.57%	1.35	12.22%	-2.65%	3.39%
Vanguard Instl Index	10.27%	1.00	10.35%	-0.08%	5.27%
Fidelity EquityIncome	10.39%	0.70	8.75%	1.65%	7.70%
S&P 500	*10.35%*	*1*	*10.35%*	*0.00%*	*5.35%*

While Fidelity Magellan and the Vanguard 500 index fund have alphas close to zero and Treynor indices that match the market, the American Funds Growth Fund earned an annual alpha of 6.94% over the period.

c. Arbitrage Pricing and Multi-factor Models

In chapter 2, we noted that the assumptions that we need to arrive at the single market beta measure of risk in the capital asset pricing model are unrealistic and that the model itself systematically under estimates the expected returns for stocks with certain characteristics – low market capitalization and low PE. We considered the alternative of the arbitrage pricing model, which allows for multiple market risk factors which are unidentified or a multi-factor model, which relates expected returns to a number of macro-economic factors such as interest rates, inflation and economic growth. These models, we argued, allow us more flexibility when it comes to estimating expected returns.

You could use either the arbitrage pricing or multi-factor model to estimate the return you would have expected to earn over a period on a portfolio and compare this return to the actual return earned. In other words, you could compute an excess return on alpha for a strategy or portfolio using these models instead of the capital asset pricing model.

To the extent that the arbitrage pricing and multi-factor models are less likely to yield biased returns for small cap and low PE stocks, you could argue that the excess returns from these models should give you better measures of performance. The biggest problem that you run into in using these models to evaluate the excess returns earned by a portfolio manager or a strategy is that the portfolios themselves may be constantly shifting. What you measure as an alpha from these models may really reflect your failure to correct for the variation in exposure to different market risk factors over time. While this is also a problem with the capital asset pricing model, it is far easier to adjust a single beta over time than it is to work with multiple betas.

d. Proxy and Composite Models

The alternative to conventional risk and return models is the use of a proxy model, where the returns on stocks are correlated with observable financial characteristics of the firm. Perhaps the best known proxy model was the one developed by Fama and French, which we presented in chapter 2. They found that between 1962 and 1990, stocks with lower market capitalization and price to book ratios consistently earned higher returns than larger market capitalization companies with higher price to book ratios. In fact, market capitalization and price to book ratio differences across firms explained far more of the variation in actual returns than betas did.

Building on this theme, traditional risk and return models may fall short when it comes to estimating expected returns for portfolios that have disproportionately large exposures to small cap or low price to book value stocks. These portfolios will look like they earn excess returns. Using a proxy model, where the returns on the portfolio are conditioned on the market cap of the stocks held in the portfolio and their price to book ratios may eliminate this bias:

Expected return on portfolio = a + b (Average Market Capitalization)$_{\text{Portfolio}}$ + c (Average Price to Book Ratio)$_{\text{Portfolio}}$

This model can even be expanded to include a conventional market beta, yielding what is often called a three-factor model:

Expected return on portfolio = a + b (Market beta) + c (Average Market Capitalization)$_{\text{Portfolio}}$ + d (Average Price to Book Ratio)$_{\text{Portfolio}}$

The perils of incorporating variables such as market capitalization and price to book ratios into expected returns is that you run a risk of creating a self-fulfilling prophecy. If markets routinely misprice certain types of companies – small companies, for instance – and we insist on including these variables in the expected return regressions, we will be biased, with a complete enough model, towards finding that markets are efficient. In fact, in recent years, researchers have added a fourth factor – price momentum – to these factor models, because of recent findings that companies that have done well in the recent past are likely to continue doing well in the future.

Closing Thoughts

There are two closing points that we would like to emphasize about the use of risk and return models and tests of market efficiency. The first is that a test of market efficiency is a joint test of market efficiency and the efficacy of the model used for expected returns. When there is evidence of excess returns in a test of market efficiency, it can indicate that markets are inefficient or that the model used to compute expected returns is wrong or both. While this may seem to present an insoluble dilemma, if the conclusions of the study are insensitive to different model specifications, it is much more likely that the results are being driven by true market inefficiencies and not just by model misspecifications.

In terms of which approach you should use to come up with expected returns, it is worth noting that each approach has its own built in biases that you need to be aware of. Table 6.3 below summarizes the alternative approaches to evaluating returns and the types of strategies and portfolios that they are likely to be biased towards and against.

Strategies for testing market efficiency

There are a number of different ways of testing for market efficiency, and the approach used will depend in great part on the investment scheme being tested. A scheme based upon trading on information events (stock splits, earnings announcements or acquisition announcements) is likely to be tested using an 'event study' where returns around the event are scrutinized for evidence of excess returns. A scheme based upon trading on a observable characteristic of a firm (price earnings ratios, price book value ratios or dividend yields) is likely to be tested using a 'portfolio' approach, where portfolios of stocks with these characteristics are created and tracked over time to see if, in fact, they make excess returns. An alternative way of testing to see if there is a relationship between an observable characteristic and returns is to run a regression of the latter on the former. This approach allows for more flexibility if you are testing for interactions among variables. The following pages summarize the key steps involved in each of these approaches, and some potential pitfalls to watch out for when conducting or using these tests.

A. Event Study

An event study is designed to examine market reactions to, and excess returns around specific information events. The information events can be market-wide, such as macro-economic announcements, or firm-specific, such as earnings or dividend announcements. The steps in an event study are as follows -

(1) <u>The event to be studied is clearly identified, and the date on which the event was announced pinpointed</u>. The presumption in event studies is that the timing of the event is known with a fair degree of certainty. Since financial markets react to the information about an event, rather than the event itself, most event studies are centered around the announcement date[9] for the event.

Announcement Date

____________________________________|____________________________________

(2) <u>Once the event dates are known, returns are collected around these dates for each of the firms in the sample</u>. In doing so, two decisions have to be made. First, you have to decide whether to collect weekly, daily or shorter-interval returns around the event. This will, in part, be decided by how precisely the event date is known (the more precise, the more likely it is that shorter return intervals can be used) and by how quickly information is reflected in

[9] In most financial transactions, the announcement date tends to precede the event date by several days and, sometimes, weeks.

prices (the faster the adjustment, the shorter the return interval to use). Second, you have to determine how many periods of returns before and after the announcement date will be considered as part of the 'event window'. That decision also will be determined by the precision of the event date, since more imprecise dates will require longer windows.

R_{-jn} .... R_{j0} ....R_{+jn}

______________|________________________|_________________________|__________

Return window: -n to +n

where,

R_{jt} = Returns on firm j for period t (t = -n, ...,0, ... +n)

(3) The returns, by period, around the announcement date, are adjusted for market performance and risk to arrive at excess returns for each firm in the sample. You could use any of the risk and return models described in the last section to estimate excess returns. For instance, if the capital asset pricing model is used to control for risk -

Excess Return on period t = Return on day t – (Riskfree rate + Beta * Return on market on day t)

ER_{-jn} .... ER_{j0} ....ER_{+jn}

______________|________________________|_________________________|__________

Return window: -n to +n

where,

ER_{jt} = Excess Returns on firm j for period t (t = -n, ...,0, ... +n)

You can also look at how a portfolio held over multiple periods would have done by measuring a cumulated abnormal return (CAR) by compounding the excess returns over the periods. Thus, if your excess return on day 1 is +2%, day 2 is –1% and day 3 is +1.5%, your cumulative excess return over all three days would be:

Cumulated Excess Return = $(!+ ER_1)\ (!+ ER_2)\ (!+ ER_3) – 1$

$= (1.02)\ (0.99)(1.015) –1 = 1.02495$ or 2.495%

(4) Once the excess returns are estimated for each firm in the sample, the average excess returns can be computed across the firms and it will almost never be equal to zero. To test to see whether this number is significantly different from zero, however, you need a statistical test. The simplest is to compute a standard deviation in the excess returns across the sampled firms, and to use this to estimate a t statistic. Thus, if you have N firms in your sample and you have computed the excess returns each day for these firms:

$$\text{Average excess return on day t} = \sum_{j=1}^{j=N} \frac{ER_{jt}}{N}$$

T statistic for excess return on day t = Average Excess Return / Standard Error

You can then check to see if the t statistics are statistically significant. For instance, if the t statistic is 2.33 or higher, there is a 99% chance that the average excess return is different from zero. If the average is positive, the event increases stock prices, whereas if it is negative, the event decreases stock prices.

Illustration 6.3: Example of an event study - Effects of Option Listing on Stock prices

Academics and practitioners have long argued about the consequences of option listing for stock price volatility. On the one hand, there are those who argue that options attract speculators and hence increase stock price volatility. This higher risk, they argue, should lead to lower stock prices. On the other hand, there are others who argue that options increase the available choices for investors and increase the flow of information to financial markets, and thus lead to lower stock price volatility and higher stock prices.

One way to test these alternative hypotheses is to do an event study, examining the effects of listing options on the underlying stocks' prices. In1989, Conrad did such a study, following these steps -

Step 1: The date on which the announcement that options would be listed on the Chicago Board of Options on a particular stock was made was collected.

Step 2: The prices of the underlying stocks(j) were collected for each of the ten days prior to the option listing announcement date, the day of the announcement, and each of the ten days after.

Step 3: The returns on the stock (R_{jt}) were computed for each of these trading days.

Step 4: The beta for the stock (β_j) was estimated using the returns from a time period outside the event window (using 100 trading days from before the event and 100 trading days after the event).

Step 5: The returns on the market index (R_{mt}) were computed for each of the 21 trading days.

Step 6: The excess returns were computed for each of the 21 trading days -

$$ER_{jt} = R_{jt} - \beta_j R_{mt} \qquad .... \quad t = -10,-9,-8,.....,+8,+9,+10$$

The excess returns are cumulated for each trading day.

Step 7: The average and standard error of excess returns across all stocks with option listings were computed for each of the 21 trading days. The t statistics are computed using the averages and standard errors for each trading day. Table 6.4 summarizes the average excess returns and t statistics around option listing announcement dates –

Table 6.4: Excess Returns around Option Listing Announcement Dates

Trading Day	*Average Excess* Return	*Cumulative* Excess Return	*T Statistic*

-10	0.17%	0.17%	1.30
-9	0.48%	0.65%	1.66
-8	-0.24%	0.41%	1.43
-7	0.28%	0.69%	1.62
-6	0.04%	0.73%	1.62
-5	-0.46%	0.27%	1.24
-4	-0.26%	0.01%	1.02
-3	-0.11%	-0.10%	0.93
-2	0.26%	0.16%	1.09
-1	0.29%	0.45%	1.28
0	0.01%	0.46%	1.27
1	0.17%	0.63%	1.37
2	0.14%	0.77%	1.44
3	0.04%	0.81%	1.44
4	0.18%	0.99%	1.54
5	0.56%	1.55%	1.88
6	0.22%	1.77%	1.99
7	0.05%	1.82%	2.00
8	-0.13%	1.69%	1.89
9	0.09%	1.78%	1.92
10	0.02%	1.80%	1.91

Based upon these excess returns, there is no evidence of an announcement effect on the announcement day alone, but there is mild[10] evidence of a positive effect over the entire announcement period.

How different is different? Statistics versus Economic Significance

If you compare two samples on any dimension, you will get different results. Thus, you could compare the average returns on portfolios of companies with tall CEOs to the average returns of portfolios of companies with short CEOs, and you would find them to be different. But what should we read into the difference? If the average return on companies headed by tall CEOs is higher, should we rush out to buy stock in those companies? Not quite yet, because the differences often arise purely from chance.

The first test that you can run is a statistical test, where you apply the laws the probability to estimate the likelihood that the difference you are observing is purely random. This is what we do, for instance, when we compute a t statistic on abnormal returns. If the t statistics is 2.33, for example, we are saying that there is only a 1% chance that the difference we are observing is random and a 99% chance that returns are higher on companies with tall CEOs. If the t statistic has been only 0.50, there would have been a 31%

[10] The t statistics are marginally significant at the 5% level.

chance that the difference was purely random. In fact, it is common to test for statistical significance at the 1% or 5% levels, i.e. only differences where the probability of randomness is less than 1% or 5% would be viewed as statistically significant.

What is economic significance? When the difference between two samples is economically significant, you can make money off the difference. In the example that we have used, you could buy stocks with tall CEOs and sell stocks with short CEOs and make excess returns on your investment. Statistical significance does not always equal economic significance. First, you may have transactions costs that are much higher than the difference in returns between the two groups. Note that the larger the sample you use, the more likely it is that even small differences can be statistically significant. Thus, a difference of 0.20% may be statistically significant, but it clearly is not sufficient to cover execution costs. Second, you have the thorny issue of causation not being equal to correlation. In other words, all you have established is a correlation between returns and CEO height, but you have not established causation. Do firms with tall CEOs earn higher returns or do firms with higher return hire tall CEOs? If it is the latter, you may very well find statistical significance but not economic profits.

B. Portfolio Study

In some investment strategies, firms with specific characteristics are viewed as more likely to be undervalued, and therefore have excess returns, than firms without these characteristics. In these cases, the strategies can be tested by creating portfolios of firms possessing these characteristics at the beginning of a time period, and examining returns over the time period. To ensure that these results are not colored by the idiosyncracies of one time period, this analysis is repeated for a number of periods. The steps in doing a portfolio study are as follows -

(1) The variable on which firms will be classified is defined, using the investment strategy as a guide. This variable has to be observable, though it does not have to be numerical. Examples would include market value of equity, bond ratings, stock price, price earnings ratios and price book value ratios.

(2) The data on the variable is collected for every firm in the defined sample[11] at the start of the testing period, and firms are classified into portfolios based upon the magnitude of the variable. Thus, if the price earnings ratio is the screening variable, firms are classified on the

[11] Though there are practicial limits on how big the sample can be, care should be taken to make sure that no biases enter at this stage of the process. An obvious one would be to pick only stocks that have done well over the time period for the universe.

basis of PE ratios into portfolios from lowest PE to highest PE classes. The number of classes will depend upon the size of the sample, since there have to be sufficient firms in each portfolio to get some measure of diversification.

(3) The returns are collected for each firm in each portfolio for the testing period, and the returns for each portfolio are computed, making an assumption about how stocks will be weighted – some studies use equal weightings whereas others are value weighted.

(4) The beta (if using a single factor model like the CAPM) or betas (if using a multifactor model like the arbitrage pricing model) of each portfolio are estimated, either by taking the average of the betas of the individual stocks in the portfolio or by regressing the portfolio's returns against market returns over a prior time period (for instance, the year before the testing period).

(5) The excess returns earned by each portfolio are computed, in conjunction with the standard error of the excess returns.

(6) There are a number of statistical tests available to check whether the average excess returns are, in fact, different across the portfolios. Some of these tests are parametric[12] (they make certain distributional assumptions about excess returns) and some are non-parametric[13].

(7) As a final test, the extreme portfolios can be matched against each other to see whether there are statistically significant differences across these portfolios.

Illustration 6.3: Example of a portfolio study - Price Earnings Ratios

Practitioners have claimed that low price-earnings ratio stocks are generally bargains and do much better than the market or stocks with high price earnings ratios. This hypothesis can be tested using a portfolio approach -

Step 1: Using data on price-earnings ratios from the end of 1987, firms on the New York Stock Exchange were classified into five groups, the first group consisting of stocks with the lowest PE ratios and the fifth group consisting of stocks with the highest PE ratios. Firms with negative price-earnings ratios were ignored.

Step 2: The returns on each portfolio were computed using data from 1988 to 1992. Stocks which went bankrupt or were delisted were assigned a return of -100%.

---

[12] One parametric test is an F test, which tests for equality of means across groups. This test can be conducted assuming either that the groups have the same variance, or that they have different variances.

[13] An example of a non-parametric test is a rank sum test, which ranks returns across the entire sample and then sums the ranks within each group to check whether the rankings are random or systematic.

Step 3: The betas for each stock in each portfolio were computed using monthly returns from 1983 to 1987, and the average beta for each portfolio was estimated. The portfolios were assumed to be equally-weighted.[14]

Step 4: The returns on the market index were computed each year from 1988 to 1992.

Step 5: The excess returns on each portfolio were computed each year, using the actual returns estimated from step 2, the betas estimated from step 3 and the market returns from step 4:

Excess Return in year t= Actual return on Portfolio in year t - (Riskfree Rate at the start of year t– Beta * (Return on market in year t – Riskfree rate at the start of year t))

Table 6.5 summarizes the excess returns each year from 1988 to 1992 for each portfolio.

Table 6.5: Excess Returns from 1988 to 1992 for PE Ratio Portfolios

P/E Class	*1988*	*1989*	*1990*	*1991*	*1992*	*1988-1992*
Lowest	3.84%	-0.83%	2.10%	6.68%	0.64%	2.61%
2	1.75%	2.26%	0.19%	1.09%	1.13%	1.56%
3	0.20%	-3.15%	-0.20%	0.17%	0.12%	-0.59%
4	-1.25%	-0.94%	-0.65%	-1.99%	-0.48%	-1.15%
Highest	-1.74%	-0.63%	-1.44%	-4.06%	-1.25%	-1.95%

Step 6: While the ranking of the returns across the portfolio classes seems to confirm our hypothesis that low PE stocks earn a higher return, we have to consider whether the differences across portfolios are statistically significant. There are several tests available, but these are a few:

- An F test can be used to accept or reject the hypothesis that the average returns are the same across all portfolios. A high F score would lead us to conclude that the differences are too large to be random.

  F statistic for difference across PE portfolios between 1988 and 1992 = 14.75

  This suggests that there is less than a 1% chance that the difference between the portfolios is random.
- A chi-squared test is a non-parametric test that can be used to test the hypothesis that the means are the same across the five portfolio classes.

  Chi-squared statistic for difference across PE portfolios: 1988-92 = 36.16

  This confirms our conclusion from the F test that the differences are statistically significant.

---

[14] This will be a function of your strategy. If your strategy requires market-cap weighted holdings, you would have to modify the test accordingly.

- We could isolate just the lowest PE and highest PE stocks and estimate a t statistic that the averages are different across these two portfolios. In this case, the t statistics that we obtain when we compare the returns on the lowest and highest PE ratio classes is 5.61. This difference is also statistically significant.

C. Regressions

One of the limitations of portfolio studies is that they become increasing unwieldy, as the number of variables that you use in your strategy increases. For instance, assume that you pick stocks that have low PE ratios, low institutional investment and whose stock prices have done well in the last six months. You could categorize all firms in your sample into five portfolios, based upon each variable, but you would end up with 125 portfolios overall because of the potential interactions among the variables. The other problem with portfolio studies is that you group firms into classes and ignore differences across firms within each class. Thus, the stocks in the lowest PE ratio class may have PE ratios that range from the 4 to 12. If you believe that these differences may affect the expected returns on your strategy, you could get a better measure of the relationship by running a multiple regression. Your dependent variable would be the returns on stocks and the independent variables would include the variables that form your strategy. There are four steps to running a regression and they are listed below:

Step 1: Identify your dependent variable. This is the variable that you are trying to explain. In most investment schemes, it will be a measure of the return you would make on the investment but you have to make at least two judgments. The first is whether you plan to use total returns or excess returns; with the latter, you would adjust the returns for risk and market performance, using one the measures discussed earlier in the chapter. The second decision you have to make is on the return interval you will be using – monthly, quarterly, annual or five-year, for instance. This choice will be determined both by your investment strategy – long-term strategies require long-term returns – and the ease with which you can get data on your independent variables for the intervals. For instance, if you use accounting variables such as earnings or book value as independent variables, you will be able to get updates only once every quarter for these variables.

Step 2: Decide on how you will measure the variables that will underlie you strategy. For instance, in the example cited above, you will have to define PE ratios, institutional investment and stock price momentum with more specificity. With PE ratios, you will have to choose between different measures of earnings – primary or diluted, before or after extraordinary items, current or trailing. With institutional investment, you could measure the institutional holdings as a percent of outstanding stock or as a percent of float (stock that is traded) and you will also have to decide whether you will consider all institutional investors or only certain kinds (mutual funds, pension fund etc.). With stock price momentum, you may have to choose between percent changes over the previous six months, which will bias you towards lower priced stocks or absolute changes, which will bias you towards higher priced stocks. Once you determined your independent variables, you will have to collect information on them at the beginning of each of your testing periods. For instance, if you decide that annual returns in 2000 will be your dependent variable, you will have to collect information on PE ratios and institutional holdings from January 1, 2000, and stock price momentum from June 30, 1999 to January 1, 2000.[15]

Step 3: You should check for the nature of the relationship between the dependent variable and each independent variable. A scatter plot provides a simple graphical tool for doing this. You are checking to see not only if there is a relationship but also for whether the relationship is linear. Figure 6.1 presents two scatter plots:

Figure 6.1: Scatter Plots – Linear and Non-linear Relationships

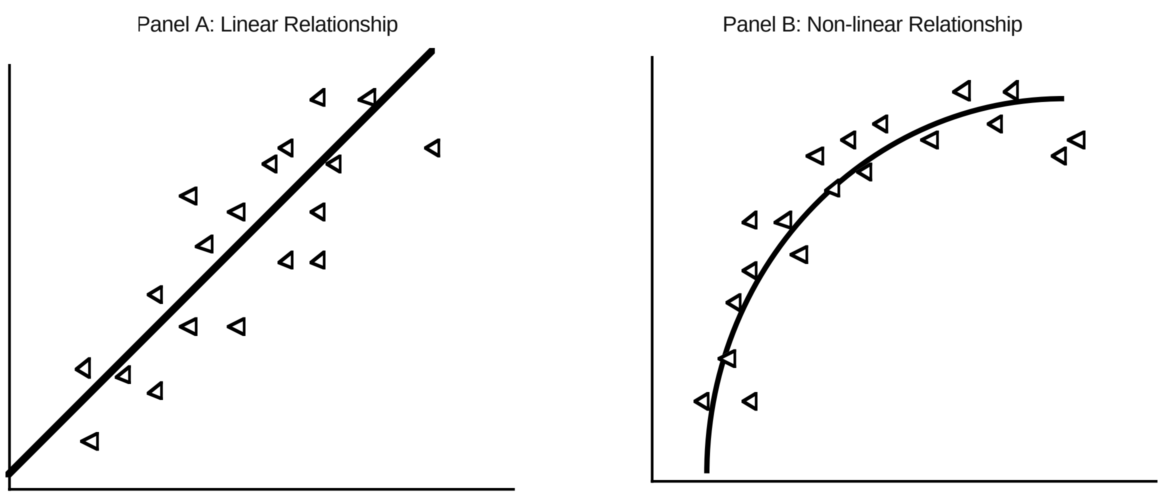

Panel a represents a scatter plot with a linear relationship but Panel b is more consistent with a non-linear relationship. If you observe the latter, you may have to transform the variable to make the relationship more linear.[16]

---

[15] For institutional holdings, you will have to use whatever values you would have been able to obtain from public sources as of January 1, 2000. Since there is a delay before institutions file with the SEC, you may not know the holdings as of January 1 until much later in the year.

[16] Transformation requires you to convert a number by taking a mathematical function of it. Some commonly used transformations include the natural log, square root and square. The natural log transformation is probably the most useful one in financial research.

Step 4: You can now run the regression of the dependent variable against the independent variables, with or without transformations. In the example noted above, for instance, you would regress returns against PE ratios, institutional holdings as a percent of the stock outstanding and the price change over the last 6 months:

Return on Stock = a + b (PE) + c (Institutional Holdings as % of Stock) + d (Stock price change over last 6 months)

If your hypothesis is right, you should expect to see the following:

b < 0 : Stocks with higher PE ratios should have lower returns

c < 0: Stocks with higher institutional holdings should have lower returns

d > 0: Stocks that have done well over the last 6 months should have higher returns

Once you run the regression, you have to pass it through the tests for statistical significance. In other words, even if all of the coefficients have the right signs, you have to check to ensure that they are significantly different from zero. In most regressions, statistical significance is estimated with a t statistic for each coefficient. This t statistic is computed by dividing the coefficient by the standard error of the coefficient. You can also compute an F statistic to measure whether the regression collectively yield statistically significant results.

The regression described above, where you look for differences across observations (firms, funds or countries) at a point in time is called a cross sectional regression. You can also use regressions to analyze how a variable changes over time as other variables change. For instance, it also long been posited that PE ratios for all stocks go up as interest rates go down and economic growth increases. You could look over at the PE ratios for the entire market each year for the last 40 years for instance, and examine whether PE ratios have changed as interest rates and economic growth have changes. This regression is called a time series regression. Some inventive analysts even combine cross sectional and time series data to and create pooled regression.

The Limits of Regressions

Regressions are incredibly powerful tools to examine relationships but they have their limits when it comes to testing market efficiency. The first problem that they share with all other tools is that they are only as good as the data that goes into them. If your data is filled with errors, you should expect the regression output to reflect that. The second problem is that you make assumptions about the nature of the relationship between the dependent and independent variables that may not be true. For instance, if you run a regression of returns against institutional holdings as a percent of outstanding stock, you are assuming a linear relationship between the two, i.e, that returns will change by the same magnitude if holdings go from 10 to 20% as they would if holdings went from 20 to 30%.

The third problem arises when you run multiple regressions. For the regression coefficients to be unbiased, the independent variables should be uncorrelated with each other. In reality, it is difficult to find independent variables that have this characteristic.

The Cardinal Sins in testing Market Efficiency

In the process of testing investment strategies, there are a number of pitfalls that have to be avoided. Some of them are listed below -

1. Using 'anecdotal evidence' to support/reject an investment strategy: Anecdotal evidence is a double-edged sword. It can be used to support or reject the same hypothesis. Since stock prices are noisy and all investment schemes (no matter how absurd) will succeed sometimes and fail at other times, there will always be cases where the scheme works.

2. Testing an investment strategy on the same data and time period from which it was extracted: This is the tool of choice for the unscrupulous investment advisor. An investment scheme is extracted from hundreds through an examination of the data for a particular time period. This investment scheme is then tested on the same time period, with predictable results. (The scheme does miraculously well and makes immense returns.) An investment scheme should always be tested out on a time period different from the one it is extracted from or on a universe different from the one used to derive the scheme.

3. Sampling Biases: Since there are thousands of stocks that could be considered part of the testable universe of investmetns, researchers often choose to use a smaller sample. When this choice is random, this does limited damage to the results of the study. If the choice is biased, it can provide results that are not true in the larger universe. Biases can enter in subtle ways. For instance, assume that you decide to examine whether stocks with low prices are good investments and you test this by estimating the returns over the last year for stocks that have low prices today. You will almost certainly find that this portfolio does badly but not because your underlying hypothesis is false. Stocks that have gone down over the last year are more likely to have low stock prices today than stocks that have gone up. By looking at stock prices today, you created a sample that is biased towards poorly performing stocks. You could very easily have avoided this bias by looking at stock prices at the start of your return period (rather than the end of the period).

4. Failure to control for market performance: A failure to control for overall market performance can lead you to conclude that your investment scheme works just because it makes good returns or does not work just because it makes poor returns. Most investment strategies will generate good returns in a period in which the market does well and few will do so when the market does badly. It is crucial therefore that investment schemes control for market performance during the period of the test.

5. Failure to control for risk: A failure to control for risk leads to a bias towards accepting high-risk investment schemes and rejecting low-risk investment schemes, since the former should make higher returns than the market and the latter lower, without implying any excess returns. For instance, a strategy of investing in the stock of bankrupt companies may generate annual returns that are much higher than returns on the S&P 500, but it is also a much riskier strategy and has to be held to a higher standard.

6. *Mistaking correlation for causation*: Statistical tests often present evidence of correlation, rather than causation. Consider the study on PE stocks cited in the earlier section. We concluded that low PE stocks have higher excess returns than high PE stocks. It would be a mistake to conclude that a low price earnings ratio, by itself, causes excess returns, since the high returns and the low PE ratio themselves might have been caused by the high risk associated with investing in the stocks. In other words, high risk is the causative factor that leads to both the observed phenomena – low PE ratios on the one hand and high returns on the other. This insight would make us more cautious about adopting a strategy of buying low PE stocks in the first place.

Some lesser sins that can be a problem

1. Data Mining: The easy access that we have to huge amounts of data on stocks today can be a double-edged sword. While it makes it far easier to test investment strategies, it also exposes us to the risk of what is called data mining. When you relate stock returns to hundreds of variables, you are bound to find some that seem to predict returns, simply by chance. This will occur even if you are careful to sample without bias and test outside your sample period.

2. Survival Bias: Most researchers start with a existing universe of publicly traded companies and work back through time to test investment strategies. This can create a subtle bias since it automatically eliminates firms that failed during the period, with obvious negative consequences for returns. If the investment scheme is particularly susceptible to picking firms that have high bankruptcy risk, this may lead to an overstatement of returns on the scheme. For example, assume that the investment scheme recommends investing in stocks that have very negative earnings, using the argument that these stocks are most likely to benefit from a turnaround. Some of the firms in this portfolio will go bankrupt, and a failure to consider these firms will overstate the returns from this strategy.

3. Not allowing for Transactions Costs: Some investment schemes are more expensive than others because of transactions costs - execution fees, bid-ask spreads and price impact. A complete test will take these into account before it passes judgment on the strategy. This is easier said than done, because different investors have different transactions costs, and it is unclear which investor's trading cost schedule should be used in the test. Most researchers who ignore transactions costs argue that individual investors can decide for themselves, given their transactions costs, whether the excess returns justify the investment strategy.

4. Not allowing for difficulties in execution: Some strategies look good on paper but are difficult to execute in practice, either because of impediments to trading or because trading creates a price impact. Thus a strategy of investing in very small companies may seem to create excess returns on paper, but these excess returns may not exist in practice because the price impact is significant.

Conclusion

The question of whether markets are efficient will always be a provocative one, given the implications that efficient markets have for investment management and research. If an efficient market is defined as one where the market price is an unbiased estimate of the true value, it is quite clear that some markets will always be more efficient than others and that markets will always be more efficient to some investors than to others. The capacity of a market to correct inefficiencies quickly will depend, in part, on the ease of trading, the transactions cost and the vigilance of profit-seeking investors in that market.

While market efficiency can be tested in a number of different ways, the three most widely used tests to test efficiency are 'event studies' which examine market reactions to information events, 'portfolio studies' which evaluate the returns of portfolios created on the basis of observable characteristics and regressions that relate returns to firm characteristics either at a point in time or across time. It does make sense to be vigilant, because bias can enter these studies, intentionally or otherwise, in a number of different ways and can lead to unwarranted conclusions, and, worse still, wasteful investment strategies.

Lessons for Investors

1. An efficient market makes mistakes but the mistakes tend to be random. In other words, you know that some stocks are undervalued and some are overvalued but you have no way of identifying which group each stock falls into.
2. You are more likely to find inefficiencies in markets that are less liquid and where information is less easily available or accessible.
3. In an inefficient market, you can use publicly available information to find under and over valued stocks and trade on them to earn returns that are consistently greater than what you would have earned on a randomly selected portfolio of equivalent risk.
4. To create portfolios of equivalent risk, you have to use models for risk and return. To the extent that your model for risk is mis-specified, you may uncover what look like inefficiencies but really represent the failures of your model.

CHAPTER 7

SMOKE AND MIRRORS? CHARTING AND TECHNICAL ANALYSIS

Charts have been around as long as there have been markets. Some investors have always believed that charts of past prices provide signals of the future, and have pored over charts looking for patterns that predict price movements. Notwithstanding the disdain with which they are viewed by other investors and many academics, easy access to data combined with an increase in computing capabilities – charting and graphing programs abound – has meant that more investors look at charts now than ever before. In addition, data on trading volume and from derivatives markets have provided chartists with new indicators of the future.

In this chapter, we will look at the basis of charting by examining the underlying premise in charting and technical analysis, which is a belief that there are systematic and often irrational patterns in investor behavior and that technical indicators and charts provide advance warning of shifts in investor behavior. While we will not attempt to describe every charting pattern and technical indicator – there are hundreds – we will categorize them based upon the view of human behavior that underlies each. In the process, we will see if there are lessons in charts that even non-believers can take away with them and cautionary notes for true believers about potential inconsistencies.

Random Walks and Price Patterns

In many ways, the antithesis of charting is the notion that prices follow a random walk. In a random walk, the stock price reflects the information in past prices, and knowing what happened yesterday is of no consequence to what will happen today. Since the random walk comes in for a fair degree of abuse, some justified and some not, from technical analysts, we will begin by looking at what the random walk is and its implications.

The Basis for Random Walks

To understand the argument for prices following a random walk, we have to begin with the presumption that investors at any point in time estimate the value of an asset based upon expectations of the future, and that these expectations are both unbiased and rational, given the information that investors have at that point in time. Under these conditions, the price of the asset changes only as new information comes out about it. If the market price at any point in time is an unbiased estimate of value, the next piece of information that comes out about the asset should be just as likely to contain good news as bad.[1] It therefore follows that the next price change is just as likely to be positive as it likely to be negative. The implication of course is that each price change will be independent of the previous one, and that knowing an asset's price history will not help form better predictions of future price changes. Figure 7.1 summarizes the assumptions.

Figure 7.1: Information and Price Changes in a Rational Market

	Current	Next period
Information	All information about the firm is publicly available and traded on.	New information comes out about the firm.
Market Expectations	Investors form unbiased expectations about the future	Since expectations are unbiased, there is a 50% chance of good or bad news.
Price Assessment	Stock price is an unbiased estimate of the value of the stock.	The price changes in accordance with the information. If it contains good (bad) news, relative to expectations, the stock price will increase (decrease).
Implications for Investors	No approach or model will allow us to identify under or over valued assets.	Reflecting the 50/50 chance of the news being good or bad, there is an equal probability of a price increase and a price decrease.

While the random walk is not magic, there are two prerequisits for it to hold. The first is that investors are rational and form unbiased expectations of the future, based upon all of the information that is available to them at the time. If expectations are set too low or set too high consistently – in other words, investors are too optimistic or pessimistic - information will no longer have an equal chance of being good or bad news, and prices will not follow a random walk. The second is that price changes are caused by new information. If investors can cause prices to change by just trading, even in the absence of information, you can have price changes in the same direction rather than a random walk.

The Basis for Price Patterns

Chartists are not alone in believing that there is information in past prices that can be useful in forecasting future price changes. There are some fundamentalists who use

---

[1] If the probability of good news is greater than the probability of bad news, the price should increase before the news comes out. Technically, it is the expected value of the next information release is zero.

technical and charting indicators, albeit as secondary factors, in picking stocks. They disagree with the fundamental assumptions made by random walkers and argue that

1. Investors are not always rational in the way they set expectations. These irrationalities may lead to expectations being set too low for some assets at some times and too high for other assets at other times. Thus, the next piece of information is more likely to contain good news for the first asset and bad news for the second.
2. Price changes themselves may provide information to markets. Thus, the fact that a stock has gone up strongly the last four days may be viewed as good news by investors, making it more likely that the price will go up today then down.

The debate about whether price changes are random or not has raged for the last 50 years, ever since researchers were able to access price data on stocks. We have to admit that the initial tests were almost all conducted by those who believed that prices follow a random walk, and not surprisingly, they found no price patterns. In the last two decades, there has been an explosion in both the amount of data available and in the points of view of researchers. One of the biggest surprises (at least to those who believed the prevailing dogma of efficient and rational markets) has been the uncovering of numerous price patterns, though it is not clear whether these are evidence of irrational markets and whether they offer potential for profits.

Empirical Evidence

As the studies of the time series properties of prices have proliferated, the evidence can be classified into two classes - studies that focus on short-term (intraday, daily and weekly price movements) price behavior and research that examines long-term (monthly, annual and five-year returns) price movements. Since the findings are contradictory, we will present them separately. We will also present evidence on seasonal patterns in stock prices that seem to persist not only over many periods but also across most markets.

a. Short Term Price Patterns

The notion that today's price change conveys information about tomorrow's price change is deep rooted in most investors' psyches. In its more sophisticated formats, the notion that there are patterns in price movements over short periods of time forms the basis for much of charting. All too often, these patterns are backed up anecdotal evidence, with the successful experiences on one or a few stocks extrapolated to form rules about all stocks and assets. Even in a market that follows a perfect random walk, you will see price patterns on some stocks that seem to defy probability. The entire market may go up ten days in a row, or down, for no other reason than pure chance. Given that this is often true, how do we test to see if there are significant price patterns? We will consider two ways in which researchers have examined this question in this section.

a. Serial correlation

If today is a big up day for a stock, what does this tell us about tomorrow? There are three different points of view. The first is that the momentum from today will carry into tomorrow, and that tomorrow is more likely to be an up day than a down day. The second is that there will be the proverbial profit taking as investors cash in their profits and that the resulting correction will make it more likely that tomorrow will be a down day. The third is that each day we begin anew, with new information and new worries, and that what happened today has no implications for what will happen tomorrow.

Statistically, the serial correlation measures the relationship between price changes in consecutive time periods, whether hourly, daily or weekly, and is a measure of how much the price change in any period depends upon the price change over the previous time period. A serial correlation of zero would therefore imply that price changes in consecutive time periods are uncorrelated with each other, and can thus be viewed as a rejection of the hypothesis that investors can learn about future price changes from past ones. A serial correlation that is positive, and statistically significant, could be viewed as evidence of price momentum in markets, and would suggest that returns in a period are more likely to be positive (negative) if the prior period's returns were positive (negative). A serial correlation which is negative, and statistically significant, could be evidence of price reversals, and would be consistent with a market where positive returns are more likely to follow negative returns and vice versa.

From the viewpoint of investment strategy, serial correlations can sometimes be exploited to earn excess returns. A positive serial correlation would be exploited by a strategy of buying after periods with positive returns and selling after periods with negative returns. A negative serial correlation would suggest a strategy of buying after periods with negative returns and selling after periods with positive returns. Since these strategies generate transactions costs, the correlations have to be large enough to allow investors to generate profits to cover these costs. It is therefore entirely possible that there be serial correlation in returns, without any opportunity to earn excess returns for most investors.

The earliest studies[2] of serial correlation all looked at large U.S. stocks and concluded that the serial correlation in stock prices was small. One of the first by Fama in

[2] Alexander (1964), Cootner (1962) and Fama (1965) all estimated serial correlation in stock prices. Given the difficulty of obtaining data, they worked with small samples over short periods.

1965, for instance, found that 8 of the 30 stocks listed in the Dow had negative serial correlations and that most of the serial correlations were less than 0.05. Other studies confirm these findings – of very low correlation, positive or negative - not only for smaller stocks in the United States, but also for other markets. For instance, Jennergren and Korsvold (1974) report low serial correlations for the Swedish equity market and Cootner (1961) concludes that serial correlations are low in commodity markets as well. While there may be statistical significance associated with some of these correlations, it is unlikely that there is enough correlation in short-period returns to generate excess returns, after you adjust for transactions costs.

The serial correlation in short period returns is affected by market liquidity and the presence of a bid-ask spread. Not all stocks in an index are liquid, and, in some cases, stocks may not trade during a period. When the stock trades in a subsequent period, the resulting price changes can create positive serial correlation. To see why, assume that the market is up strongly on day 1, but that three stocks in the index do not trade on that day. On day 2, if these stocks are traded, they are likely to go up to reflect the increase in the market the previous day. The net result is that you should expect to see positive serial correlation in daily or hourly returns in illiquid market indices. The bid-ask spread creates a bias in the opposite direction, if transactions prices are used to compute returns, since prices have a equal chance of ending up at the bid or the ask price. The bounce that this induces in prices will result in negative serial correlations in returns.[3] For very short return intervals, this bias induced in serial correlations might dominate and create the mistaken view that price changes in consecutive time periods are negatively correlated.

There are some recent studies that find evidence of serial correlation in returns over short time periods, but the correlation is different for high volume and low volume stocks. With high volume stocks, stock prices are more likely to reverse themselves over short periods, i.e., have negative serial correlation. With low volume stocks, stock prices are more likely to continue to move in the same direction – i.e., have positive serial correlation.[4] None of these studies suggest that you can make money of these correlations.

---

[3] Roll (1984) provides a simple measure of this relationship,

$$\text{Bid-Ask Spread} = -\sqrt{2}\ (\text{Serial Covariance in returns})$$

where the serial covariance in returns measures the covariance between return changes in consecutive time periods.

[4] See “Volume and Autocovariances in Short-Horizon Individual Security Returns”, Conrad, Hameed and Niden, in Journal of Finance, 1994.

b. Runs Tests

Once in a while a stock has an extended run where stock prices go up several days in a row or down several days in a row. While this, by itself, is completely compatible with a random walk, you can examine a stock's history to see if these runs happen more frequently or less frequently than they should. A runs test is based upon a count of the number of runs, i.e., sequences of price increases or decreases, in price changes over time. Thus, the following time series of price changes, where U is an increase and D is a decrease would result in the following runs -

UUU DD U DDD UU DD U D UU DD U DD UUU DD UU D UU D

There were 18 runs in this price series of 33 periods. The actual number of runs in the price series is compared against the number that can be expected[5] in a series of this length, assuming that price changes are random. If the actual number of runs is greater than the expected number, there is evidence of negative correlation in price changes. If it is lower, there is evidence of positive correlation. A study of price changes in the Dow 30 stocks, assuming daily, four-day, nine-day and sixteen day return intervals provided the following results -

	DIFFERENCING INTERVAL			
	Daily	Four-day	Nine-day	Sixteen-day
Actual runs	735.1	175.7	74.6	41.6
Expected runs	759.8	175.8	75.3	41.7

The actual number of runs in four-day returns (175.8) is almost exactly what you would expect in a random process. Tthere is slight evidence of positive correlation in daily returns but no evidence of deviations from normality for longer return intervals.

Again, while the evidence is dated, it serves to illustrate the point that long strings of positive and negative changes are, by themselves, insufficient evidence that markets are not random, since such behavior is consistent with price changes following a random walk. It is the recurrence of these strings that can be viewed as evidence against randomness in price behavior.

b. Long Term Price Patterns

While most of the earlier studies of price behavior focused on shorter return intervals, more attention has been paid to price movements over longer periods (six months

[5] There are statistical tables that summarize the expected number of runs, assuming randomness, in a series of any length.

to five-year) in recent years. Here, there is an interesting dichotomy in the results. When long term is defined as months rather than years, there seems to be a tendency towards positive serial correlation. Jegadeesh and Titman present evidence of what they call "price momentum" in stock prices over time periods of up to eight months – stocks that have gone up in the last six months tend to continue to go up whereas stocks that have gone down in the last six months tend to continue to go down. The momentum effect is just as strong in the European markets, though it seems to be weaker in emerging markets.[6] What may cause this momentum? One potential explanation is that mutual funds are more likely to buy past winners and dump past losers, thus generating price continuity.[7]

However, when long term is defined in terms of years, there is substantial negative correlation in returns, suggesting that markets reverse themselves over very long periods. Fama and French examined five-year returns on stocks from 1941 to 1985 and present evidence of this phenomenon. They found that serial correlation is more negative in five-year returns than in one-year returns, and is much more negative for smaller stocks rather than larger stocks. Figure 7.2 summarizes one-year and five-years serial correlation by size class for stocks on the New York Stock Exchange.

---

[6] Rouwenhorst (1998) studied 12 European markets and finds evidence of momentum in each market. In 1999, he presented evidence of momentum in emerging markets. Another paper by Bekaert, Erb, Harvey and Viskanta (1997) finds that momentum investing is not consistently profitable in emerging markets.

[7] Grinblatt, Titman and Wermers (1995) present evidence that is consistent with this explanation.

Figure 7.2: One year and Five year Correlations: Market Value Class: 1941- 1985

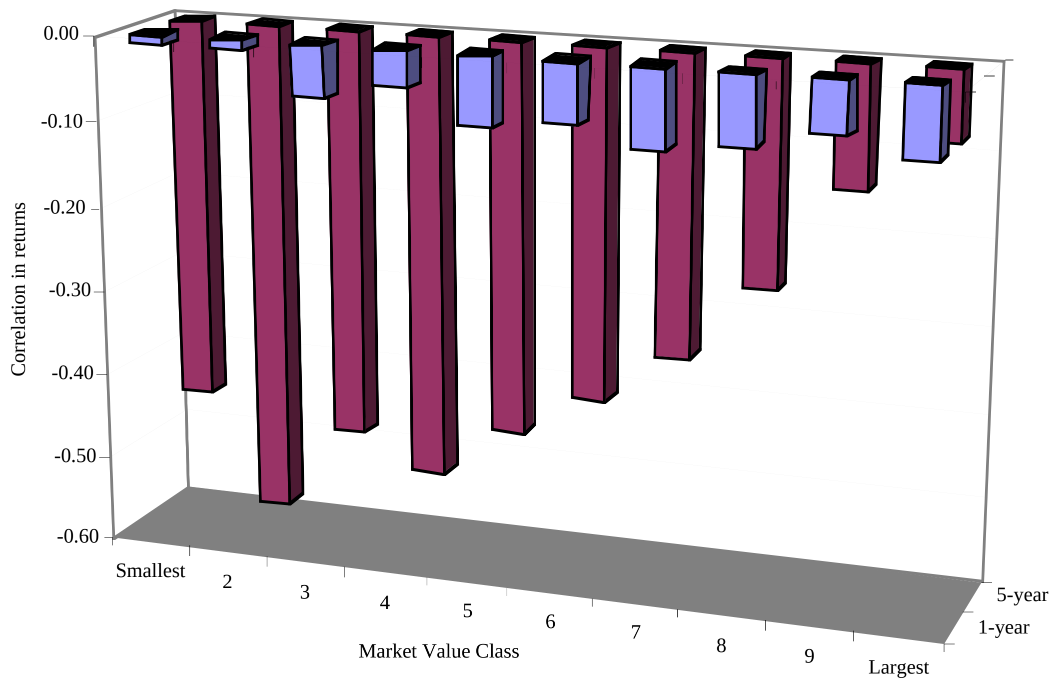

Source: Fama and French (1988)

This phenomenon has also been examined in other markets, and the findings have been similar. There is evidence that returns reverse themselves over long time periods.

Given the findings of little or no correlation in the short term and substantial correlation in the long term, it is interesting that so many technical analysts focus on predicting intraday or daily prices. The bigger payoff seems to be in looking at price patterns over much longer periods, though there are caveats we will present in the next chapter on these long term strategies.

c. Seasonal and Temporal Patterns in Prices

One of the most puzzling phenomena in asset prices is the existence of seasonal and temporal patterns in stock prices that seem to cut across all types of asset markets. As we will see in this section, stock prices seem to go down more on Mondays than on any other day of the week and do better in January than in any other month of the year. What is so surprising about this phenomenon, you might ask? It is very difficult to justify the existence of patterns such as these in a rational market – after all, if investors know that stocks do better in January than in any other month, they should start buying the stock in December and shift the positive returns over the course of the year. Similarly, if investors know that stocks are likely to be marked down on Monday, they are likely to begin marking them down on Friday and hence shift the negative returns over the course of the week.

The January Effect

Studies of returns in the United States and other major financial markets consistently reveal strong differences in return behavior across the months of the year. Figure 7.3 reports average returns by month of the year from 1927 to 2001.

Figure 7.3: Returns by Month of the year - 1927 - 2001

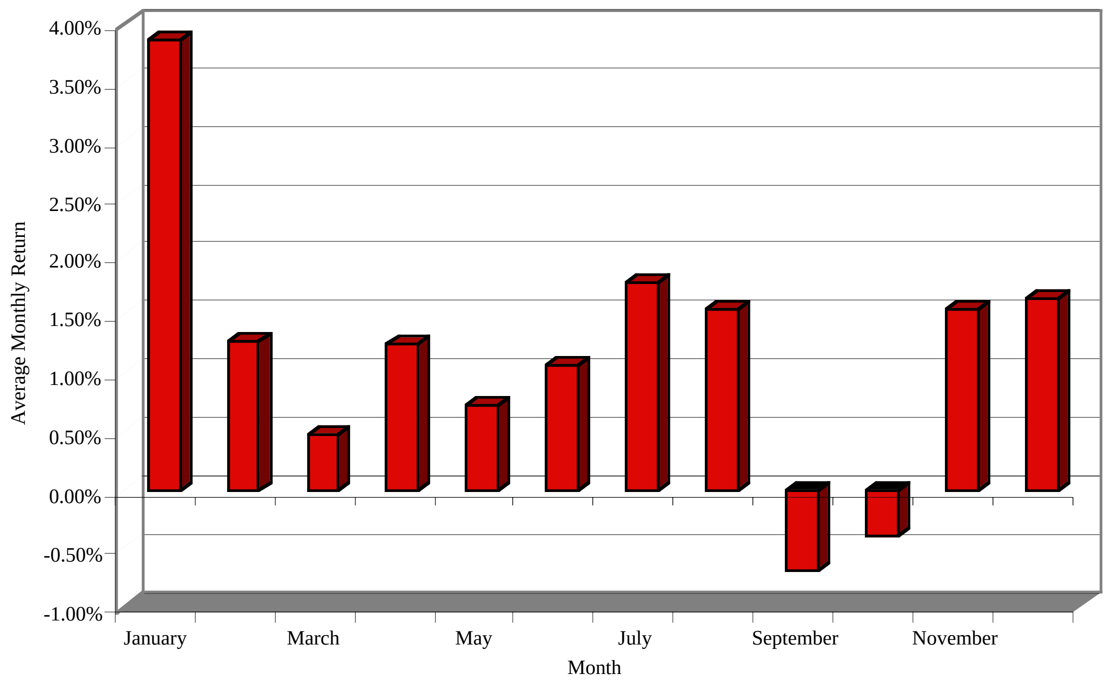

Returns in January are significantly higher than returns in any other month of the year. This phenomenon is called the year-end or January effect, and it can be traced to the first two weeks in January.

The January effect is much more pronounced for small firms than for larger firms, and roughly half of the small firm premium, which is the additional return earned by small firms relative to large firms, is earned in the first few days of January. Figure 7.4 graphs returns in January by size and risk class for data from 1935 to 1986.[8]

[8] This finding is from Haugen, R. and J. Lakonishok, The Incredible January Effect, Dow-Jones Irwin.

Figure 7.4: Returns in January by Size and Risk Class

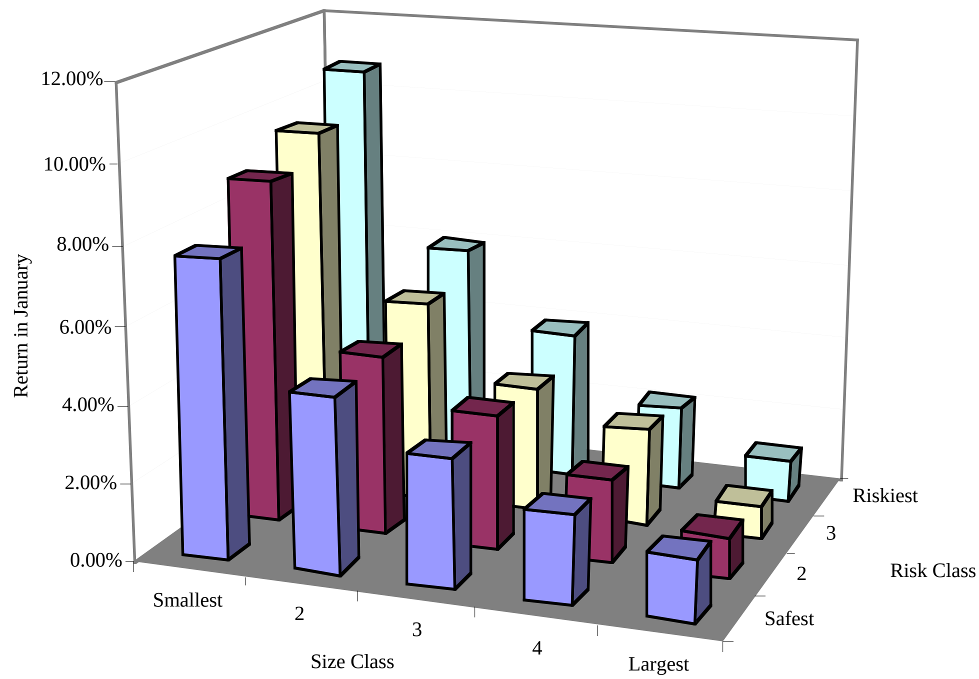

Source: Chopra and Ritter

Note that the January effect is most pronounced for the smallest, riskiest firms in the market and least pronounced for larger, safer firms.

A number of explanations have been advanced for the January effect, but few hold up to serious scrutiny. One is that there is tax loss selling by investors at the end of the year on stocks which have gone down to capture the capital gain, driving prices down, presumably below true value, in December, and a buying back of the same stocks[9] in January, resulting in the high returns. The fact that the January effect is accentuated for stocks that have done worse over the prior year is offered as evidence for this explanation. There are several pieces of evidence that contradict it, though. First, there are countries, like Australia, which have a different tax year, but continue to have a January effect. Second, the January effect is no greater, on average, in years following bad years for the stock market, than in other years.

---

[9] It is to prevent this type of trading that the internal revenue service has a “wash sale rule” that prevent you from selling and buying back the same stock within 45 days. To get around this rule, there has to be some substitution among the stocks. Thus investor 1 sells stock A and investor 2 sells stock B, but when it comes time to buy back the stock, investor 1 buys stock B and investor 2 buys stock A.

A second rationale is that the January effect is related to institutional trading behavior around the turn of the years. It has been noted, for instance, that ratio of buys to sells for institutions drops significantly below average in the days before the turn of the year and picks to above average in the months that follow.[10] It is argued that the absence of institutional buying pushes down prices in the days before the turn of the year and pushes up prices in the days after. Again, while this may be true, it is not clear why other investors do not step in and take advantage of these quirks in institutional behavior.

The universality of the January effect is illustrated in Figure 7.5 where we examine returns in January versus the other months of the year in several major financial markets, and finds strong evidence of a January effect in every market.[11]

Figure 7.5: The International January Effect

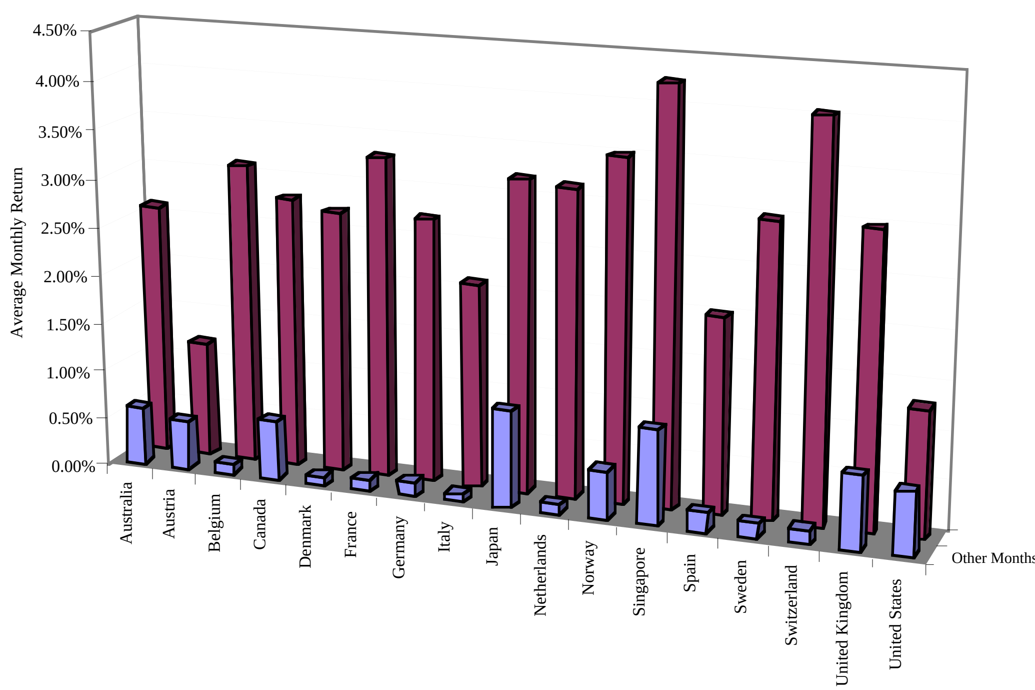

Source: Haugen and Lakonishok

In fact, researchers have unearthed evidence of a January effect in bond and commodity markets as well.

---

[10] Institutional buying drops off in the last 10 days of the calendar year, and picks up again in the first 10 days of the next calendar year.

[11] This is also from Haugen, R. and j. Lakonishok, The Incredible January Effect, Dow Jones Irwin.

The Weekend Effect

Are stock returns consistently higher on some days of the week than others? A surprising feature of stock returns is the existence of what is called the weekend effect, another return phenomenon that has persisted over extraordinary long periods and over a number of international markets. It refers to the differences in returns between Mondays and other days of the week. The significance of the return difference is brought out in Figure 7.6, which graphs returns by days of the week from 1927 to 2001.

Figure 7.6: Returns by Day of the Week - 1927-2001

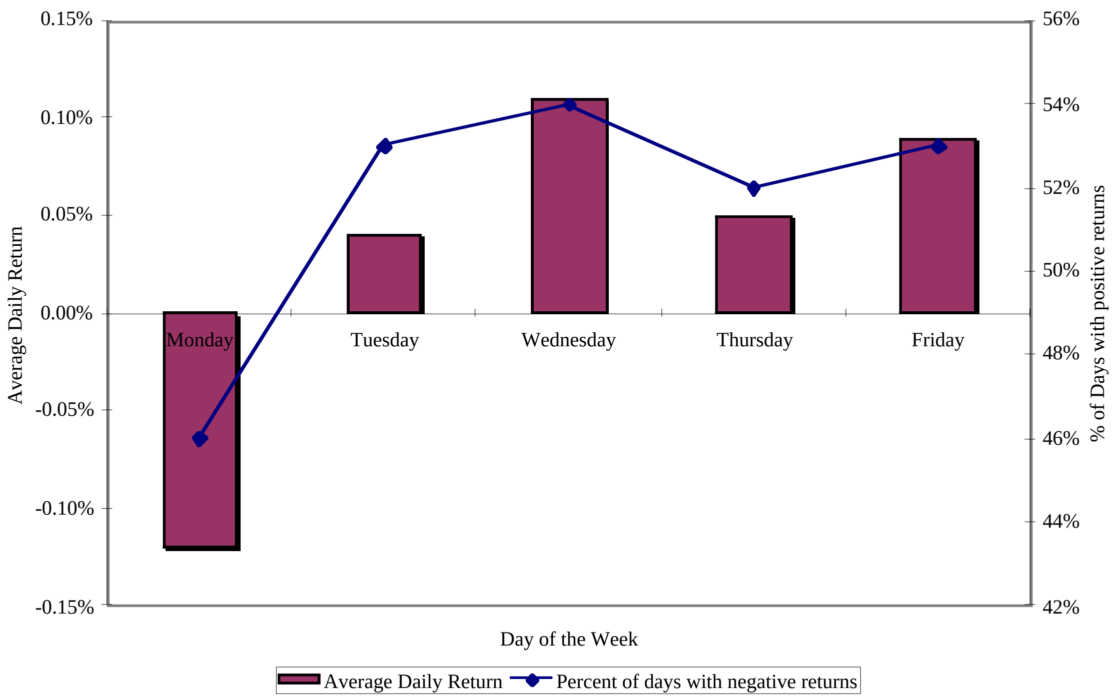

Source: Raw data from CRSP

The returns on Mondays are, on average, negative, whereas the returns on every day of the week are not. In addition, returns on Mondays are negative more often than returns on any other trading day. There are a number of other findings on the Monday effect that researchers have fleshed out.

- The Monday effect is really a weekend effect since the bulk of the negative returns are manifested in the Friday close to Monday open returns. In other words, the negative returns on Monday are generated by the fact that stocks tend to open lower on Mondays than from what happens during the day. The returns from intraday returns on Monday (the price changes from open to close on Monday) are not the culprits in creating the negative returns.

- The Monday effect is worse for small stocks than for larger stocks. This mirrors our findings on the January effect.
- The Monday effect is no worse following three-day weekends than two-day weekends.
- Monday returns are more likely to be negative if the returns on the previous Friday were negative. In fact, Monday returns are, on average, positive following positive Friday returns, and are negative 80% of the time following negative Friday returns.[12]

There are some who have argued that the weekend effect is the result of bad news being revealed after the close of trading on Friday and during the weekend. They point to the fact that more negative earnings reports are revealed after close of trading on Friday. Even if this were a widespread phenomenon, the return behavior would be inconsistent with a rational market, since rational investors would build in the expectation of the bad news over the weekend into the price before the weekend, leading to an elimination of the weekend effect.

The weekend effect is strong in most major international markets, as shown in Figure 7.7.

Figure 7.7: Weekend Effect in International Markets

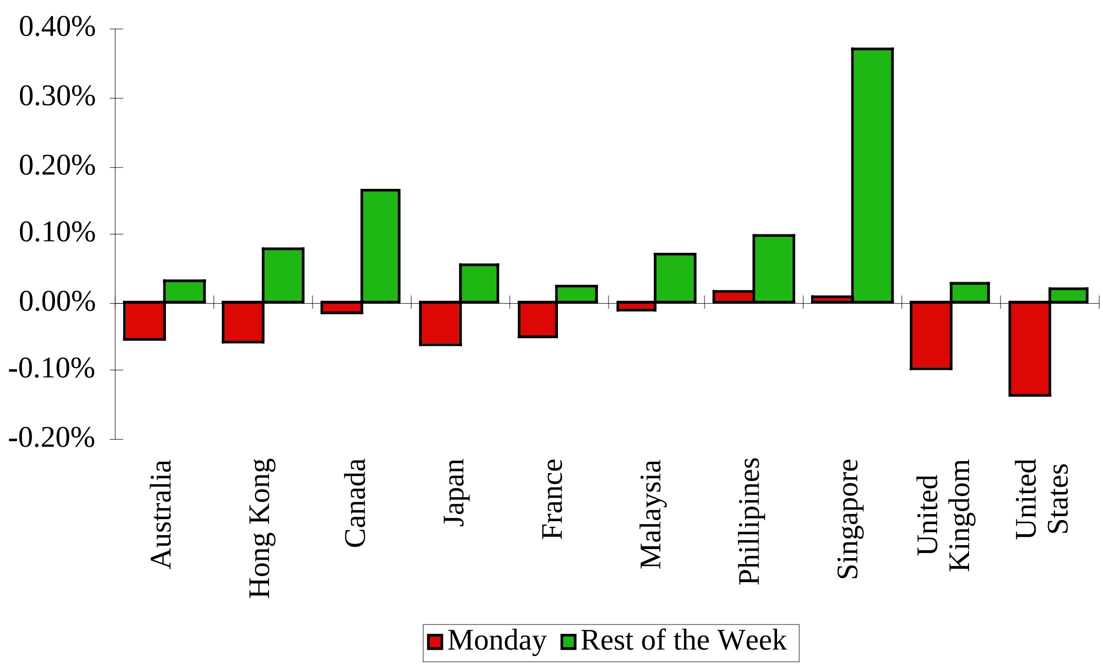

[12] See "The Individual Investor and the Weekend Effect", Abraham and Ikenberry, Journal of Financial and Quantitative Analysis.

The returns on Monday are lower than returns on other days of the week for every international market examined. The presence of a strong weekend effect in Japan, which allowed Saturday trading for a portion of the period studies here indicates that there might be a more direct reason for negative returns on Mondays than bad information over the weekend.

As a final note, the negative returns on Mondays cannot be just attributed to the absence of trading over the weekend. The returns on days following trading holidays, in general, are characterized by abnormally positive, not negative, returns. Figure 7.8 summarizes returns on trading days following major holidays and confirms this pattern.

Figure 7.8: Return on first trading day after

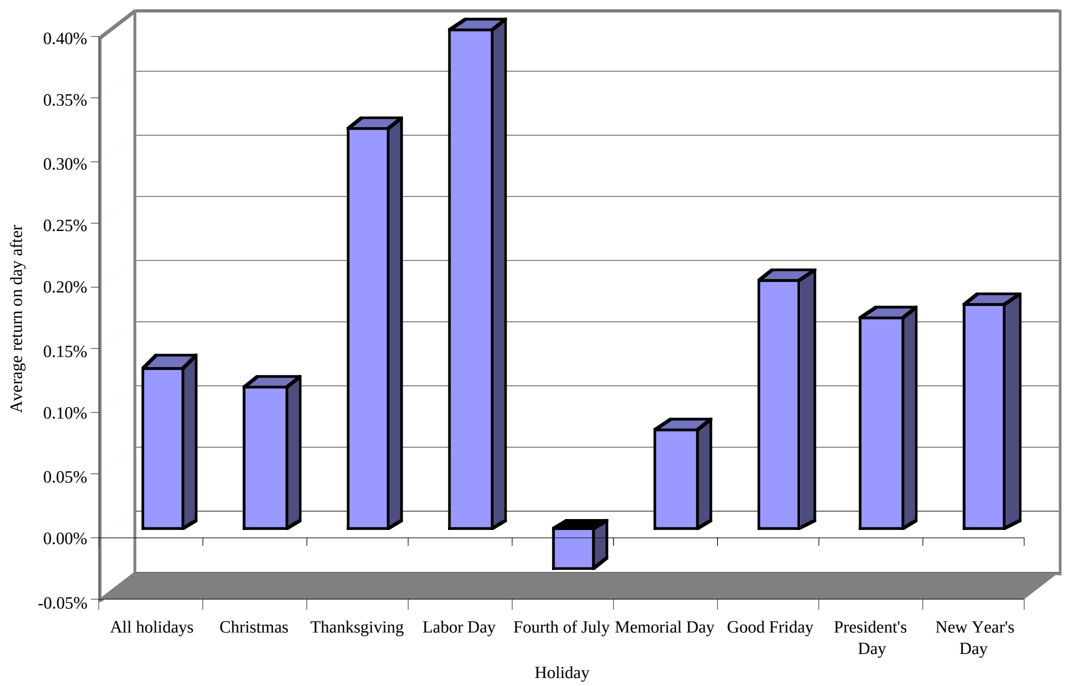

In fact, the returns on the first trading day after a holiday tend to be much more positive than returns on other trading days.[13]

d. Volume Patters

While the random walk hypothesis is silent about the relationship between trading volume and prices, it does assume that all available information is incorporated in the current price. Since trading volume is part of publicly available information, there should therefore be no information value to knowing how many shares were traded yesterday or the day before.

---

[13] See “High Stock Returns before holidays” R.A. Ariel.

As with prices, there is evidence that trading volume carries information about future stock price changes. In a study in 1998, Datar, Naik and Radcliffe show that low volume stocks earn higher returns than high volume stocks, though they attribute the differential return to a liquidity premium on the former. A more surprising result comes from Lee and Swaminathan (1998) who look at the interrelationship between price and trading volume. In particular, they examine the price momentum effect that was documented by Jegadeesh and Titman – that stocks that go up are more likely to keep going up and stocks that go down are more likely to keep dropping in the months after - and show that it is much more pronounced for high volume stocks. Figure 7.9 classifies stocks based upon how well or badly they have done in the last six months (winners, average and loser stocks) and their trading volume (low, average and high) and looks at returns on these stocks in the following six months.

Figure 7.9: Volume and Price Intreaction- NYSE and AMEX stocks - 1965-95

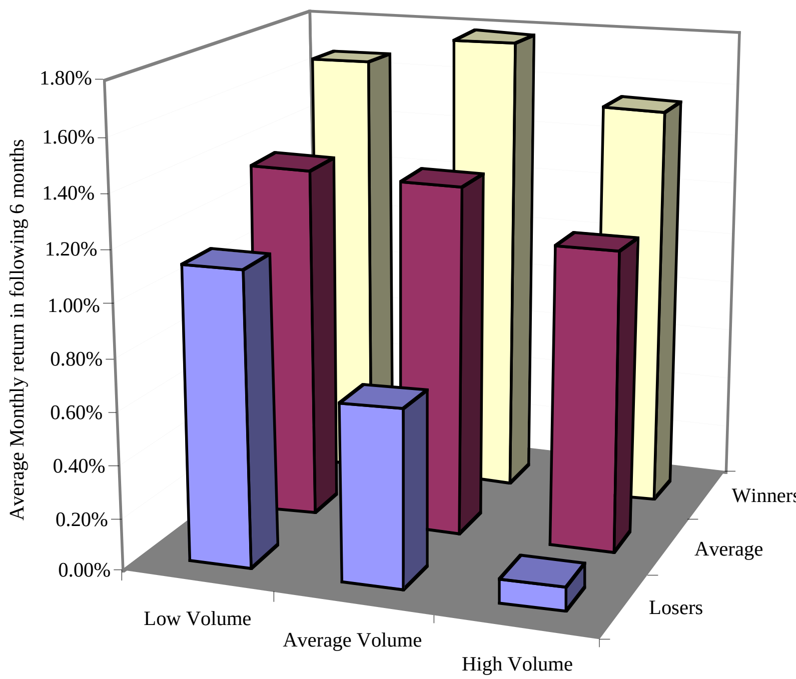

Source: Lee and Swaminathan

Note that the price momentum effect is strongest for stocks with high trading volume. In other words, a price increase or decrease that is accompanied by strong volume is more likely to continue into the next period. Stickel and Verecchia confirm this result with shorter period returns – they conclude that increases in stock prices that are accompanied by high trading volume are more likely to carry over into the next trading day.

In summary, the level of trading volume in a stock, changes in volume and volume accompanied by price changes all seem to provide information that investors can use to pick stocks. It is not surprising that trading volume is an integral part of technical analysis.

Data Mining or Anomalies

When looking at the evidence on seasonal and temporal anomalies in stock price data, we are faced with an interesting dilemma. As stock price data has become both richer (we have gone from annual to intraday data and from just equity markets to bond and derivatives markets) and easier to access and use, it is not surprising that the number of inefficiencies and anomalies discovered have also increased. You could argue that some of these findings can be attributed to the sheer volume of data that is available to us. As hundreds of researchers pore over this data, using finer and finer microscopes, they will find patterns depending upon the portion of the data that they are looking at. In a spirited defense of efficient markets, Fama presents the argument that almost of the anomalies and inefficiencies that researchers have detected over the last 40 years can be attributed purely to chance, rather than irrational or inefficient investors. In fact, he makes the interesting point that those researchers who claim to find inefficiencies cannot seem to agree on whether the inefficiencies indicate a market that over reacts or one that under reacts to new information.[15]

Investor Irrationality

Historians who have examined the behavior of financial markets over time have challenged the assumption of rationality that underlies much of efficient market theory. They point to the frequency with speculative bubbles have formed in financial markers, as investors buy into fads or get-rich-quick schemes, and the crashes with which these bubbles have ended, and suggest that there is nothing to prevent the recurrence of this phenomenon in today's financial markets. In fact, the evidence on price patterns, in the short and long term, in different calendar months and on different weekdays suggests that there is much about markets that we cannot explain with a rational investor model. In this section, we will begin by considering some of the evidence accumulated by psychologists on human behavior and then consider financial market phenomena that seem more consistent with an irrational market than a rational one.

---

[15] See “Market Efficiency, Long Term Returns and Behavioral Finance” by E.F. Fama, Journal of Financial Economics, v 49, 0g 283-306.

Psychological Studies

At the risk of stating the obvious, investors are human and it is not surprising that financial markets reflect human frailties. In an extraordinary book (at least for an academic economist), Robert Schiller presented some of the evidence accumulated of human behavior by psychologists that may help us understand financial market behavior. He categorizes these findings into several areas and we will consider each below.

The Need for Anchors

When confronted with decisions, it is human nature to begin with the familiar and use it to make judgments. Kahnemann and Tversky, whose research has helped illuminate much of what is called behavioral finance, ran an experiment where they used a wheel of fortune with numbers from 1 to 100 to illustrate this point. With a group of subjects, they spun the wheel to get a number and then asked the subjects numerical questions about obscure percentages – the percent of the ancient Egyptians who ate meat, for instance. The subjects would have to guess whether the right answer was higher or lower than the number on the wheel and then provide an estimate of the actual number. They found that the answer given by subjects was consistently influenced by the outcome of the wheel spin. Thus, if the number on the wheel was 10, the answer was more likely to be 15 or 20%, whereas if the number on the wheel was 60%, it was more likely to be 45 or 50. Shiller argues that market prices provide a similar anchor with publicly traded assets. Thus, an investor asked to estimate the value of a share is likely to be influenced by the market price, with the value increasing as the market price rises.

The Power of the Story

For better or worse, human actions tend to be based not on quantitative factors but on story telling. People tend to look for simple reasons for their decisions, and will often base their decision on whether these reasons exist. In a study of this phenomenon, Shafir, Simonson and Tversky gave subjects a choice on which parent they would choose for sole custody of a child. One parent was described as average in every aspect of behavior and standing whereas the other was described more completely with both positive (very close relationship with child, above-average income) and negative characteristics (health problems, travels a lot). Of the subjects studied, 64% picked the second. Another group of subjects was given the same choice but asked which one they would deny custody to. That group also picked the second parent. While the results seem inconsistent – the first group chose the second parent as the custodian and the second group rejected the same parent, given the same facts – they suggest that investors are more comfortable with investment decisions that can be justified with a strong story than one without.

Overconfidence and Intuitive Thinking

As you have undoubtedly become aware from your interactions with friends, relatives and even strangers over time, human beings tend to be opinionated about things they are not well informed on and to make decisions based upon these opinions. In an illustrative study, Fischhhof, Slovic and Lichtenstein asked people factual questions, and found that people gave an answer and consistently overestimated the probability that they are right. In fact, they were right only about 80% of the time that they thought they were. What are the sources of this overconfidence? One might just be evolutionary. The confidence, often in the face of poor odds, may have been what allowed us to survive and dominate as a species. The other may be more psychological. Human beings seem to have a propensity to hindsight bias, i.e., they observe what happens and act as it they knew it was coming all the time. Thus, you have investors that claim to have seen the crash in dot.com companies in the late 1990s coming during earlier years, thought nothing in their behavior suggests that they did.

Herd Behavior

The tendency of human beings to be swayed by crowds has been long documented and used by tyrants over time to impose their will on us. In a fascinating experiment, Asch illustrated this by putting a subject into a group of people, asking them a question to which the answer was obvious and then inducing other people in the group to provide the wrong answer deliberately. Asch noted that the subject changed his answer one-third of the time to reflect the incorrect answer given in the group. While Asch attributed this to peer pressure, subsequent studies found the same phenomenon even when the subject could not see or interact with others in the group. This would suggest that the desire to be part of the crowd is due to more than peer pressure.

While there is a tendency to describe herd behavior as irrational, it is worth noting that you can have the same phenomena occur in perfectly rational markets through a process called information cascade. Schiller provides an example with two restaurants, where people come into town one after another. Assume that the first person to come in picks the first restaurant and assume that the choice is random. The second person who comes into town will observe the first person sitting in the first restaurant, and is more likely to pick the same restaurant. As the number of subjects entering the market increases, you are likely to see the crowd at the first restaurant pick up, while business at the second restaurant will be minimal. Thus, a random choice by the first customer in the market creates enough momentum to make it the dominant restaurant. All to often, in investing, investors at early stages in the process (initial public offering) pile into specific initial public offerings and push their prices up. Other initial public offerings are ignored and languish at low prices. It is entirely possible that the first group of stocks will be overvalued, while the latter are undervalued. Since herd behavior is made worse by rumors by the spreading of rumors, you could argue that the coming together of the available data and media sites such as CNBC and MSNBC has made it more possible for herd behavior to spend and not less.

Unwillingness to admit mistakes

It may be human to err, but it is also human to claim not to err. In other words, we are much more willing to claim our successes than we are willing to face up to our failures. Kahneman and Tversky, in their experiments on human behavior, noticed that subjects when presented with choices relative to the status quo often made choices based upon unrealistic expectations. They noted that a person who has not made peace with his losses is likely to accept gambles that would otherwise be unacceptable to him. Anyone who has visited a casino will attest to this finding.

In investing, Shefrin and Statman call this the disposition effect, i.e, the tendency to hold on to losers too long and to sell winners too soon.[16] They argue that it is widespread and can cause systematic mispricing of some stocks. Terrance Odean used the trading records of over 10000 customers at a discount brokerage house to examine whether there is evidence of this behavior among investors.[17] He notes that investors realized only 9.8% of their losses each year, whereas they realize 14.8% of their gains.[18] He also finds that investors seem to sell winners too soon, since winning stocks that get sold continue to go up for months after the sale. Overall, he argues that there is evidence of the disposition effect among investors.

Empirical Evidence

While it is evident that human beings do not always behave rationally, it does not necessarily follow that markets will also be irrational. In fact, you could argue (as some believers in market efficiency do) that markets can be efficient even with irrational investors for several reasons. First, it is possible that there is a selection process that occurs in markets where irrational investors lose consistently to rational investors and eventually get pushed out of the market. Second, it is also possible that irrationalities cut in both directions

---

[16] Shefrin, H. and M. Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and Evidence, Journal of Finance, v40, p777-790.

[17] Odean, T., 1997,Are investors reluctant to realize their losses?, Working paper, University of California, Davis.

[18] The only month in which more losses are realized than gains is December.

– some leading investors to buy when they should not and others leading them to sell when they should not; if these actions offset each other, you could still have a market price that is unaffected by rational investors. The only way to resolve this debate is to look at the empirical evidence on the presence or absence of irrationality in market behavior. In this section, we will begin by looking at experimental studies that claim to document irrational investors, and then consider the evidence accumulated through the centuries on bubbles and whether their existence alone indicates irrational investors. In fact, this is a discussion we will revisit in the chapters to come, since investment philosophies are based upon specific and often contradictory views of human irrationality.

Experimental Studies

One of the problems we face when we test for irrationality in financial markets is the number of variables that cannot be controlled for. Investors enter and leave markets, new information arrives constantly and the macroeconomic environment changes frequently, making it impossible to construct a controlled experiment. A few researchers have attempted to get around this problem by constructing experimental studies, similar to those used by psychologists and sociologists in the previous section, to examine how investors behave in financial markets.

One such study was done at the University of Arizona. In this study, groups of students were chosen as subjects and asked to play the role of traders in a single asset for 15 trading days. They were told at the start of the experiment that a payout would be declared on this asset after each trading day, and that it would take one of four values- 0, 8, 28 or 60 cents – with equal probability. Consider how a rational investor would value this very simple asset. Since the average payout is 24 cents, the asset's expected value on the first trading day of a fifteen day experiment should be $3.60 (24*15), the second day should be $3.36 and so on. The traders were allowed to trade each day and the entire experiment was repeated 60 times. The resulting market prices each day, averaged across all 60 experiments, are reported in figure 7.10 and contrasted with the expected values to a rational investor.

Figure 7.10: Prices from Behavioral Experiment

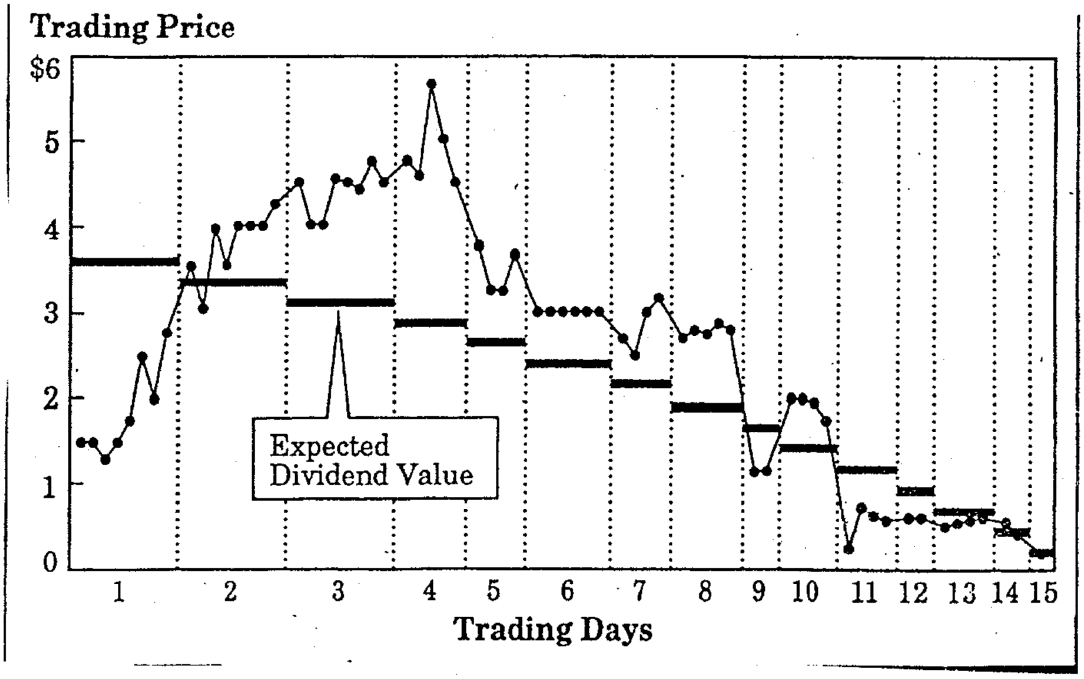

There is clear evidence here of a 'speculative bubble' forming during periods 3 to 5, where prices exceed expected values significantly, The bubble ultimately bursts, and prices approach the expected value by the end of the 15^{th} period.[19] Furthermore, when price curbs of 15 cents were introduced, the booms lasted even longer because traders knew that prices would not fall by more than 15 cents in a period. Thus, the notion that price limits can control speculative bubbles seems misguided. Does this experiment conclusively prove that investors are irrational? Of course not. It is worth noting, though, that if bubbles are feasible in as simple a market as this one, where every investor obtains the same information, it is clearly feasible in real financial markets, where there is much more differential information and much greater uncertainty about expected value.

In fairness, it should be noted that the evidence from other experimental studies is largely supportive of rationality. Investors do seem to make reasonable judgments based upon the information they have, and markets do a good job of aggregating this information in the market price.

---

[19] Some of the experiments were run with students, and some with Tucson businessmen, with 'real world' experience. The results were similar for both groups.

Market Bubbles

Proponents of market irrationality have pointed to market bubbles as a primary exhibit in their case against efficient markets. Through the centuries, markets have boomed and busted, and in the aftermath of every bust, irrational investors have been blamed for the crash. As we will see in this section, it is not that simple. You can have bubbles in markets with only rational investors, and assessing whether a bubble is due to irrational investors is significantly more difficult than it looks from the outside.

A Short History of Bubbles

As long as there have been markets, there have been bubbles. Two of the earliest bubbles to be chronicled occurred in the 1600s in Europe. One was the amazing boom in prices of tulip bulbs in Holland that began in 1634. A single Tulip bulb (Semper Augustus was one variety) sold for more than 5000 guilders (the equivalent of more than $ 60,000 today) at the peak of the market. Stories abound, though many of them may have been concocted after the fact, of investors selling their houses and investing the money in tulip bulbs. As new investors entered the market in 1636, the frenzy pushed up bulb prices even more until the price peaked in early February. Figure 7.11 presents the price of one type of bulb (Switzers) in January and February of 1637.[20]

[20] This graph is based upon data provided by Garber(1990) in “Crashes and Panics: The lessons of History”, Dow Jones Irwin. It should be pointed out that he does not believe that the pricing of tulip bulbs was irrational for much of the period.

Figure 7.11: Price of a Tulip Bulb (Switser) - January - February 1637

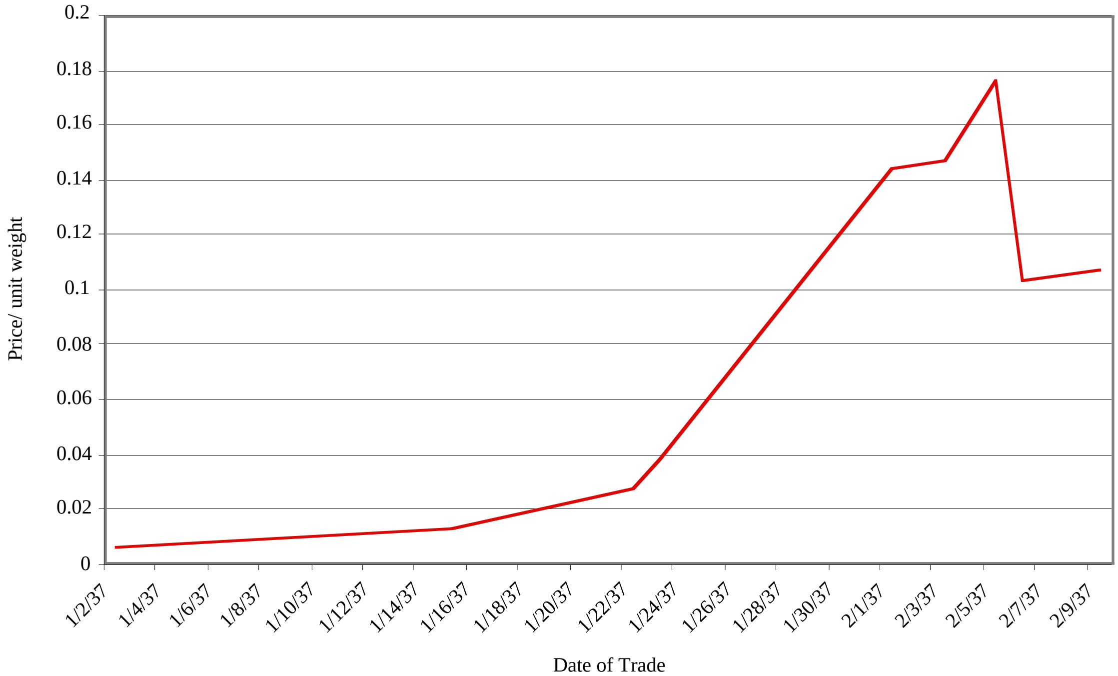

Source: Raw data from Garber

Note that the price peaked on February 5, 1637, but an investor who bought tulip bulbs at the beginning of the year would have seen his or her investment increase almost 30 fold over the next few weeks.

A little later in England, a far more conventional bubble was created in securities of a firm called the South Seas Corporation, a firm with no assets that claimed to have the license to mint untold riches in the South Seas. The stock price was bid up over the years before the price plummeted. The crash, which is described in vivid detail in Charles Mackay's classic book titled "Extraordinary Delusions and the Popular Madness of Crowds", left many investors in England poorer.[21]

Through the 1800s, there were several episodes of boom and bust in the financial markets in the United States and many of these were accompanied by banking panics.[22] As markets became broader and more liquid in the 1900s, there was a renewed hope that

---

[21] To get a flavor of financial markets in England at the time of the South Sea bubble, you should look at "A Conspiracy of Paper", a novel set in the era by David Liss. Edward Chancellor's "Devil takes the hindmost" provides historical perspective on the bubble.

[22] The crash of 1873 was precipitated by the failure of firm called Jay Cooke, a financial-service firm in Philadelphia. The New York Stock Exchange was closed for ten days and several banks closed their doors in the aftermath.

liquidity and more savvy investors would make bubbles a phenomenon of the past, but it was not to be. In 1907, J.P. Morgan had to intervene in financial markets to prevent panic selling, a feat that made his reputation as the financier of the world. The 1920s saw a sustained boom in U.S. equities and this boom was fed by a number of intermediaries ranging from stockbrokers to commercial banks and sustained by lax regulation. The crash of 1929 precipitated the great depression, and created perhaps the largest raft of regulatory changes in the United States, ranging from restrictions on banks (the Glass-Steagall Act) to the creation of a Securities Exchange Commission.

The period after the second world war ushered in a long period of stability for the United States, and while there was an extended period of stock market malaise in the 1970s, the bubbles in asset prices tended to be tame relative to past crashes. In emerging markets, though, bubbles continued to form and burst. In the late 1970s, speculation and attempts by some in the United States to corner the precious metals markets did create a brief boom and bust in gold and silver prices. By the mid-1980s, there were some investors who were willing to consign market bubbles to history. On October 19, 1987, the U.S. equities market lost more than 20% of its' value in one day, the worst single day in market history, suggesting that investors, notwithstanding technological improvements and more liquidity, still shared a great deal with their counterparts in the 1600s. In the 1990s, we witnessed the latest in this cycle of market bubbles in the dramatic rise and fall of the "dot-com" sector. New technology companies with limited revenues and large operating losses went public at staggering prices (given their fundamentals) and kept increasing. After peaking with a market value of $ 1.4 trillion in early 2000, this market too ran out of steam and lost almost all of this value in the subsequent year or two. Figure 7.12 summarizes the Internet index and the NASDAQ from 1994 to 2001:

Figure 7.12: The Tech Boom

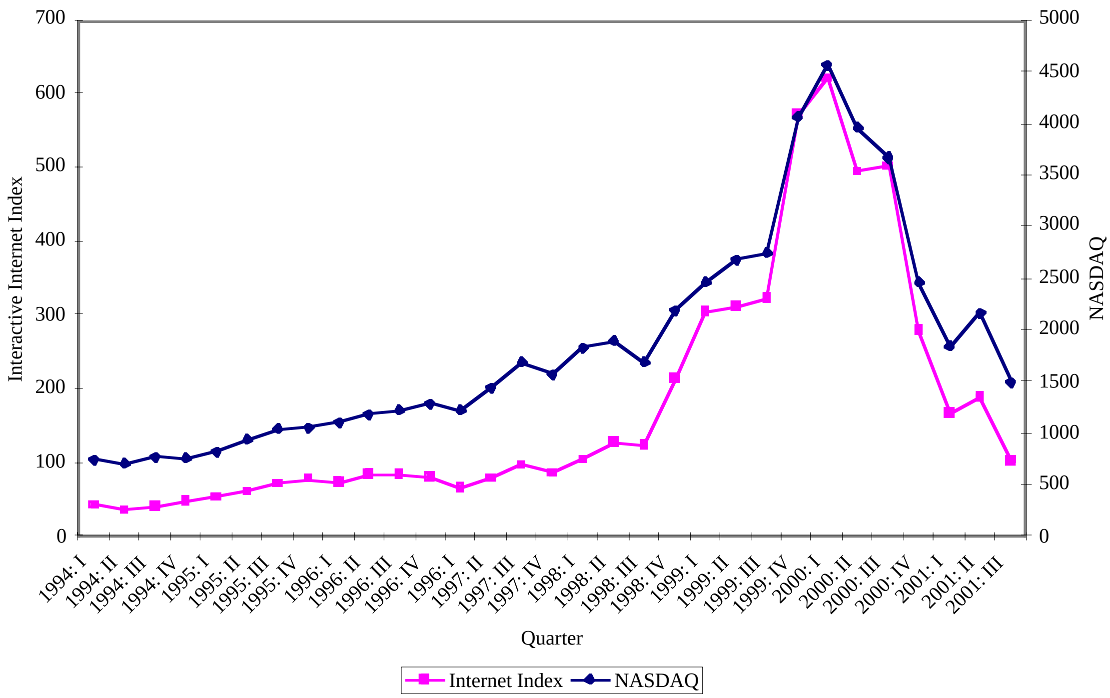

Source: Raw data from Bloomberg

The chart again has the makings of a bubble, as the value of the index internet index increased almost ten fold over the period, dragging the tech-heavy NASDAQ up with it.

Rational Bubbles?

A rational bubble sounds like an oxymoron, but it is well within the realms of possibility. Perhaps the simplest way to think of a rational bubble is to consider a series of coin tosses, with a head indicating a plus day and a tail a minus day. You would conceivably get a series of plus days pushing the stock price above the fair value, and the eventual correction is nothing more than a reversion back to a reasonable value. Note too that it is difficult to tell a bubble from a blunder. Investors in making their assessments for the future can make mistakes in pricing individual assets, either because they have poor information or because the actual outcomes (in terms of growth and returns) do not match expected values. If this is the case, you would expect to see a surge in prices followed by an adjustment to a fair value. In fact, consider what happened to gold prices in the late 1970s. As inflation increased, many investors assumed (incorrectly in hindsight) that high inflation was here to stay and pushed up gold prices accordingly. Figure 7.13, which graphs gold prices from 1970 to 1986, looks very much like a classic bubble, but may just indicate our tendencies to look at things in the rear view mirror, after they happen.

Figure 7.13: Gold Prices: 1970-86

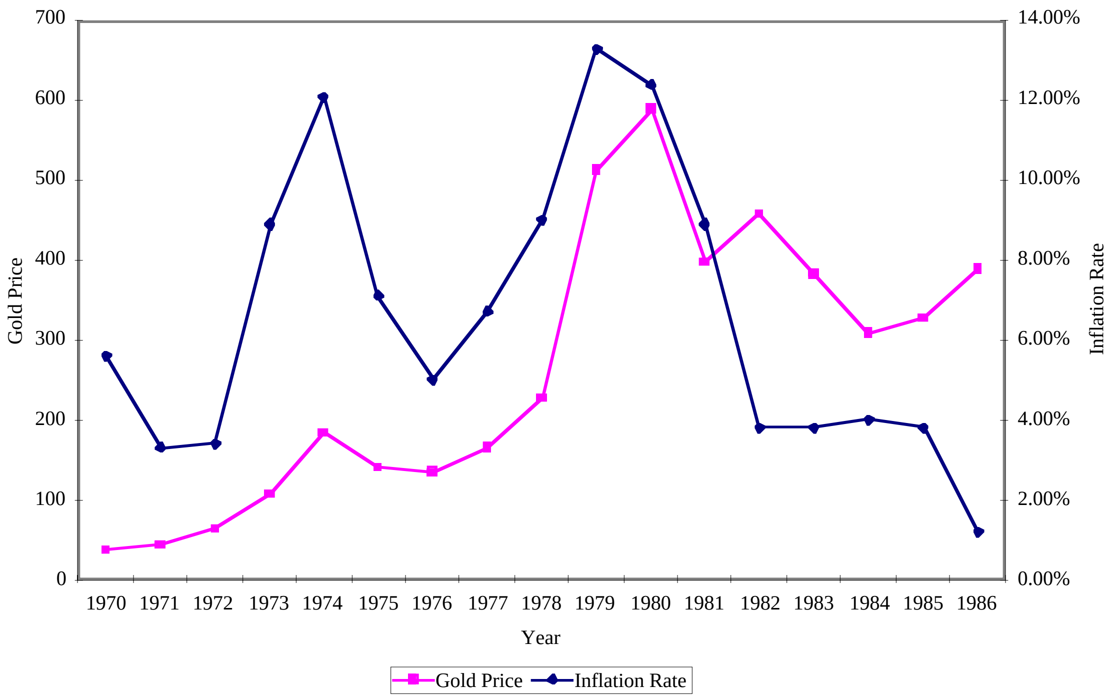

Source: Raw data from Bloomberg

Note that the surge in gold prices closely followed the increase in inflation in the late 1970s, reflecting its value as a hedge against inflation. As inflation declined in the 1980s, gold prices followed. It is an open question, therefore, whether this should be even considered a bubble.

Bubble or Blunder: Tests

There are some researchers who argue that you can separate bubbles from blunders by looking at how prices build up over time. Santoni and Dwyer (1990), for instance, argue that you need two elements for a bubble – positive serial correlation in returns and a delinking of prices and fundamentals as the bubble forms. They test the periods prior to 1929 and 1987 crashes to examine whether there is evidence of bubbles forming in those periods. Based upon their analysis, there is no evidence of positive serial correlation in returns or of a reduction in the correlation between prices and fundamentals (which they define as dividends) in either period. Therefore, they argue that neither period can be used as an example of a bubble.

While there is truth to the underlying premise, these tests may be too weak to capture bubbles that form over long periods. For instance, Santoni and Dwyer's conclusion of no serial correlation seems to be sensitive to both the time periods examined and the return interval used. In addition, detecting a delinking of prices and fundamentals statistically may be difficult to do if it happens gradually over time. In short, these may be useful indicators but they are not conclusive.

Bubbles: From Inception to Crash

One or the more fascinating questions in economics examines how and why bubbles form and what precipitates their bursting. While each bubble has its own characteristics, there seem to four phases to every bubble.

Phase 1: The Birth of the Bubble

Most bubbles have their genesis in a kernel of truth. In other words, at the heart of most bubbles is a perfectly sensible story. Consider, for instance, the dot.com bubble. At its center was a reasonable argument that as more and more individuals and businesses gained online access, they would also be buying more goods and services online. The bubble builds as the market provides positive reinforcement to some investors and businesses for irrational or ill-thought out actions. Using the dot.com phenomenon again, you could point to the numerous start-up companies with half-baked ideas for e-commerce that were able to go public with untenable market capitalizations and the investors who made profits along the way.

A critical component of bubbles building is the propagation of the news of the success to other investors in the market, who on hearing the news, also try to partake in the bubble. In the process, they push prices up and provide even more success stories that can be used to attract more investors, thus providing the basis for a self-fulfilling prophecy. In the days of the tulip bulb craze, this would have had to be word of mouth, as successful investors spread the word, with the success being exaggerated in each retelling of the story. Even in this century, until very recently, the news of the success would have reached investors through newspapers, financial newsmagazines and the occasional business show on television. In the dot.com bubble, we saw two additional phenomena that allowed news and rumors to spread even more quickly. The first was the internet itself, where chat rooms and web sites allowed investors to tell their success stories (or make them up as they went along). The second was the creation of cable stations such as CNBC, where analysts and money managers could present their views to millions of investors.

Phase 2: The Sustenance of the Bubble

Once a bubble forms, it needs sustenance. Part of the sustenance is provided by the institutional parasites that make money of the bubble and develop vested interests in preserving and expanding the bubbles. Among these parasites, you could include:

- Investment banks: Bubbles in financial markets bring with them a number of benefits to investment banks, starting with a surge in initial public offerings of firms but expanding to include further security issues and restructurings on the part of established firms that do not want to be shut out of the party.
- Brokers and analysts: A bubble generates opportunities for brokers and analysts selling assets related to the bubble. In fact, the ease with which investors make money as asset prices go up, often with no substantial reason, relegates analysis to the backburner.
- Portfolio Managers: As a bubble forms, portfolio managers initially watch in disdain as investors they view as naïve push up asset prices. At some point, though,, even the most prudent of portfolio managers seem to get caught up in the craze and partake of the bubble, partly out of greed and partly out of fear.
- Media: Bubbles make for exciting business news and avid investors. While this is especially noticeable in the dot.com bubble, with new books, television shows and magazines directly aimed at investors in these stocks, even the earliest bubbles had their own versions of CNBC.

In addition to the institutional support that is provided for bubbles to grow, intellectual support is usually also forthcoming. There are both academics and practitioners who argue, when confronted with evidence of over pricing, that the old rules no longer apply. New paradigms are presented justifying the high prices, and those who disagree are disparaged as old fashioned and out of step with reality.

Phase 3: The Bursting of the Bubble

All bubbles eventually burst, though there seems to be no single precipitating event that causes the reassessment. Instead, there is a confluence of factors that seem to lead to the price implosion. The first is that bubbles need ever more new investors (or at least new investment money) flowing in for sustenance. At some point, you run out of suckers as the investors who are the best targets for the sales pitch become fully invested. The second is that each new entrant into the bubble is more outrageous than the previous one. Consider, for instance, the dot.com bubble. While the initial entrants like America Online and even Amazon.com might have had a possibility of reaching their stated goals, the new dot.com companies that were listed in the late 1990s were often idea companies with no vision of how to generate commercial success. As these new firms flood the market, even those who are apologists for high prices find themselves exhausted trying to explain the unexplainable.

The first hint of doubt among the true believers turns quickly to panic as reality sets in. Well devised exit strategies break down as everyone heads for the exit doors at the same time. The same forces that created the bubble cause its demise and the speed and magnitude of the crash mirror the formation of the bubble in the first place.

Phase 4: The Aftermath

In the aftermath of the bursting of the bubble, you initially find investors in complete denial. In fact, one of the amazing features of post-bubble markets is the difficulty of finding investors who lost money in the bubble. Investors either claim that they were one of the prudent ones who never invested in the bubble in the first place or that they were one of the smart ones who saw the correction coming and got out in time.

As time passes and the investment losses from the bursting of the bubble become too large to ignore, the search for scapegoats begins. Investors point fingers at brokers, investment banks and the intellectuals who nurtured the bubble, arguing that they were mislead.

Finally, investors draw lessons that they swear they will adhere to from this point on. “I will never invest in a tulip bulb again” or “I will never invest in a dot.com company again” becomes the refrain you hear. Given these resolutions, you may wonder why price bubbles show up over and over. The reason is simple. No two bubbles look alike. Thus, investors, wary about repeating past mistakes, make new ones, which in turn create new bubbles in new asset classes.

Upside versus Downside bubbles

Note that most investors think of bubbles in terms of asset prices rising well above fair value and then crashing. In fact, all of the bubbles we have referenced from the tulip bulb craze to the dot-com phenomenon were upside bubbles. But can asset prices fall well below fair market value and keep falling? In other words, can you have bubbles on the downside? In theory, there is no reason why you could not, and this makes the absence of downside bubbles, at least in the popular literature, surprising. One reason may be that investors are more likely to blame external forces – the bubble, for instance – for the money they lose when they buy assets at the peak of an upside bubble and more likely to claim the returns they make when they buy stocks when they are at the bottom of a downside bubble as evidence of their investment prowess.

Another may be that it is far easier to create investment strategies to take advantage of under priced assets (in a downside bubble) than it is to take advantage of over priced assets. With the former, you can always buy the asset and hold until the market rebounds. With the latter, your choices are both more limited and more likely to be time limited. You can borrow the asset and sell it (short the asset), but not for as long as you want – most short selling is for a few months. If there are options traded on the asset, you may be able to buy puts on the asset though, until recently, only of a few months duration. In fact, there is a regulatory bias in most markets against such investors who are often likely to be categorized as speculators. As a consequence of these restrictions on betting against overpriced assets, bubbles on the upside are more likely to persist and become bigger over time, whereas bargain hunters operate as a floor for downside bubbles.

A Closing Assessment

Based upon our reading of history, it seems reasonable to conclude that there are bubbles in asset prices, though only some of them can be attributed to market irrationality. Whether investors can take advantage of bubbles to make money seems to be a more difficult question to answer. Part of the reason for the failure to exploit bubbles seems to stem from greed; even investors who believe that assets are over priced want to make money off the bubble due to the difficulty of determining when a bubble will burst. Over valued assets may get even more over valued and these overvaluations can stretch over years, thus imperiling the financial well being of any investor who has bet against the bubble. There is also an institutional interest on the part of investment banks, the media and portfolio managers, all of whom feed of the bubble, to perpetuate the bubble.

The Foundations of Technical Analysis

It is best to let technical analysts provide the basis for their approach in their own words. Magee in his classic book on technical analysis made the following argument: "*It is futile to assign an intrinsic value to a stock certificate. One share of US Steel, for example, was worth $261 in the early fall of 1929, but you could buy it for only $22 in June 1932. By March 1937 it was selling for $126 and just one year later for $38. ... This sort of thing, this wide divergence between presumed value and intrinsic value, is not the exception; it is the rule; it is going on all the time. The fact is that the real value of US Steel is determined at any give time solely, definitely and inexorably by supply and demand, which are accurately reflected in the transactions consummated on the floor of the exchange*."

If we were to summarize the assumptions that underlie technical analysis, we would list the following:

(1) *Market value is determined solely by the interaction of supply and demand*. We do not think that non-chartists would have any quarrels with this assumption, which describes how prices are set in any market.

(2) *Supply and demand are governed by numerous factors, both rational and irrational.* The market continually and automatically weighs all these factors. Note that a random walker would have no qualms about this assumption either. He would point out that

any irrational factors are just as likely to be on one side of the market as on the other.

(3) *Disregarding minor fluctuations in the market, stock prices tend to move in trends that persist for an appreciable length of time*. This is where random walkers would part ways with chartists. In a rational market, any trend that can be discerned by investors using charts should provide profit opportunities that when taken advantage of should eliminate the trend.

(4) *Changes in trend are caused by shifts in demand and supply*. These shifts, no matter why they occur, can be detected sooner or later in the action of the market itself. This is at the core of technical analysis. Charts, the believers argue, send advance warning of shifts in demand and supply in the form of price and volume patterns.

The views of technical analysts are best described by another quote from Magee: "*The market price reflects not only the differing fears and guesses and moods, rational and irrational, of hundreds of potential buyers and sellers, but it also reflects their needs and resources- in total, factors which defy analysis and for which no statistics are obtainable. These are nevertheless all synthesized, weighted and finally expressed in the one precise figure at which a buyer and seller get together and make a deal. The resulting price is the only figure that counts.*"

Both the anecdotal and the empirical evidence seem to suggest that investors often are irrational, at least based upon the economic definition of rationality. Whether this irrationality results in systematic price patterns is a little more difficult to assess, though the serial correlation in prices, both over short and long periods, and the periodic appearance of price bubbles in asset markets seems to indicate that irrational behavior has price effects. Finally, even if there are systematic price patterns caused by irrationality, there is the question of whether you can take advantage of these price patterns. It is entirely possible that the price patterns are so unpredictable that no investor can take advantage of them to earn excess returns. Technical analysts and chartists would disagree.

Technical Indicators and Charting Patterns

Over the years, technical analysts have developed hundreds of technical indicators and detected dozens of chart patterns that they contend help them forecast future price changes. While we cannot describe or even list all of them, we can categorize them based upon the nature of irrationality that we attribute to markets. Consolidating all of the irrationalities that have been attributed to financial markets, we have created five groupings:

- Market participants over react to new information: If this is true – prices rise too much on good news and fall too much on bad news – you would draw on contrarian indicators which would help you to gauge the direction in which the crowd is going and to go against it.

- Market participants are slow learners: In many ways, this is the polar opposite of the first grouping. If investors are slow learners, prices will under react to new information and you would expect price direction to persist and use momentum strategies, which would gauge market direction and move with it.
- Investors change their minds frequently and often irrationally, causing significant shifts in demand and supply, causing prices to move. If you believe that this is the way markets work, you would use technical indicators and charting patterns to detect these shifts.
- There are a group of investors who lead markets, and finding out when and what they are buying and selling can provide a useful leading indicator of future price movements. If this is what you believe about markets, you would track the trading of these leading investors and try to follow them.
- There are external forces that govern up and down movements in markets that override fundamentals and investor preferences. Technical indicators and charting patterns that allow up to see their larger cycles in stock prices can allow us to get ahead of other investors.

Within each, we can consider different technical indicators that we can broadly categorize into three groups – price indicators, which are based upon past price movements, volume indicators, that look at trading volume and sentiment indicators, that use qualitative measures of how bullish or bearish investors feel about stocks.

Markets overreaction - Contrarian Indicators

There are many practitioners and some economists, especially in the behavioral school, who believe that investors overreact to new information. This, in turn, can create patterns in stock prices that can be exploited by investors to earn excess returns. In this section, we consider some of the indicators, which we label contrarian, that have been developed by analysts who subscribe to this view.

The Basis for Overreaction and Implications

Why would markets over react to new information? Some researchers in experimental psychology suggest that people tend to overweight recent information and underweight prior data in revising their beliefs when confronted with new information. Others argue that a few investors tend to panic when confronted with new information, and that they take the rest of the market with them. As evidence, you could point to the strong evidence of price reversals over long periods that we presented earlier in this chapter.

If markets overreact, it follows that large price movements in one direction will be followed by large price movements in the opposite direction. In addition, the more extreme the initial price movement, the greater will be the subsequent adjustment. If markets overreact, the road to investment success seems clear. You buy assets when others are most bearish about it and selling, and sell assets when other investors are most optimistic and buying. If your assumption about market overreaction is correct, you will earn excess returns as markets correct themselves over time.

Technical Trading Rules based upon Contrarian Opinion

There are a number of indicators, some based upon price patterns, some based upon trading volume and some on market views that are designed to provide you with a sense of market direction. The objective is to not follow the market direction but to go against it and these are contrarian indicators. We will consider three widely used indicators in this section, each of which focused on a different subset of investors.

Trades that are in lots of less than a 100 are called odd-lots and are usually made by small investors. There are data services that track the number of odd-lot trades – both buys and sells - in individual stocks and in the market. As small investors become more enthusiastic about a stock, odd lot buys increase relative to sells. When they become pessimistic, the reverse occurs. To the extent that you view small investors as more likely to over react to information, you would sell as odd lot buying increases and buy as odd lot selling decrease.

But what if you believe that it is institutional investors who panic and not small investors? After all, large price movements are usually caused by institutional buying and selling, rather than by individual traders. There are indicators that track the stocks that institutions are selling and buying, with the objective of doing the opposite. There are also indicators that track the percent of mutual fund portfolios that is invested in cash and near cash investments, a good indicator of how bullish or bearish mutual fund investors are. When mutual funds are optimistic about the market, cash holdings tend to fall, whereas cash holdings increase as they become more pessimistic. If you believe that mutual fund managers over react, you would buy when they are bearish and sell when they are bullish.

Finally, you could look at investment advisors who claim to have divined the future. Investment advisory services often have their lists of most desirable and least desirable stocks. Value Line and Standard and Poor's categorize stocks into classes based upon their perceived attractiveness as investments. In keeping with the notion that the market is usually wrong, you would sell those stocks that investments advisors are most bullish on and buy those stocks where they are most bearish.

Shifting Demand

Technical analysts often argue that the greatest profits are to be made at what can be called inflection points – a fancy term for shifts in price trends from positive to negative or vice versa. Since price is ultimately determined by demand and supply, analysts often look for leading indicators of shifts in demand, especially when they are caused my emotion rather than fundamentals. If they succeed, they will make money.

The Basis for Shifting Demand and Implications

The basis for the shifting demand argument is that demand shifts cause price changes and that these demand shifts often have no basis in economic fundamentals. The anecdotal evidence seems to bear out this view. Markets often move for no discernible reason and the volatility in stock prices seems to vastly exceed the volatility in underlying value. The empirical evidence also backs up the view that prices are more volatile than fundamental value. Shiller compared stock price movements over time to movements in the present value of dividends (which he viewed as a measure of fundamental value) and concluded that stock prices were significantly more volatile (See figure 7.14)

Figure 7.14: Are markets too volatile?

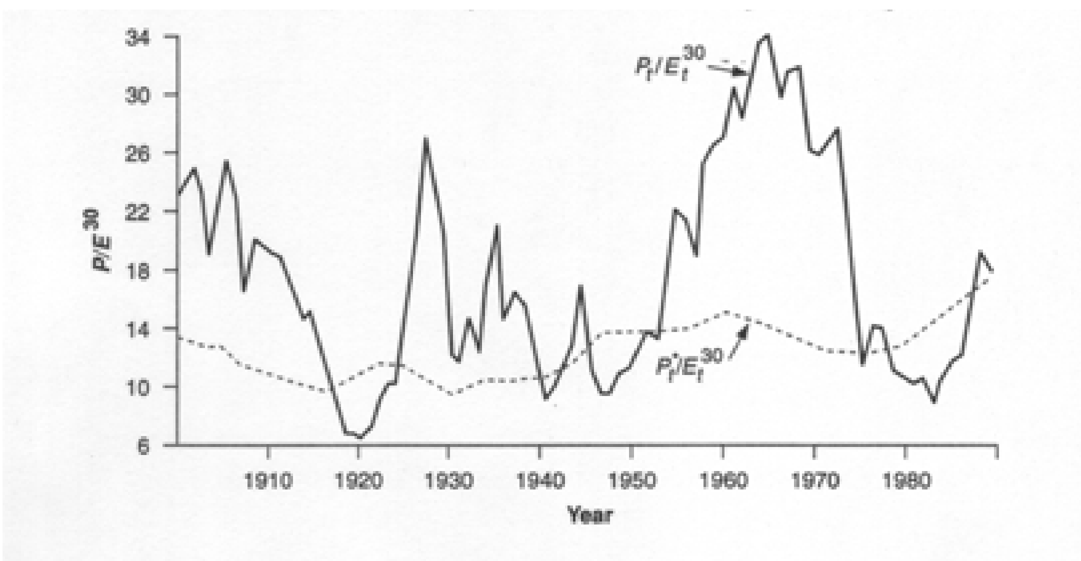

Source: Shiller

Note that the smoothed out line is the present value of dividends, whereas the volatile line represents the S&P 500.

It should be noted, though, that neither the anecdotal evidence nor Shiller's study conclusively proves emotional volatility, In fact, some researchers have argued that if the value of a stock is based upon expectations, small news announcements can cause big shifts in expectations and stock prices.

Technical Trading Rules aimed at detecting Shifting Demand

There are numerous pricing patterns and indicators that chartists claim provide advance warning of shifting demand. We will consider four broad measures here. The first relate to the entire market, and measure the breadth of the market by looking at the number of stocks that advance relative to those that decline. The argument here is that a market that goes up with limited breadth (a few stocks are creating much of the upward momentum, while the rest are flat or declining) is a market where demand (and prices) are likely to decline soon. In fact, an extension of this measure is the advance/decline line, which is reported in many financial newspapers, where you graph the ratio of the number of stocks that have gone up to the number of stocks that have dropped. Here again, analysts argue that a divergence between index levels and the advance/decline line – a drop in the index accompanied by an improvement in the advance/decline line may indicate an upcoming shift towards buying.

The second is the presence (at least perceived presence) of support and resistance lines in prices. A resistance line is an upper bound on the price whereas a support line represents a lower bound on the price. Both are extracted by looking at past prices. Thus, a stock that has tended to move between \$ 20 and \$ 40 over the last few periods has a support line at \$ 20 and a resistance line at \$ 40. It may be pure coincidence though we think not but support and resistance lines often are nice round numbers – you very seldom see a resistance line at \$ 39.88 and a support line at \$ 21.13. Figure 7.15 provides a chart with support and resistance lines.

Figure 7.15: Support and Resistance Lines

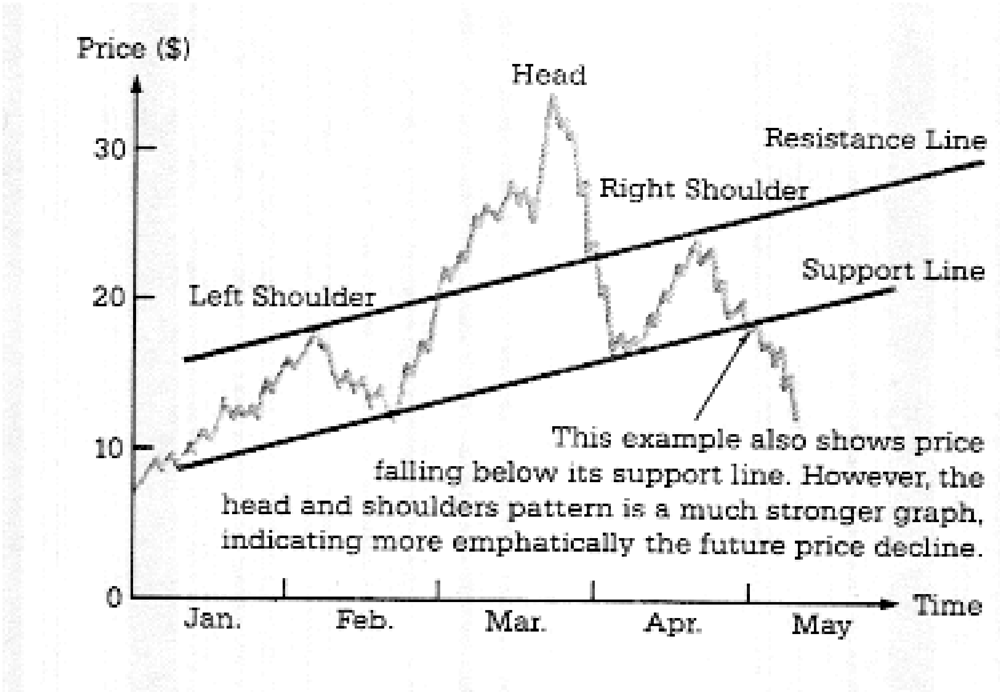

The fact that the stock stays below the resistance line and above the support line is not news, but a stock that breaks through either gets attention. When a stock breaks through the resistance line, technical analysts view it as a sign of a shift in demand upwards and the beginning of a sustained upward movement in prices. Conversely, when a stock falls below the support line, analysts view it as a breakdown in demand and the precursor of a further decline in prices. While the notion of arbitrary support and resistance lines strikes us as fanciful, if enough investors buy into their existence, there can be a self-fulfilling prophecy. To see why, assume that a stock with a resistance line of $ 40 millions sees its stock price go up to $40.50. Investors who believe that this is a beginning of a surge in prices will all try to buy the stock on the event, causing the stock price to go up. Whether such a price increase can be sustained for more than a few days is an open question. In the graph, you can also see another widely followed chart pattern, called “head and shoulders”. In fact, there are hundreds of patterns that chartists have uncovered over time that have been offered as leading indicators of price changes.[23]

Central to much of technical analysis is a reverence for moving averages, i.e., averages of stock prices over the last few months or weeks. Often, you will see price charts

---

[23] For a comprehensive listing of indicators, see “The Encyclopedia of Technical Market Indicators” by Robert Colby and Thomas Myers, Irwin.

with a moving average line superimposed on actual prices. Again, analysts view any deviation of stock prices from a moving average line as an indication of an underlying shift in demand that can be exploited for profits.

Analysts have also long used a charting technique called point and figure to detect trends in prices. The essential feature of a point and figure chart is that it is composed of a series of Xs and Os. Each X represents a price movement of a given size called a box size. As long as prices continue to rise, Xs are added to the column. If there is a price decline of more than a given magnitude (called the reversal size), a new column of Os is opened. Figure 7.16 presents a point and figure chart.

Figure 7.16: Point and Figure Chart

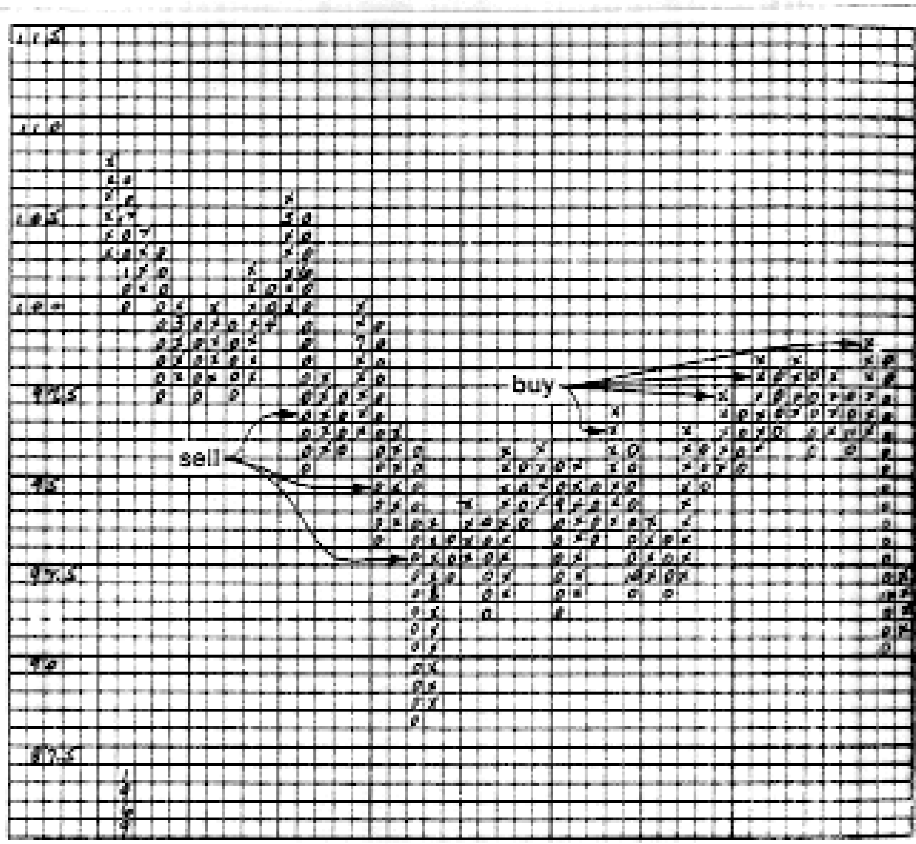

FIG. 23-5 Point and figure chart of Standard & Poor's composite 500. (From A. W. Cohen, *Technical Indicator Analysis*, Chartcraft, Inc., Larchmont, N.Y.)

In recent years, information on <u>trading volume</u> for individual stocks has become increasingly accessible. Technical analysts now routinely look at trading volume for clues of future price movements, either in conjunction with price changes or by itself. For instance, an increase in the stock price that is accompanied by heavy trading volume is considered a more positive prognosticator of future price increases than one generated with light volume.

Empirical Evidence on Technical Indicators

There is not much empirical evidence for or against many of the individual charting patterns. Part of the reason for this is that many of these patterns are so subjectively defined – different analysts use different and often shifting definitions of what comprises a support or a resistance line, for instance - that they cannot be tested empirically, which serves both sides of the argument very well. Supporters of charting can then use their own tests which are often biased to offer proof that their patterns works. Opponents of technical analysis can rest secure in their absolute conviction that charting is for the naïve and the misguided and not worry about evidence to the contrary.

It is quite ironic that some of the best defenses of technical analysis have been offered by academics who would not categorize themselves as chartists or technical analysts. Lo, Wang and Mamaysky (2000) present a fairly convincing defense of technical analysis from the perspective of financial economists. They use daily returns of stocks on the New York Stock Exchange and NASDAQ from 1962 and 1996 and use the most sophisticated computational techniques (rather than human visualization) to look for pricing patterns. They find that the most common patterns in stocks are double tops and bottoms, followed by the widely used head and shoulders pattern. In other words, they find evidence that some of the most common patterns used by technical analysts exist in prices. Lest this be cause for too much celebration among chartists, they also point out that these patterns offer only marginal incremental returns (an academic code word for really small) and offer the caveat that these returns may not survive transactions cots.

Are currency markets different?

While there is little empirical evidence to back the use of charts in the stock market, a number of studies claim to find that technical indicators may work in currency markets. To name a few:

- Filter rules, where you buy a currency if it goes up by x% and sell if it goes down by the same amount earned substantial profits in the Deutsche mark, yen and sterling markets between 1973 and 1981.[24]
- Moving average rules would have generated excess returns in foreign currency markets.[25]
- Head and Shoulder patterns would have generated excess returns in the pound sterling, Canadian dollar, French franc and Swiss franc markets between 1973 and1994. [26]

Though there are dissenting voices, there clearly seem to be more opportunities for technical analysis in currency markets. Some attribute it to central bank intervention. When central banks target exchange rates, they can generate speculative profits for investors. Another possibility is that the foreign currency market is less efficient that the stock market.

Slow Learning Markets: Momentum Indicators

If investors are slow to assess the effects of new information on stock prices, you can see sustained up or down movements in stock prices after news comes out about the stock – up movements after good news and down movements after bad news. There are analysts who contend that this is indeed the case and create trading rules that take advantage of this slow learning process. Since these rules are based upon the assumption that trends in prices tend to continue for long periods, they can be categorized as momentum rules.

The Basis for Slow Learning and Implications

What is the evidence that markets learn slowly? The best support for slow learning markets comes from studies that look at information events such as earnings announcements or acquisitions. As we will see later in this book, there is evidence that markets continue to adjust to the information well after it has come out. For instance, a firm that reports much better than expected earnings will generally see its stock price jump on the announcement and continue to drift upwards for the next few days. The same seems to occur to a target firm in an acquisition. While there are alternative explanations for price

---

[24] See "Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981" by Dooley, M.P. and J.R. Shafer in Exchange Rate and Trade Instability, Causes, Consequences and Remedies, 1983, Ballinger.

[25] See "Time Varying Risk Premia, Volatility and Technical Trading Rules" by B.C. Kho, Journal of Financial Economics, v41, 246-290.

[26] See "Head and Shoulders: Not a flaky pattern", by Osler,C.L. and P.H.K. Chang, Staff Paper, 1995, Federal Reserve Bank of New York.

drifts, one potential explanation is that markets learn slowly and that it takes them a while to assimilate the information.

If markets learn slowly, you should expect to see prices move in the same direction after a precipitating action. If the initial news was good – a good earnings report or an earnings upgrade from an analyst – you should expect to see upward price momentum. If the news was bad, you should expect to see the opposite. In fact, recent empirical studies (referenced in the earlier part of this chapter) have found evidence of price momentum in equity markets in the United States at least in the short term.

Technical Indicators to take advantage of slow learning markets

Momentum investors firmly believe that the trend is your friend and that it is critical that you look past the day-to-day movements in stock prices at the underlying long-term trends. The simplest measure of trend is a trend line. Figure 7.17 contains two trend lines – the graph on the left is for a silver futures contracts over the few months of its existence and the graph on the right is for cocoa futures over a much longer time period.

Figure 7.17: Trend Lines

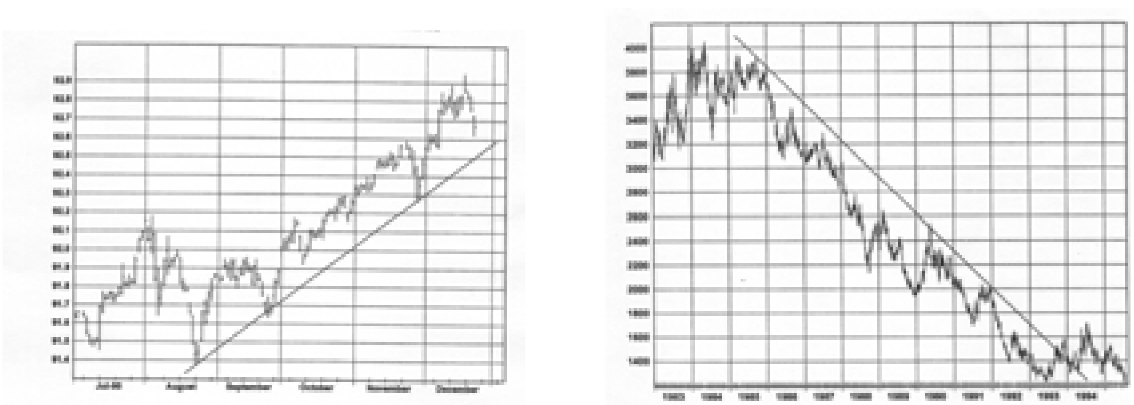

In this silver futures contract to the left, you see an uptrend line, drawn by connecting a series of lows in prices, each one higher than the other. On the right, cocoa prices had been declining over the period in question and a down trend line is drawn by connecting a series of lower highs. As momentum investors, you would buy stocks that are going up and staying above the uptrend line. If the price falls below the uptrend line, it is viewed as a negative sign. Conversely, if the price rises above a down trend line, it is considered a bullish sign.

Stocks with highest relative strength: Take a look at the 50 stocks with the highest relative strength over the last 6 months.

A closely followed momentum measure is called relative strength, which is the ratio of the current price to an average over a longer period (say six months or a year). Stocks that score high on relative strength are therefore stocks that have gone up the most over the period, whereas those that score low are stocks that have gone down. The relative strength can be used either in absolute terms, where only stocks that have gone up over the period would be considered good investments. Alternatively, the relative strength can be compared across stocks, and you invest in stocks that show the highest relative strength – i.e, have gone up the most, relative to other stocks.

Following the Informed Investors: Leading Indicators

This approach is the flip side of the contrarian approach. Instead of assuming that investors, on average, are likely to be wrong, you assume that they are right To make this assumption more palatable, you do not look at all investors but only at the investors who presumably know more than the rest of the market.

The Basis for Following Smart Investors and Implications

Are some investors smarter and better informed than others? Undoubtedly. Do they make higher returns as a consequence? Not necessarily. As Keynes was fond of pointing out, a stock market is a beauty contest, where the prize goes to the person who best gauges who the other judges in the contest will pick as the winner. In investment terminology, the high returns often go to the investor who can best pick the stocks that other investors will buy.

There are two keys to making a strategy of following other investors work. The first is identifying the smart investors, who may not always be the largest or best known. It stands to reason that investors who have access to the best information are most likely to beat the market and would be the ones that you should follow. The second is to find out when and what these smart investors are trading in a timely fashion, so that you can imitate them. This is often difficult to do. Even though insiders and institutions have to file with the Securities and Exchange Commission (SEC), providing details about their trades, the filings are made several weeks after the trades occur.

Technical Indicators for Followers

There are several technical indicators that attempt to pinpoint what better informed investors are buying and selling. Here, we consider two. The first looks at short sales made by market specialists. Since these specialists are close to the action and have access to information that the rest of us cannot see (such as the order book and trading on the floor), it can be argued that they should have an inside track on over priced and under priced stocks. Thus, a surge in specialist short sales in a stock would be a precursor for bad news on the stock and a big price drop. Some analysts look at all short sales made on a stock, arguing that only larger, more sophisticated investors can short stock in the first place. A study by Senchack and Starks in 1993 provides some support for this indicator by noting that stock returns tend to me more negative for stocks where the short interest (short sales as a percent of the outstanding stock) is higher.

In the last few years, as the SEC has speeded up the process of recording transactions by insiders and has made this data more easily accessible to the public. You can therefore look up stocks where insider buying or selling has increased the most. In fact, the ratio of insider buying to selling is often tracked for stocks with the idea that insiders who are buying must have positive information about a stock whereas insiders who are selling are likely to have negative information.

Long Term Cycles: Mystical Indicators

The final set of technical indicators are based upon long term cycles in prices that exercise an inexorable hold on how prices move. Since these long-term cycles operate independently of fundamentals, it is very difficult to explain them without resorting to mysticism.

Basis for long term cycles and Implications

There are two ways in which you can defend the use of long-term cycles. One is to abandon any basis in rationality and argue that there are a number of phenomena in nature that cannot be explained with models.[27] You can think of such investors as subscribers to the karmic theory of investing. In other words, everything that happens has already been pre-destined and there is nothing that we can do to stop it. This requires an almost religious

[27] Scientists would undoubtedly disagree.

belief that cycles will replicate themselves. The other defense is based on market behavior. You can argue that investors, even though they might be separated over time, behaved in very much the same way in the South Sea Bubble as they did in the dot-come bubble. Consequently, long term cycles reflect the pricing mistakes that investors make and remake over time. As a cautionary note, you should realize that if you look for patterns too intently in charts, you will find them, especially if you use visual techniques (rather than statistical ones).

Technical Indicators based upon Cycles

While there are numerous cycles that analysts see in stock prices, we will consider two in this section. In the first, the Dow Theory, the market is considered as having three movements, all going at the same time. The first is the narrow movement (daily fluctuations) from day to day. The second is the short swing (secondary movements) running from two weeks to a few months and the third is the main movement (primary trends) covering at several years in its duration. Proponents of the theory claim that you can tell where you are in the primary cycle by charting the industrial and transportation components of the Dow Index and looking for confirmation (i.e, both indices moving in the same direction). In figure 7.18, the Dow Theory is presented:

Figure 7.18: The Dow Theory

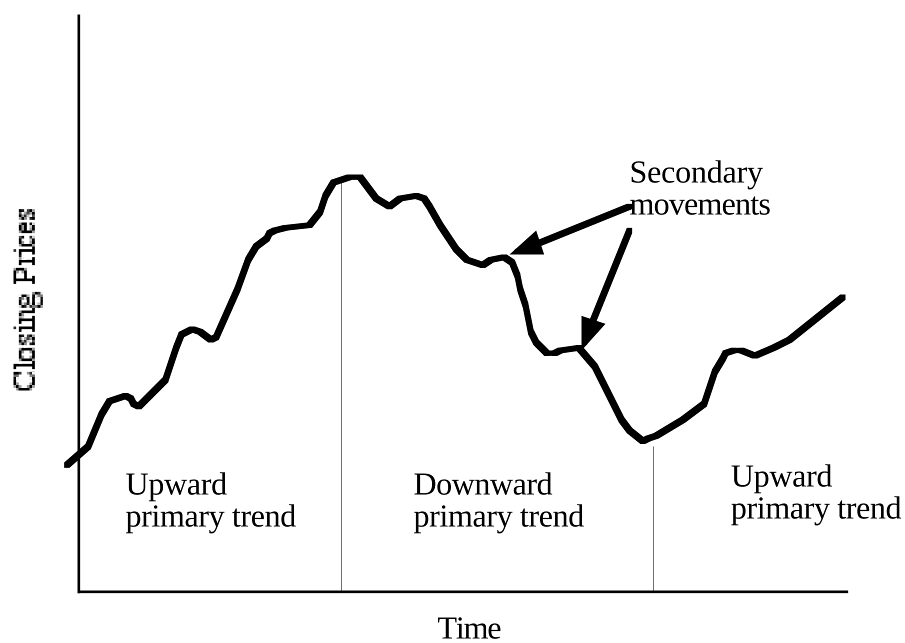

In 1922, William Hamilton wrote a book titled "The Stock Market Barometer"about the Dow Theory, where he presented evidence on the measure's efficacy at predicting market movements. A recent study[28] appraised Hamilton's predictions in the Wall Street Journal between 1901 and 1929 and concluded that he had far too many correct calls than could be attributed to chance and that you would have earned excess returns following his advice.

While the Dow Theory has been around for decades, the Elliott Wave acquired a wide following in the 1980s. Elliot's theory was that the market moves in waves of various sizes, from those encompassing only individual trades to those lasting centuries, perhaps longer. In the classic Elliot wave, a cycle lasts 200 years and has 8 waves – five up and three down – with smaller cycles within each of these waves. By classifying these waves and counting the various classifications, he argued that it was possible to determine the relative positions of the market at all times.

Figure 7.19: The Elliott Wave

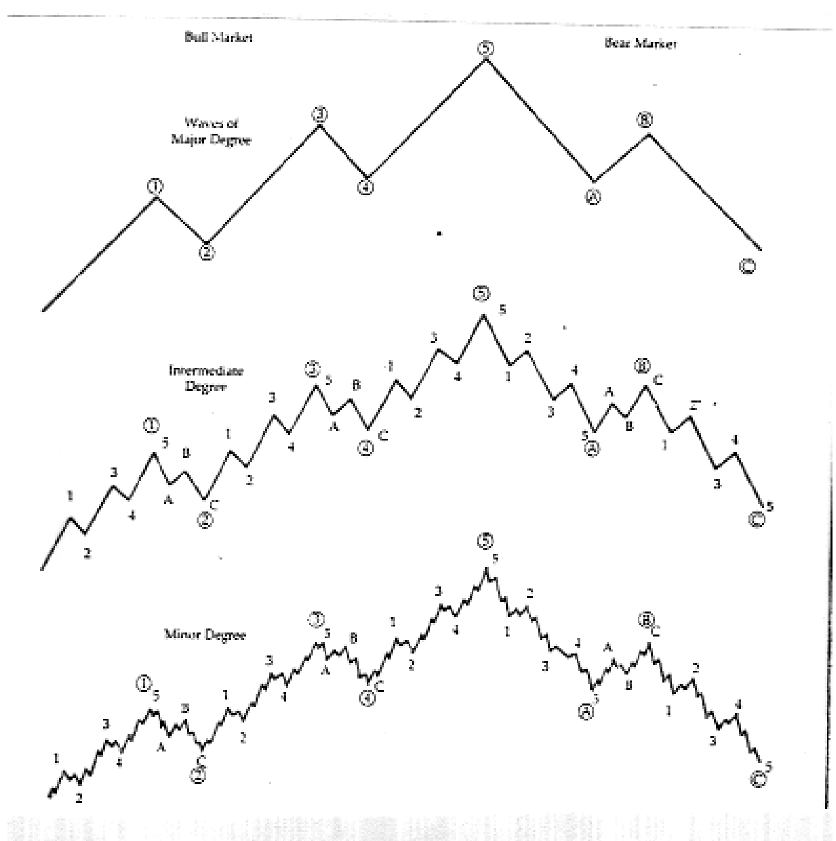

---

[28] See "The Dow Theory: William Peter Hamilton's Track Record Reconsidered", by Brown, Goetzmann and Kumar. They conclude that following Hamilton's advice would have generated excess returns of about 4.04% a year.

In the aftermath of the 1987 crash, there were several newsletters that based upon the Elliott Wave.[29] Most of them faded in the years after, as the predictive power of the model was found to be wanting.

Other cycles include: the *Kitchen cycle* (inventories, 3-5 years); the *Juglar Cycl***e** (fixed investment patterns, 7-11 years); and *Kuznets Cycle* (building patterns, 15-25 years). Other more controversial theories include: the *Kondratyev Cycle* (also called "the long economic cycle," about 54 years) in three stages of upswing, crisis, and depression. The *Babson chart* of business barometers uses statistics and charts to model a 20-year cycle in four stages: overexpansion, decline, depression, and improvement.

Determinants of Success at Charting and Technical Analysis

Can you succeed with technical indicators and charts? The answer that has been long given by academics and fundamentals is no, but that answer may need to be reassessed in light of the research on price patterns (especially price momentum) and trading volume in recent years. There seems to be enough evidence now for us to conclude that it is foolhardy to ignore recent price movements and changes in trading volume when investing in a stock. So what are the essential ingredients for success with technical analysis? These seem to be a few:

- If you decide to use a charting pattern or technical indicator, you need to be aware of the investor behavior that gives rise to its success. This is not just to satisfy your curiosity but also to ensure that you can modify or abandon the indicator if the underlying behavior changes.
- It is important that you back-test your indicator to ensure that it delivers the returns that are promised. In running these tests, you should pay particular attention to the volatility in performance over time and how sensitive the returns are to holding periods. There are some strategies that work only in bull markets, for example, and only for specific holding periods – say 1 month or less.
- The excess returns on many of the strategies that we described in this chapter seem to depend upon timely trading. In other words, to succeed at some of these strategies, you may need to monitor prices continuously, looking for the patterns that would trigger trading.

---

[29] The best known book on the Elliott Wave was written by Frost and Prechter and is titled “The Elliott Wave Principle”. Gehm (1983) provides a critical look at the Elliott Wave and argues that its ambiguity makes it impossible to test.

- Building on the theme of time horizons, success at charting can be very sensitive to how long you hold an investment. Recall, for instance, that momentum indicators seem to work for a few months and that reversals seem to occur beyond that time period. Finding the optimal holding period and staying disciplined seem to be key to earning the returns that we sometimes see on paper.
- The strategies that come from technical indicators are generally short-term strategies that require frequent and timely trading. With some strategies, you may need to trade several times during the course of a day or a week. Not surprisingly, these strategies also generate large trading costs that can very quickly eat into any excess returns you may have.

In summary, investors who can track markets continuously and trade cheaply may be able to take advantage of price patterns and volume indicators to earn excess returns, if they can pinpoint the right indicators and stay disciplined. As price and volume data becomes increasingly available to all investors, though, it is likely that these strategies will be more useful as secondary strategies, used to augment returns on a primary strategy. For instance, a growth investor who buys stocks with rising earnings may also consider adding price momentum to the mix of variables that she looks at before making her investment choices. Investors who cannot or do not want to track markets continuously are unlikely to earn enough returns on these strategies to cover transactions costs.

Conclusion

Investors have always claimed to find patterns in charts that help them make better investment decisions. Skeptics have viewed these claims as fiction and have argued that there is no basis to technical analysis. In recent years, evidence has steadily accumulated that there is information in past price movements and trading volume and that there may be a foundation for some of the claims made by chartists. In particular, stocks that have done well in the recent past seem to be more likely to do well in the near future (price momentum) and trading volume changes seem to lead price changes in some markets.

All technical indicators have their basis in quirks in human behavior. We categorize technical trading indicators based upon the type of behavior that may lead to their success. Contrarian indicators such as mutual fund holdings or odd lot ratios, where you track what investors are buying and selling with the intention of doing the opposite, are grounded in the belief that markets over react. A number of technical indicators are built on the presumption that investors often change their views collectively, causing shifts in demand and prices, and that patterns in charts – support and resistance lines, price relative to a moving average- can predict these changes. With momentum indicators, such as relative strength and trend lines, you are assuming that markets often learn slowly and that it takes time for prices to adjust to true values. If you believe that there are some traders who trade ahead of the market, either because they have better analysis tools or information, your indicators will follow these traders – specialist short sales and insider buying/selling, for instance – with the objective of piggy-backing on their trades. Finally, if you believe that there are long-term cycles in stock prices, your investment strategy may be driven by the cycle you subscribe to and where you believe you are in the cycle.

If you are a short-term investor with the discipline to stick with a tested indicator, low trading costs and continuous access to information, you may be able to use technical indicators as the basis for your investment strategy. Even those who do not want to build their entire strategy around price patterns and trading volume may still find them useful to augment returns on their primary strategies.

Lessons for Investors

To be a successful technical analyst, you need to:

1. Understand human nature: Investors are human and display all off the foibles of human nature. Some of them tend to be over confident and to over react and move in herds. At the same time, others display too little confidence, learn too slowly and are born contrarians. What happens in markets represents the tug and the pull between these groups. When you use an indicator, you need to understand the assumption about human behavior that underlies it.
2. Not mistake random price movements for price patterns: Even when prices move randomly, you can generate charts that look like they have patterns. Even bubbles and crashes, which are used by many analysts as evidence of irrationality, can exist in rational markets.
3. Have a time horizon that matches your indicator: Some indicators require time horizons of a few hours, others require a few weeks and some may even stretch for a few months.
4. Be disciplined: If you decide to use a technical indicator to pick stocks, assuming you have back-tested the indicator, you will need to stay within your specified strategy.

CHAPTER 8

GRAHAM'S DISCIPLES: VALUE INVESTING

Value investors are bargain hunters and many investors describe themselves as such. But who is a value investor? In this chapter, we begin by addressing this question, and argue that value investors come in many forms. Some value investors use specific criteria to screen for what they categorize as undervalued stocks and invest in these stocks for the long term. Other value investors believe that bargains are best found in the aftermath of a sell-off, and that the best time to buy a stock is when it is down. Still others adopt a more activist approach, where they buy large stakes in companies that they believe are under valued and push for changes that they believe will unleash this value.

Value investing is backed by empirical evidence from financial theorists and by anecdotal evidence – the success of value investors like Ben Graham and Warren Buffett are part of investment mythology – but it is not for all investors. We will consider what investors need to bring to the table to succeed at value investing.

Who is a value investor?

Morningstar is a widely used source of mutual fund information, and it categorized 38% of mutual funds as value funds in 2001. But how did it make this categorization? While it did look at the way these funds described themselves in their prospectuses, the ultimate categorization was based on a far simpler measure. Any fund that invested in stocks with low price to book value ratios or low price earnings ratios, relative to the market, was categorized as a value fund. This is a fairly conventional categorization, but we believe that it is too narrow a definition of value investing and misses the essence of value investing.

Another widely used definition of value investors suggests that they are investors interested in buying stocks for less that what they are worth. But that is too broad a definition since you could potentially categorize most active investors as value investors on this basis. After all, growth investors (who are often viewed as competing with value investors) also want to buy stocks for less than what they are worth. So what is the essence of value investing? To understand value investing, we have to begin with the proposition that the value of a firm is derived from two sources – investments that the firm has already made (assets in place) and expected future investments (growth opportunities). What sets value investors apart is their desire to buy firms for less than what their assets-in-place are worth. Consequently, value investors tend to be leery of large premiums paid by markets for growth opportunities and try to find their best bargains in more mature companies that are out of favor.

Even with this definition of value investing, there are three distinct strands that we see in value investing. The first and perhaps simplest form of value investing is passive screening, where companies are put through a number of investment screens – low PE ratios, assets which are easily marketable, low risk etc. – and those that pass the screens are categorized as good investments. In its second form, you have contrarian value investing, where you buy assets that are viewed as untouchable by other investors because of poor past performance or bad news about them. In its third form, you become an activist value investor, who buys equity in under valued or poorly managed companies but then use the power of your position (which has to be a significant one) to push for change that will unlock this value.

The Passive Screener

There are many investors who believe that stocks with specific characteristics – good management, low risk and high quality - outperform other stocks, and that the key to investment success is to identify what these characteristics are. While investors have always searched for these characteristics, it was Ben Graham, in his classic books on security analysis (with David Dodd), who converted these qualitative factors into quantitative screens that could be used to find promising investments. In recent years, as data has become more easily accessible and computing power has expanded, these screens have been refined and extended, and variations are used by many portfolio managers and investors to pick stocks.

Ben Graham: The Father of Screening

Many value investors claim to trace their antecedents to Ben Graham and to use the book on Security Analysis that he co-authored with David Dodd, in 1934 as their investment bible. But who was Ben Graham and what were his views on investing? Did he invent screening and do his screens still work?

Graham's screens

Ben Graham started life as a financial analyst and later was part of an investment partnership on Wall Street. While he was successful on both counts, his reputation was made in the classroom. He taught at Columbia and the New York Institute of Finance for more than three decades and during that period developed a loyal following among his students. In fact, much of Ben's fame comes from the success enjoyed by his students in the market.

It was in the first edition of "Security Analysis" that Ben Graham put his mind to converting his views on markets to specific screens that could be used to find under valued stocks. While the numbers in the screens did change slightly from edition to edition, they preserved their original form and are summarized below:

1. Earnings to price ratio that is double the AAA bond yield.
2. PE of the stock has to less than 40% of the average PE for all stocks over the last 5 years.
3. Dividend Yield > Two-thirds of the AAA Corporate Bond Yield
4. Price < Two-thirds of Tangible Book Value[1]
5. Price < Two-thirds of Net Current Asset Value (NCAV), where net current asset value is defined as liquid current assets including cash minus current liabilities
6. Debt-Equity Ratio (Book Value) has to be less than one.
7. Current Assets > Twice Current Liabilities
8. Debt < Twice Net Current Assets
9. Historical Growth in EPS (over last 10 years) > 7%
10. No more than two years of declining earnings over the previous ten years.

Any stock that passes all 10 screens, Graham argued, would make a worthwhile investment. It is worth noting that while there have been a number of screens that have been developed by practitioners since these first appeared, many of them are derived from or are subsets of these original screens.

The Performance

How well do Ben Graham's screens work when it comes picking stocks? Henry Oppenheimer: studied the portfolios obtained from these screens from 1974 to 1981 and concluded that you could have made an annual return well in excess of the market. As we will see later in this section, academics have tested individual screens – low PE ratios and high dividend yields to name two – in recent years and have found that they indeed yield portfolios that deliver higher returns. Mark Hulbert, who evaluates the performance of investment newsletters, found newsletters than espoused to follow Graham did much better than other newsletters.

> *Stocks that pass the Graham screens:* Take a look at the stocks that currently pass the Graham Screens.

The only jarring note is that an attempt to convert the screens into a mutual fund that would deliver high returns did fail. In the 1970s, an investor name James Rea was convinced enough of the value of these screens that he founded a fund called the Rea-Graham fund,

---

[1] Tangible Book value is computed by subtracting the value of intangible assets such as goodwill from the total book value.

which would invest in stocks based upon the Graham screens. While it had some initial successes, the fund floundered during the 1980s and early 1990s and was ranked in the bottom quartile for performance.

The best support for Graham's views on value investing don't come from academic studies or the Rea-Graham fund but from the success of many of his students at Columbia. While they chose diverse paths, many of them ended up managing money and posting records of extraordinary success. In the section that follows, we will look at the most famous of his students – Warren Buffett.

Graham's maxims on investing

Janet Lowe, in her biography of Ben Graham, notes that while his lectures were based upon practical examples, he had a series of maxims that he emphasized on investing. Since these maxims can be viewed as the equivalent of the ten commandments of value investing, they are worth revisiting.

1. Be an investor, not a speculator. Graham believed that investors bought companies for the long term, but speculators looked for short term profits.
2. Know the asking price. Even the best company can be a poor investment at the wrong (too high) price.
3. Rake the market for bargains. Markets make mistakes.
4. Stay disciplined and buy the formula:

$$E\ (2g + 8.5) * \text{T.Bond rate}/Y$$

where E = Earnings per share, g= Expected growth rate in earnings, Y is the yield on AAA rated corporate bonds and 8.5 is the appropriate multiple for a firm with no growth. For example consider a stock with $ 2 in earnings in 2002 and 10% growth rate, when the treasury bond rate was 5% and the AAA bond rate was 6%. The formula would have yielded the following price:

Price = \$2.00 (2 (10)+8.5)* (5/6) = \$47.5

If the stock traded at less than this price, you would buy the stock.

5. Regard corporate figures with suspicion, advice that carries resonance in the aftermath of recent accounting scandals.
6. Diversify. Don't bet it all on one or a few stocks.
7. When in doubt, stick to quality.
8. Defend your shareholder's rights. This was another issue on which Graham was ahead of his time. He was one of the first advocates of corporate governance.
9. Be patient. This follows directly from the first maxim.

It was Ben Graham who created the figure of Mr. Market which was later much referenced by Warren Buffett. As described by Mr. Graham, Mr Market was a manic-depressive who does not mind being ignored, and is there to serve and not to lead you. Investors, he argued, could take advantage of Mr. Market's volatile disposition to make money.

Warren Buffett: Sage from Omaha

No investor is more lionized or more relentlessly followed than Warren Buffet. The reason for the fascination is not difficult to fathom. He has risen to become one of the wealthiest men in the world with his investment acumen, and the pithy comments on the markets that he makes at stockholder meetings and in annual reports for his companies are widely read. In this section, we will consider briefly Buffett's rise to the top of the investment world, and examine how he got there.

Buffett's History

How does one become an investment legend? Warren Buffett started a partnership with seven limited partners in 1956, when he was 25, with $105,000 in funds. He generated a 29% return over the next 13 years, developing his own brand of value investing during the period. One of his most successful investments during the period was an investment in American Express, after the company's stock price tumbled in the early 1960s. Buffett justified the investment by pointing out that the stock was trading at far less than what the American Express generated in cash flows for the company for a couple of years. By 1965, the partnership was at $26 million and was widely viewed as successful.

The moment that made Buffett's reputation was his disbanding of the partnership in 1969 because he could not find any stocks to buy with his value investing approach. At the time of the disbanding, he said "*On one point, I am clear. I will not abandon a previous approach whose logic I understand, although I might find it difficult to apply, even though it may mean foregoing large and apparently easy profits to embrace an approach which I don't fully understand, have not practiced successfully and which possibly could lead to substantial permanent loss of capital*" The fact that a money manager would actually put his investment philosophy above short term profits, and the drop in stock prices in the years following this action played a large role in creating the Buffett legend.

Buffett then put his share of partnership ((about $25 million) into Berkshire Hathaway, a textile company whose best days seemed to be in the past. He used Berkshire Hathaway as a vehicle to acquire companies (GEICO in the insurance business and non-insurance companies such as See's candy, Blue Chip Stamps and Buffalo News) and to make investments in other companies (Am Ex, Washington Post, Coca Cola, Disney). His golden touch seemed to carry over and Berkshire Hathaway's stock price reflected his success (See figure 8.1 below):

Figure 8.1: Berkshire Hathaway

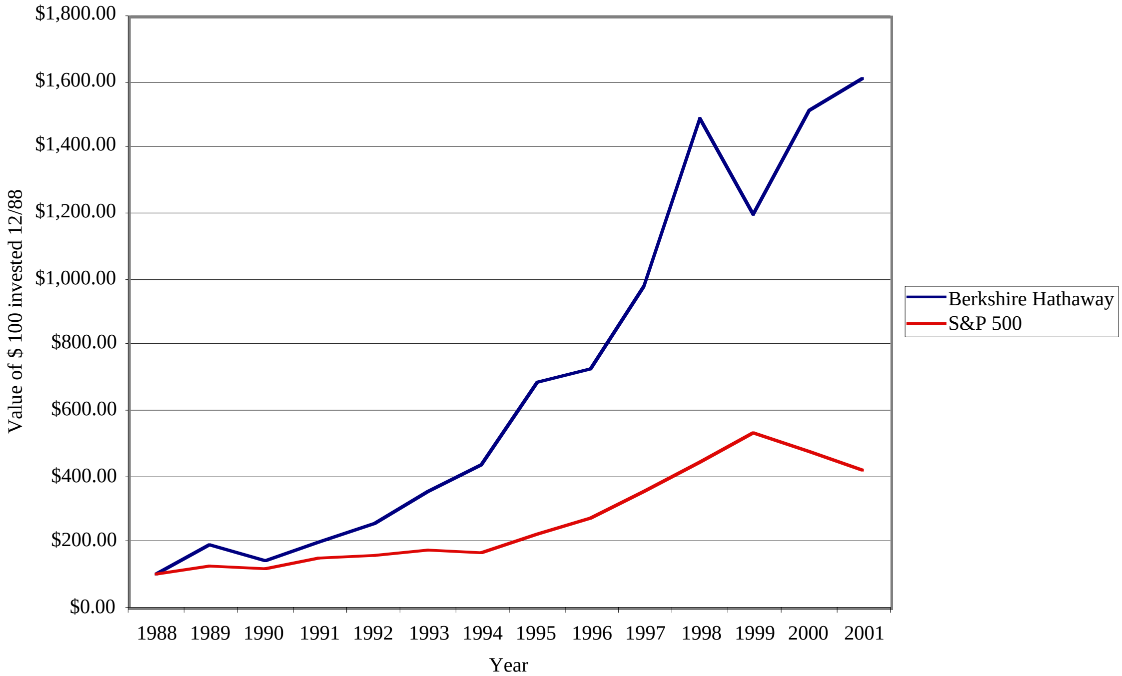

Source: Raw data from Bloomberg

An investment of $ 100 in Berkshire Hathaway in December 1988 would have outstripped the S&P 500 four-fold over the next thirteen years.

As CEO of the company, Buffett broke with the established practices of other firms in many ways. He refused to fund the purchase of expensive corporate jets and chose to keep the company in spartan offices in Omaha, Nebraska. He also refused to split the stock as the price went ever higher to the point that relatively few individual investors could afford to buy a round lot in the company. On December 31, 2001, a share of Berkshire Hathaway stock was trading at $75,600, making it by far the highest priced listed stock in the United States. He insisted on releasing annual reports that were transparent and included his views on investing and the market, stated in terms that could be understood by all investors.

Buffett's Tenets

Roger Lowenstein, in his excellent book on Buffett, suggests that his success can be traced to his adherence to the basic notion that when you buy a stock, you are buying an underlying business and the following tenets:

Business Tenets:

- The business the company is in should be <u>simple and understandable</u>. In fact, one of the few critiques of Buffett was his refusal to buy technology companies, whose business he said was difficult to understand.
- The firm should have a <u>consistent operating history</u>, manifested in operating earnings that are stable and predictable.
- The firm should be in a business with <u>favorable long term prospects</u>.

Management Tenets:

- The managers of the company should be <u>candid</u>. As evidenced by the way he treated his own stockholders, Buffett put a premium on managers he trusted. Part of the reason he made an investment in Washington Post was the high regard that he had for Katherine Graham, who inherited the paper from her husband.
- The managers of the company should be <u>leaders and not followers</u>. In practical terms, Buffett was looking for companies that mapped out their own long term strategies rather than imitating other firms.

Financial Tenets:

- The company should have a <u>high return on equity</u>, but rather than base the return one equity on accounting net income, Buffett used a modified version of what he called owner earnings

Owner Earnings = Net income + Depreciation & Amortization – Capital Expenditures

Harking back to chapter 5, where we looked at valuation, note that this is very close to a free cash flow to equity.

- The company should have <u>high and stable profit margins</u>.and a history of creating value for its stockholders.

Market Tenets:

- In determining value, much has been made of Buffett's use of a riskfree rate to discount cash flows. Since he is known to use conservative estimates of earnings and since the firms he invests in tend to be stable firms, it looks to us like he makes his risk adjustment in the cashflows rather than the discount rate.[2]
- In keeping with Buffett's views of Mr. Market as capricious and moody, even valuable companies can be bought at attractive prices when investors turn away from them.

---

[2] In traditional capital budgeting, this approach is called the certainty equivalent approach, where each expected cash flow is replaced with a lower cash flow, representing its certainty equivalent.

Assessing Buffett

It may be presumptuous of us to assess an investor who has acquired mythic status but is Warren Buffett the greatest investor ever? If so, what accounts for his success and can it be replicated? We believe that his reputation is well deserved and that his extended run of success cannot be attributed to luck. While he has had his bad years, he has always bounced back in subsequent years. The secret to his success seems to rest on the long view he brings to companies and his discipline – the unwillingness to change investment philosophies even in the midst of short term failure.

Much has been made about the fact that Buffett was a student of Graham at Columbia University, and their adherence to value investing. Warren Buffett's investment strategy is more complex than Graham's original passive screening approach. Unlike Graham, whose investment strategy was inherently conservative, Buffett's strategy seems to extend across a far more diverse range of companies, from high growth firms like Coca Cola to staid firms such as Blue Chip Stamps. While they both may use screens to find stocks, the key difference, as we see it, between the two men is that Graham strictly adhered to quantitative screens whereas Buffett has been more willing to consider qualitative screens. For instance, he has always put a significant weight on both the credibility and competence of top managers when investing in a company.

In more recent years, he has had to struggle with two by-products of his success. His record of picking winners has attracted a crowd of imitators who follow his every move and buy everything be buys, making it difficult for him to accumulate large positions at attractive prices. At the same time the larger funds at his disposal imply that he is investing far more than he did two or three decades ago in each of the companies that he takes a position in, which makes it more difficult for him to be a passive investor. It should come as no surprise, therefore, that he is a much more activist investor than he used to be, serving on boards of the Washington Post and other companies and even operating as interim chairman of Salomon Brothers during the early 1990s.

Be like Buffett?

Warren Buffett's approach to investing has been examined in detail and it is not a complicated one. Given his track record, you would expect a large number of imitators. Why, then, do we not see other investors, using his approach, replicate his success? There are three reasons:

- Markets have changed since Buffett started his first partnership. His greatest successes did occur in the 1960s and the 1970s, when relatively few investors had access to information about the market and institutional money management was not dominant. Even Warren Buffett would have difficulty replicating his success in today's market, where information on companies is widely available and dozens of money managers claim to be looking for bargains in value stocks.

- In recent years, Buffett has adopted a more activist investment style and has succeeded with it. To succeed with this style as an investor, though, you would need substantial resources and have the credibility that comes with investment success. There are few investors, even among successful money managers, who can claim this combination.
- The third ingredient of Buffett's success has been patience. As he has pointed out, he does not buy stocks for the short term but businesses for the long term. He has often been willing to hold stocks that he believes to be under valued through disappointing years. In those same years, he has faced no pressure from impatient investors, since stockholders in Berkshire Hathaway have such high regard for him. Many money managers who claim to have the same long time horizon that Buffett have come under pressure from investors wanting quick results.

In short, it is easy to see what Warren Buffett did right over the last half century but it will be very difficult for an investor to replicate that success. In the sections that follow, we will examine both the original value investing approach that brought him success in the early part of his investing life and the more activist value investing that has brought him success in recent years.

Value Screens

The Graham approach to value investing is a screening approach, where investors adhere to strict screens (like the ones described earlier in the chapter) and pick stocks that pass those screens. Since the data needed to screen stocks is widely available today, the key to success with this strategy seems to be picking the right screens. In this section, we will consider a number of screens used to pick value stocks and the efficacy of these screens.

Book Value Multiples

The book value of equity measures what accountants consider to be the value of equity in a company. The market value of equity is what investors attach as a value to the same equity. Investors have used the relationship between price and book value in a number of investment strategies, ranging from the simple to the sophisticated. In this section, we will begin by looking at a number of these strategies and the empirical evidence on their success.

Buy low price to book value companies

Some investors argue that stocks that trade at low price-book value ratios are under valued and there are several studies that seem to back this strategy. Rosenberg, Reid and Lanstein looked at stock returns in the United States between 1973 and 1984 found that the strategy of picking stocks with high book/price ratios (low price-book values) would have yielded an excess return of about 4.5% a year. In another study of stock returns between 1963 and 1990[3], firms were classified on the basis of book-to-price ratios into twelve portfolios, and firms in the lowest book-to-price (higher P/BV) class earned an average annual return of 3.7% a year , while firms in the highest book-to-price (lowest P/BV) class earned an average annual return of 24.31% for the 1963-90 period. We updated these studies to consider how well a strategy of buying low price to book value stocks would have done from 1991 to 2001 and compared these returns to returns in earlier time periods. The results are summarized in figure 8.2.

Figure 8.2: PBV Classes and Returns - 1927-2001

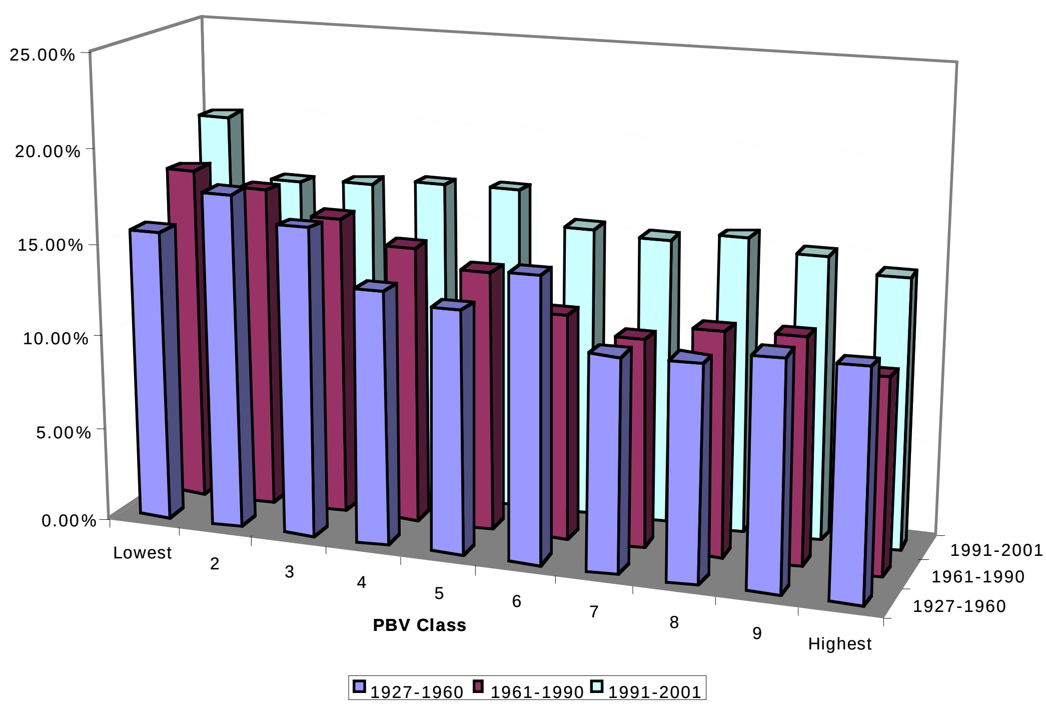

Source: Raw data from French

---

[3] This study was done by Fama and French in 1992, in the course of an examination of the effectiveness of different risk and return models in finance. They found that price to book explained more of the variation across stock returns than any other fundamental variable, including market capitalization.

The lowest price to book value stocks continued to earn higher annual returns than the high price to book value stocks during the 1990s.

These findings are not unique to the United States. A 1991 study found that the book-to-market ratio had a strong role in explaining the cross-section of average returns on Japanese stocks[4]. Another study extended the analysis of price-book value ratios across other international markets, and found that stocks with low price-book value ratios earned excess returns in every market that analyzed, between 1981 and 1992[5]. The annualized estimates of the return differential earned by stocks with low price-book value ratios, over the market index, were as follows in each of the markets studied:

Country	*Added Return to low P/BV portfolio*
France	3.26%
Germany	1.39%
Switzerland	1.17%
U.K	1.09%
Japan	3.43%
U.S.	1.06%
Europe	1.30%
Global	1.88%

Thus, a strategy of buying low price to book value stocks seems to hold out much promise. Why don't more investors use it then, you might ask? We will consider some of the possible problems with this strategy in the next section and screens that can be added on to remove these problems.

Stocks with lowest price to book ratios: Take a look at the 50 stocks with the lowest price to book value ratios in the U.S.

What can go wrong?

Stocks with low price to book value ratios earn excess returns relative to high price to book stocks, if we use conventional measures of risk and return, such as betas. But, as noted in earlier chapters, these conventional measures of risk are imperfect and incomplete. Low price-book value ratios may operate as a measure of risk, since firms with prices well below book value are more likely to be in financial trouble and go out of business. Investors

4 Chan, Hamao and Lakonishok (1991) did this study and concluded that low price to book value stocks inJ Japan earned a considerable premum over high price to book value stocks.

5 Capaul, Rowley and Sharpe (1993) did this study on international markets.

therefore have to evaluate whether the additional returns made by such firms justifies the additional risk taken on by investing in them.

The other limitation of a strategy of buying low price to book value stocks is that the low book value multiples may be well deserved if companies earn and are expected to continue earning low returns on equity. In fact, we considered the relationship between price to book value ratios and returns on equity in chapter 5. For a stable growth firm, for instance, the price to book value ratio can be written as follows:

$$\text{Price/Book} = \frac{(\text{Return on Equity - Expected Growth Rate})}{(\text{Return on Equity - Cost of Equity})}$$

Stocks with low returns on equity should trade a low price to book value ratios. In fact, a firm that is expected to earn a return on equity that is less than its cost of equity in the long term should trade at a discount on book value. In summary, then, as an investor you would want stocks with low price to book ratios that also had reasonable (if not high) returns on equity and limited exposure to risk.

Composite Screens

If low price to book value ratios may yield riskier stocks than average or stocks that have lower returns on equity, a more discerning strategy would require us to find mismatches – stocks with low price to book ratios, low default risk and high returns on equity. If we used debt ratios as a proxy for default risk and the accounting return on equity in the last year as the proxy for the returns that will be earned on equity in the future, we would expect companies with low price to book value ratios, low default risk and high return on equity to be under valued. This proposition was partially tested by screening all NYSE stocks from 1981 to 1990, on the basis of price-book value ratios and returns on equity at the end of each year and creating two portfolios - an 'undervalued' portfolio with low price-book value ratios (in bottom quartile of all stocks) and high returns on equity (in top quartile of all stocks) and an overvalued portfolio with high price-book value ratios (in top quartile of all stocks) and low returns on equity (in bottom quartile of all stocks)- each year, and then estimating excess returns on each portfolio in the following year. Table 8.1 summarizes returns on these two portfolios for each year from 1982 to 1991.

> *Stocks with low price to book and high returns on equity:* Take a look at the stocks that are in the bottom quartile for price to book and the top for ROE.

Table 8.1: Returns on Mismatched Portfolios: Price to Book and ROE

Year	*Undervalued Portfolio*	*Overvalued Portfolio*	*S & P 500*

1982	37.64%	14.64%	40.35%
1983	34.89%	3.07%	0.68%
1984	20.52%	-28.82%	15.43%
1985	46.55%	30.22%	30.97%
1986	33.61%	0.60%	24.44%
1987	-8.80%	-0.56%	-2.69%
1988	23.52%	7.21%	9.67%
1989	37.50%	16.55%	18.11%
1990	-26.71%	-10.98%	6.18%
1991	74.22%	28.76%	31.74%
1982-91	*25.60%*	*10.61%*	*17.49%*

The undervalued portfolios significantly outperformed the overvalued portfolios in eight out of ten years, earning an average of 14.99% more per year between 1982 and 1991, and also had an average return significantly higher than the S&P 500. While we did not adjust for default risk in this test, you could easily add it as a third variable in the screening process.

Market Value to Replacement Cost – Tobin's Q

Tobin's Q provides an alternative to the price-book value ratio, by relating the market value of the firm to the replacement value of the assets in place. When inflation has pushed up the price of the assets, or where technology has reduced the price of the assets, this measure may provide a better measure of undervaluation.

Tobin's Q = Market value of assets / Replacement Value of Assets in place

While this measure has some advantages in theory, it does have practical problems. The first is that the replacement value of some assets may be difficult to estimate, largely because they are so specific to each firm. The second is that, even where replacement values are available, substantially more information is needed to construct this measure than the traditional price-book value ratio. In practice, analysts often use short cuts to arrive at Tobin's Q, using book value of assets as a proxy for replacement value. In these cases, the only distinction between this measure and the price/book value ratio is that this ratio is stated in terms of the entire firm (rather than just the equity).

The value obtained from Tobin's Q is determined by two variables - the market value of the firm and the replacement cost of assets in place. In inflationary times, where the cost of replacing assets increases significantly, Tobin's Q will generally be lower than the unadjusted price-book value ratio. Conversely, if the cost of replacing assets declines much faster than the book value (computers might be a good example), Tobin's Q will generally be higher than the unadjusted price-book value ratio.

Many studies, in recent years, have suggested that a low Tobin's Q is indicative of an undervalued or a poorly managed firm, which is more likely to be taken over. One study concludes that firms with low Tobin's Q are more likely to be taken over for purposes of restructuring and increasing value.[6] They also find that shareholders of high q bidders gain significantly more from successful tender offers than shareholders of low q bidders.

Earnings Multiples

Investors have long argued that stocks with low price earnings ratios are more likely to be undervalued and earn excess returns. In fact, it was the first of Ben Graham’s ten screens for undervalued stocks. In this section, we will examine whether it stands up to the promises made by its proponents.

Empirical Evidence on Low PE Stocks

Studies that have looked at the relationship between PE ratios and excess returns have consistently found that stocks with low PE ratios earn significantly higher returns than stocks with high PE ratios over long time horizons. Figure 8.3 summarizes annual returns by PE ratio classes for stocks from 1952 to 2001. The classes were created based upon PE ratios at the beginning of each year and returns were measured during the course of the year.

Stocks with highest low PE ratios: Take a look at the 50 stocks with the lowest PE ratios in the U.S.

[6] Lang, Stulz and Walkling (1989) looked at the relationship between Tobin’s Q and acquisitions.

Figure 8.3: Returns on PE Ratio Classes - 1952 - 2001

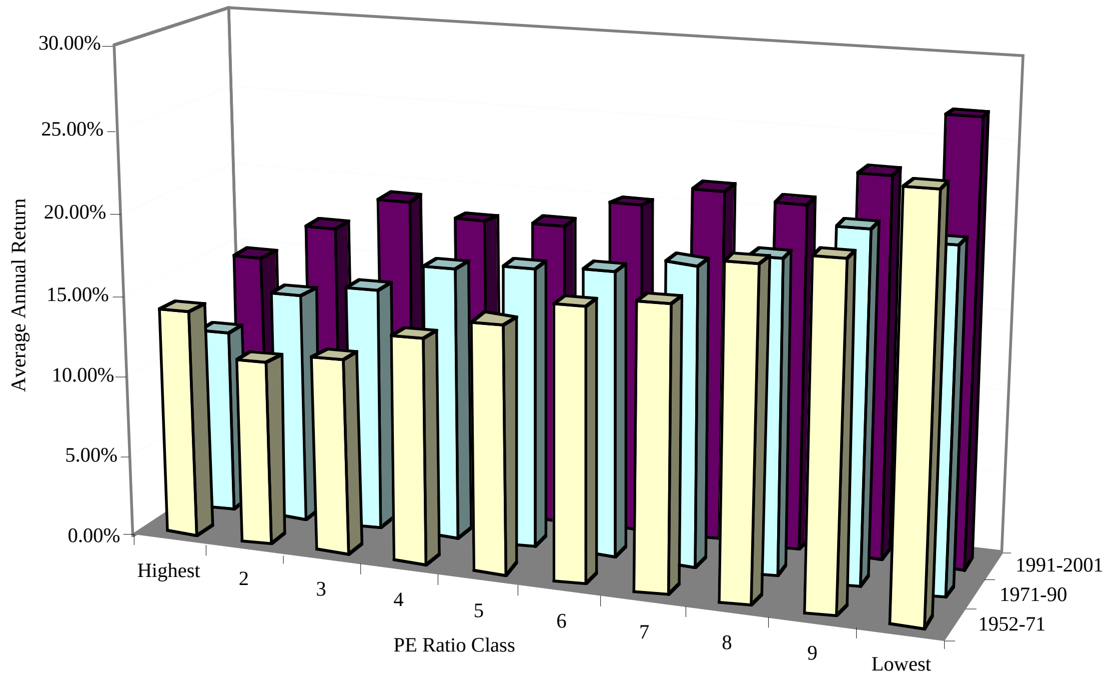

Source: Raw data from French

Firms in the lowest PE ratio class earned 10% more each year than the stocks in the highest PE class between 1952 and 1971, about 9% more each year between 1971 and 1990 and about 12% more each year between 1991 and 2001.

The excess returns earned by low PE ratio stocks also persist in other international markets. Table 8.2 summarizes the results of studies looking at this phenomenon in markets outside the United States.

Table 8.2: Excess Returns on Low P/E Ratio Stocks by Country: 1989-1994

Country	*Annual Premium earned by lowest P/E Stocks (bottom quintile)*
Australia	3.03%
France	6.40%
Germany	1.06%
Hong Kong	6.60%
Italy	14.16%
Japan	7.30%
Switzerland	9.02%
U.K.	2.40%

Annual premium: Premium earned over an index of equally weighted stocks in that market between January 1, 1989 and December 31, 1994. These numbers were obtained from a Merrill Lynch Survey of Proprietary Indices.

Thus, the results seem to hold up as we go across time and markets, notwithstanding the fact the findings have been widely disseminated for more than 20 years.

What can go wrong?

Given the types of returns that low PE ratio stocks earn, should we rush out and buy such stocks? While such a portfolio may include a number of under valued companies, it may also contain other less desirable companies.

a. Companies with high-risk earnings: The excess returns earned by low price earnings ratio stocks can be explained using a variation of the argument used for small stocks, i.e., that the risk of low PE ratios stocks is understated in the CAPM. It is entirely possible that a portfolio of low PE stocks will include stocks where there is a great deal of uncertainty about future operating earnings. A related explanation, especially in the aftermath of the accounting scandals of recent years, is that accounting earnings is susceptible to manipulation. If earnings are high not because of a firm's operating efficiency but because of one-time items such as gains from divestiture or questionable items such as income from pension funds, you may discount these earnings more (leading to a lower PE ratio).

b. Tax Costs: A second possible explanation that can be given for this phenomenon, which is consistent with an efficient market, is that low PE ratio stocks generally have large dividend yields, which would have created a larger tax burden for investors since dividends were taxed at higher rates during much of this period.

c. Low Growth: A third possibility is that the price earnings ratio is low because the market expects future growth in earnings to be low or even negative. Many low PE ratio companies are in mature businesses where the potential for growth is minimal. As an investor, therefore, you have to consider whether the trade off of a lower PE ratio for lower growth works in your favor.

Finally, many of the issues we raised about how accountants measure earnings will also be issues when you use PE ratios. For instance, the fact that research and development is expensed at technology firms rather than capitalized may bias their earnings down (and their PE ratios upwards).

Modified Earnings Multiples

The price earnings ratio is computed by dividing the current price by the current earnings per share. The latter is both volatile and subject to measurement error. Are there ways in which we can modify the ratio to make it a better tool for investment analysis? There are several variations that have been suggested by analysts:

1. Price to Normalized Earnings: When your primary concern is volatility in earnings, as is often the case with cyclical and commodity companies, you can average earnings across a cycle (an economic cycle for a cyclical firm or a price cycle for a commodity firm) and use it as a measure of normalized earnings. Only firms that have low price to normalized earnings would be considered cheap.
2. Price to Adjusted Earnings: When your concern is with accounting standards and measurement issues, you may need to restate earnings to reflect your concerns. For instance, Standard and Poor's recently came up with a measure of operating earnings for companies where they adjust the earnings for the option grants to management and remove earnings from pension funds.
3. Price to Cash Earnings: When you have non-cash items (such as depreciation and amortization) significantly affecting measured earnings, you could argue that looking at the price as a multiple of cash earnings may give you a better measure of value. In the simplest form, you add back non-cash charges to earnings -

   Price/ Cash Earnings = Price / (Earnings + Depreciation & Amortization)

   In its' more complex forms, you adjust for changes in non-cash working capital to convert accrual earnings to cash earnings.

Once you have the modified earnings multiples for firms, you can screen to find the stocks with the lowest multiples of earnings. There are two final tests that you need to run on this list to ensure that your portfolio is not composed of low growth, high-risk companies:

a. Check for risk: You may want to introduce a screen for risk, using either market variables (such as standard deviation in stock prices) or accounting variables (such as debt to equity ratios), and only invest in stocks with below-average risk.

b. Assess growth: While it would be unrealistic to expect low PE stocks to have high growth, you can still apply minimal screens for growth. For instance, you may want to eliminate firms where earnings have been declining for the last few years (with no end in sight) or are growing at rates lower than their sectors.

Enterprise Value to EBITDA Multiples

The earnings per share of a firm reflect not just the earnings from operations of a firm but all other income as well. Thus, a firm with substantial holdings of cash and marketable securities may generate enough income on these investments to push up earnings. In addition, earnings per share and equity multiples are affected by how much debt a firm has and what its interest expenses are. These concerns, in conjunction with the volatility induced in earnings by non-cash expenses (such as depreciation) and varying tax rates has led some investors to seek a more stable, cash-based measure of pre-debt earnings.

One measure that has acquired a following is called the enterprise value to EBITDA multiple, and is defined as follows:

Enterprise Value to EBITDA =

$$\frac{\text{(Market Value of Equity + Market Value of Debt - Cash \& Marketable Securities)}}{\text{Earnings before interest, taxes, depreciation and amortization}}$$

Why, you might wonder, do we add back debt and subtract out cash? Since EBITDA is before interest expenses, you would be remiss if you did not add back debt. Analysts who look at Price/EBITDA will conclude, for instance, that highly levered firms are cheap. Since we do not count the income from the cash and marketable securities in EBITDA, we net it out of the numerator as well.

The sectors where this multiple makes the most sense tend to be heavy infrastructure businesses – steel, telecommunications and cable are good examples. In these sectors, you can screen for stocks with low enterprise value to EBITDA. As a note of caution, though, in many cases firms that look cheap on an enterprise value to EBITDA basis often have huge reinvestment needs – capital expenditures eat up much of the EBITDA – and poor returns on capital. Thus, we would recommend adding two more screens when you use this multiple – low reinvestment needs and high return on capital.

Revenue Multiples

As investors have become more wary about trusting accounting earnings, an increasing number have started moving up the income statement looking for numbers that are less susceptible to accounting decisions. Not surprisingly, many have ended up screening for stocks that trade at low multiples of revenues. But how well have revenue multiples worked at picking under valued stocks? In this section, we will begin by looking at that evidence and then consider some of the limitations of this strategy.

Empirical Evidence on Price to Sales Ratios

There is far less empirical evidence, either for or against, on price to sales ratios than there is on price earnings or price to book value ratios. In one of the few direct tests of the price-sales ratio, Senchack and Martin in 1987 compared the performance of low price-sales ratio portfolios with low price-earnings ratio portfolios, and concluded that the low price-sales ratio portfolio outperformed the market but not the low price-earnings ratio portfolio. They also found that the low price-earnings ratio strategy earned more consistent returns than a low price-sales ratio strategy, and that a low price-sales ratio strategy was more biased towards picking smaller firms. In 1988, Jacobs and Levy tested the value of

Stocks with lowest price to sales ratios: Take a look at the 50 stocks with the lowest low price-sales ratios (standardized by the price-sales ratio of the industries in which the firms operated) as part of a general effort to disentangle the forces influencing equity returns. They concluded that low price-sales ratios, by themselves, yielded an excess return of about 2% a year between 1978 and 1986. Even when other factors were thrown into the analysis[7], the price-sales ratios remained a significant factor in explaining excess returns (together with price-earnings ratio and size).

We considered how a portfolio of low price to sales ratios would have done relative to a portfolio of high price to sales ratios from 1991 to 2001. We found that the returns on the low price to sales ratio portfolio were no greater than the returns earned on a high price to sales ratio portfolio over this decade, reflecting the surge of new economy companies that entered the market during the period with huge price to sales ratios.

What can go wrong?

While firms with low price to sales ratios may deliver excess returns over long periods, it should be noted, as with low price to book and price earnings ratios, that there are firms that trade at low price to sales ratios that deserve to trade at those values. In addition to risk being the culprit again – higher risk companies should have lower price to sales ratios – there are other possible explanations.

> *Stocks with highest low price to sales and high margins*: Take a look at the stocks that are in the bottom quartile for price to sales and the top quartile for margin.

1. High Leverage: One of the problems with using price to sales ratios is that you are dividing the market value of equity by the revenues of the firm. When a firm has borrowed substantial amounts, it is entirely possible that it's market value will trade at a low multiple of revenues. If you pick stocks with low price to sales ratios, you may very well end up with a portfolio of the most highly levered firms in each sector.
2. Low Margins: Firms that operate in businesses with little pricing power and poor profit margins will trade at low multiples of revenues. The reason is intuitive. Your value ultimately comes not from your capacity to generate revenues but from the earnings that you have on those revenues.

The simplest way to deal with the first problem is to redefine the revenue multiple. If you use enterprise value (which adds debt to the numerator and subtracts out cash) instead of

---

[7]Jacobs and Levy considered 25 different anomaly measures, based upon past studies, including size, PE, P/BV, earnings momentum measures, relative strength and neglect.

market value of equity in the numerator, you will remove the bias towards highly levered firms.

Composite Revenue Multiples

The significance of profit margins in explaining price-sales ratios suggests that screening on the basis of both price-sales ratios and profit margins should be more successful at identifying undervalued securities. To test this proposition, the stocks on the New York Stock Exchange were screened on the basis of price-sales ratios and profit margins to create 'undervalued' portfolios (price-sales ratios in the lowest quartile and profit margins in the highest quartile) and 'overvalued' portfolios (price-sales ratios in the highest quartile and profit margins in the lowest quartile) at the end of each year from 1981 to 1990. The returns on these portfolios in the following year are summarized in the Table 8.3:

Table 8.3: Returns on Mismatched Portfolios – PS and Net Margins

Year	*Undervalued Portfolio*	*Overvalued Portfolio*	*S & P 500*
1982	50.34%	17.72%	40.35%
1983	31.04%	6.18%	0.68%
1984	12.33%	-25.81%	15.43%
1985	53.75%	28.21%	30.97%
1986	27.54%	3.48%	24.44%
1987	-2.28%	8.63%	-2.69%
1988	24.96%	16.24%	9.67%
1989	16.64%	17.00%	18.11%
1990	-30.35%	-17.46%	6.18%
1991	91.20%	55.13%	31.74%
1982-91	23.76%	15.48%	17.49%

During the period, the undervalued portfolios outperformed the overvalued portfolios in six out of the ten years, earning an average of 8.28% more per year, and averaged a significantly higher return than the S&P 500.

Dividend Yields

While PE ratios, price to book ratios and price to sales ratios might be the most widely used value screens, there are some investors who view the dividend yield as the only secure measure of returns. Earnings, they argue, are not only illusory but they are out of reach for most investor in stocks since a significant portion may get reinvested. Following up on this logic, stocks with high dividend yields should be better investments that stocks with low dividend yields.

Does this approach yield results? Between 1952 and 2001, for instance, stocks with high dividend yields earned higher annual returns than stocks with low dividend yields, but the relationship is neither as strong or as consistent as the results obtained from the PE ratio or the PBV ratio screens. Figure 8.4 summarizes returns earned by dividend yield class from 1952 to 2001, broken down by sub-periods:

Figure 8.4: Returns on Dividend Yield Classes - 1952 - 2001

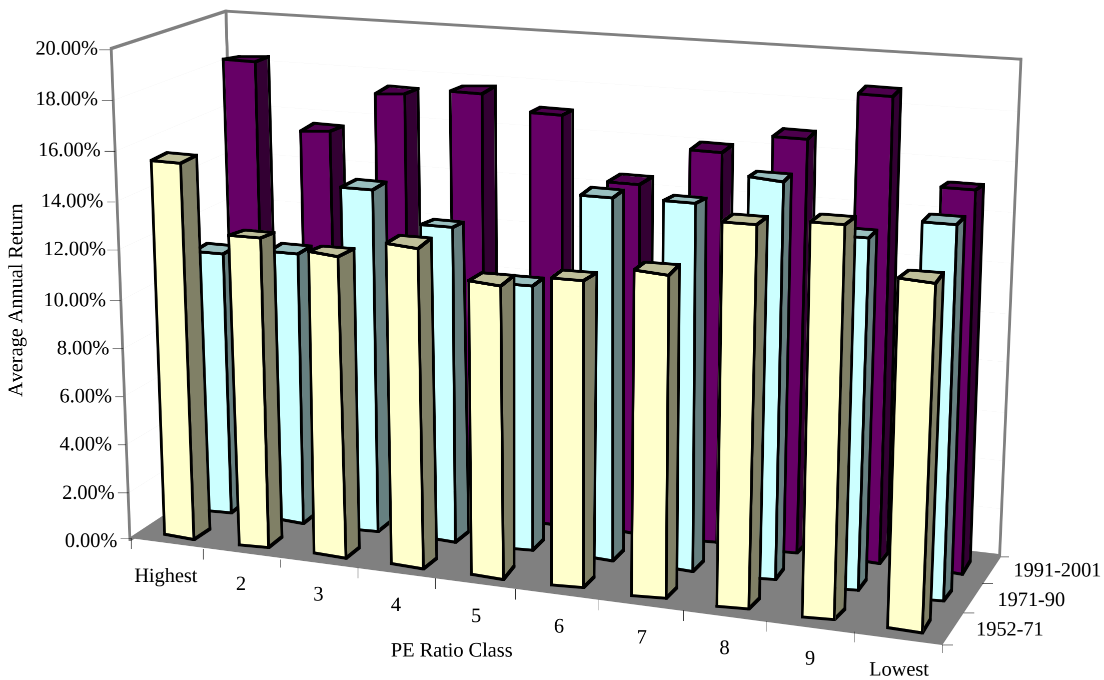

Source: Raw data from French

The highest dividend yield stocks earned higher returns than lower dividend yield stocks in the 1952-71 and the 1991-2001 time periods, but the stocks with the lowest returns are the stocks with average dividends. In the 1971-90 time period, stocks with lower dividend yields outperformed stocks with higher dividend yields.

An extreme version of this portfolio is the strategy of investing in the "Dow Dogs", the ten stocks with the highest dividend yields in the Dow 30. Proponents of this strategy claim that they generate excess returns from it, but they compare the returns to what you would have made on the Dow 30 and the S&P 500 and do not adequately adjust for risk. A portfolio with only 10 stocks in it is likely to have a substantial amount of firm-specific risk. An study by McQueen, Shields

Stocks with highest dividend yields: Take a look at the 50 stocks with the highest dividend yields in the United States, and Thorley in 1997 examined this strategy and concluded that while the raw returns from buying the top dividend paying stocks is higher than the rest of the index, adjusting for risk and taxes eliminates all of the excess return. A study by Hirschey in 2000 also indicates that there are no excess returns from this strategy after you adjust for risk.

There are three final considerations in a high-dividend strategy. The first is that you will have a much greater tax cost on this strategy, since dividends are taxed at a higher rate than capital gains. The second is that some stocks with high dividend yields currently may be paying much more in dividends than they can afford. It is only a matter of time, then, before the dividends are cut. The third is that any stock that pays a substantial portion of its earnings as dividends is reinvesting less and can therefore expect to grow at a much lower rate.

Determinants of Success

If all we have to do to earn excess returns is invest in stocks that trade at low multiples of earnings, book value or revenues, shouldn't more investors employ these screens to pick their portfolio? And assuming that they do, should they not beat the market by a healthy amount?

To answer the first question, there are a large number of portfolio managers and individual investors who employ either the screens we have referred to in this section or variants of these screens to pick stocks. Unfortunately, their performance does not seem to match up to the returns that we see earned on the hypothetical portfolios. Why might that be? We can think of several reasons.

- *Time Horizon:* All the studies quoted above look at returns over time horizons of five years or greater. In fact, low price-book value stocks have underperformed high price-book value stocks over shorter time periods. The same can be said about PE ratios and price to sales ratios.
- *Dueling Screens:* If one screen earns you excess returns, three should do even better seems to be the attitude of some investors who proceed to multiply the screens they use. They are assisted in this process by the easy access to both data and screening technology. There are web sites (many of which are free) that allow you to screen stocks (at least in the United States) using multiple criteria.[8] The problem, though, is that the use of one screen seems to undercut the effectiveness of others, leading to worse rather than better portfolios.

---

[8] Stockscreener.com. run by Hoover, is one example. You can screen all listed stocks in the United States using multiple criteria, including all of the criteria discussed in this chapter.

- *Absence of Diversification:* In their enthusiasm for screens, investors sometimes forget the first principles of diversification. For instance, it is not uncommon to see stocks from one sector disproportionately represented in portfolios created using screens. A screen from low PE stocks may deliver a portfolio of banks and utilities, whereas a screen of low price to book ratios and high returns on equity may deliver stocks from a sector with high infrastructure investments that has had bad sector-specific news come out about it. In 2001, for instance, many telecom stocks traded at a discount on their book value.
- *Taxes and Transactions costs:* As in any investment strategy, taxes and transactions costs can take a bite out of returns, although the effect should become smaller as your time horizon lengthens. Some screens, though, can increase the effect of taxes and transactions costs. For instance, screening for stocks with high dividends and low PE ratios will yield a portfolio that has much higher tax liabilities (because of the dividends).
- *Success and Imitation:* In some ways, the worst thing that can occur to a screen (at least from the viewpoint of investors using the screen) is that its success is publicized and that a large number of investors begin using that same screen at the same time. In the process of creating portfolios of the stocks they perceive to be undervalued, they may very well eliminate the excess returns that drew them to the screen in the first place.

To be a successful screener, you would need to be able to avoid or manage these problems. In particular, you need to have a long time horizon, pick your combination of screens well, and ensure that you are reasonably diversified. If a screen succeeds, you will probably need to revisit it at regular intervals to ensure that market learning has not reduced the efficacy of the screen.

The Contrarian Value Investor

The second strand of value investing that we will examine is contrarian value investing. In this manifestation of value investing, you begin with the belief that stocks that are beaten down because of the perception that they are poor investments (because of poor investments, default risk or bad management) tend to get punished too much by markets just as stocks that are viewed as good investments get pushed up too much. Within contrarian investing, we would include several strategies ranging from relatively unsophisticated ones like buying the biggest losers in the market in the prior period to vulture and distressed security investing, where you use sophisticated quantitative techniques to highlight securities (both stocks and bonds) issued by troubled firms that may be undervalued.

Basis for contrarian investing

Do markets overreact to new information and systematically over price stocks when the news is good and under price stocks when the news is bad? There is some evidence that suggests that markets do overreact to both good and bad news, especially in the long term, and that stocks that have done exceptionally well or badly in a period tend to reverse course in the following period, but only if the period is defined in terms of years rather than weeks or months.

Strategies and Evidence

While contrarian investing takes many forms, we will consider three strategies in this section. We will begin with the simple strategy of buying stocks that have gone down the most over the previous period, move on to a slightly more sophisticated process of playing the expectations game, buying stocks where expectations have been set too low and selling stocks where expectations are too high and end the section by looking at a strategy of investing in securities issued by firms in significant operating and financial trouble.

a. Buying the Losers

In chapter 7, we presented evidence that stocks reverse themselves over long periods in the form of negative serial correlation – i.e. stocks that have gone up the most over the last 5 years are more likely to go down over the next 5 years. Conversely, stocks that have gone down the most over the last 5 years are more likely to go up. In this section, we will consider a strategy of buying the latter and selling or avoiding the former.

The Evidence

How would a strategy of buying the stocks that have gone down the most over the last few years perform? To isolate the effect of price reversals on the extreme portfolios, DeBondt and Thaler constructed a winner portfolio of 35 stocks, which had gone up the most over the prior year, and a loser portfolio of 35 stocks, which had gone down the most over the prior year, each year from 1933 to 1978, They examined returns on these portfolios for the sixty months following the creation of the portfolio. Figure 8.5 graphs the returns on both the loser and winner portfolios:

Figure 8.5: Cumulative Abnormal Returns - Winners versus Losers

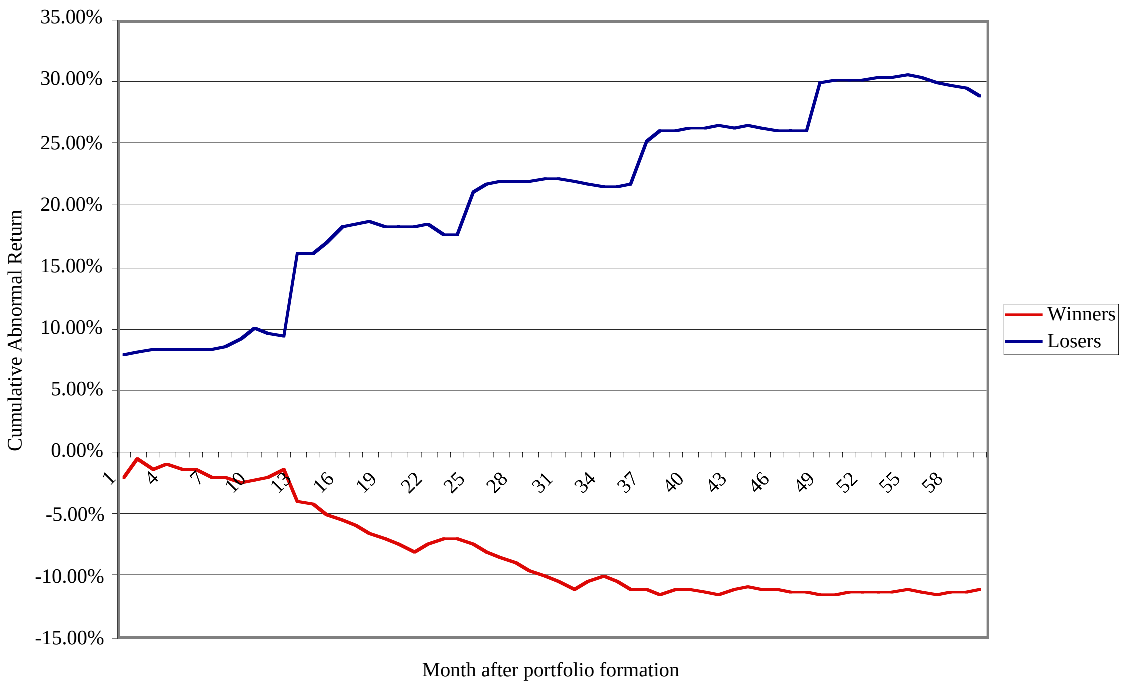

Source: DeBondt and Thaler

This analysis suggests that an investor who bought the 35 biggest losers over the previous year and held for five years would have generated a cumulative abnormal return of approximately 30% over the market and about 40% relative to an investor who bought the winner portfolio.

This evidence is consistent with market overreaction and suggests that a simple strategy of buying stocks that have gone down the most over the last year or years may yield excess returns over the long term. Since the strategy relies entirely on past prices, you could argue that this strategy shares more with charting – consider it a long term contrarian indicator - than it does with value investing.

Loser Stocks: Take a look at the 50 stocks that went down the most over the last year.

Caveats

There are many, academics as well as practitioners, who suggest that these findings may be interesting but that they overstate potential returns on 'loser' portfolios for several reasons:

- There is evidence that loser portfolios are more likely to contain low priced stocks (selling for less than $5), which generate higher transactions costs and are also more likely to offer heavily skewed returns, i.e., the excess returns come from a few stocks making phenomenal returns rather than from consistent performance.

- Studies also seem to find loser portfolios created every December earn significantly higher returns than portfolios created every June. This suggests an interaction between this strategy and tax loss selling by investors. Since stocks that have gone down the most are likely to be sold towards the end of each tax year (which ends in December for most individuals) by investors, their prices may be pushed down by the tax loss selling.
- There seems to be a size effect when it comes to the differential returns. When you do not control for firm size, the loser stocks outperform the winner stocks, but when you match losers and winners of comparable market value, the only month in which the loser stocks outperform the winner stocks is January.[9]
- The final point to be made relates to time horizon. As we noted in the last chapter, while there may be evidence of price reversals in long periods (3 to 5 years), there is evidence of price momentum – losing stocks are more likely to keep losing and winning stocks to keep winning – if you consider shorter periods (six months to a year). An earlier study that we referenced, by Jegadeesh and Titman tracked the difference between winner and loser portfolios[10] by the number of months that you held the portfolios. Their findings are summarized in Figure 8.5:

---

[9] See “Size, Seasonality and Stock Market Overreaction”, by Zarowin (1990)

[10] The definition of winner and loser portfolios is slightly different in this study. The portfolios were created, based upon returns over

Figure 8.6: Differential Returns - Winner versus Loser Portfolios

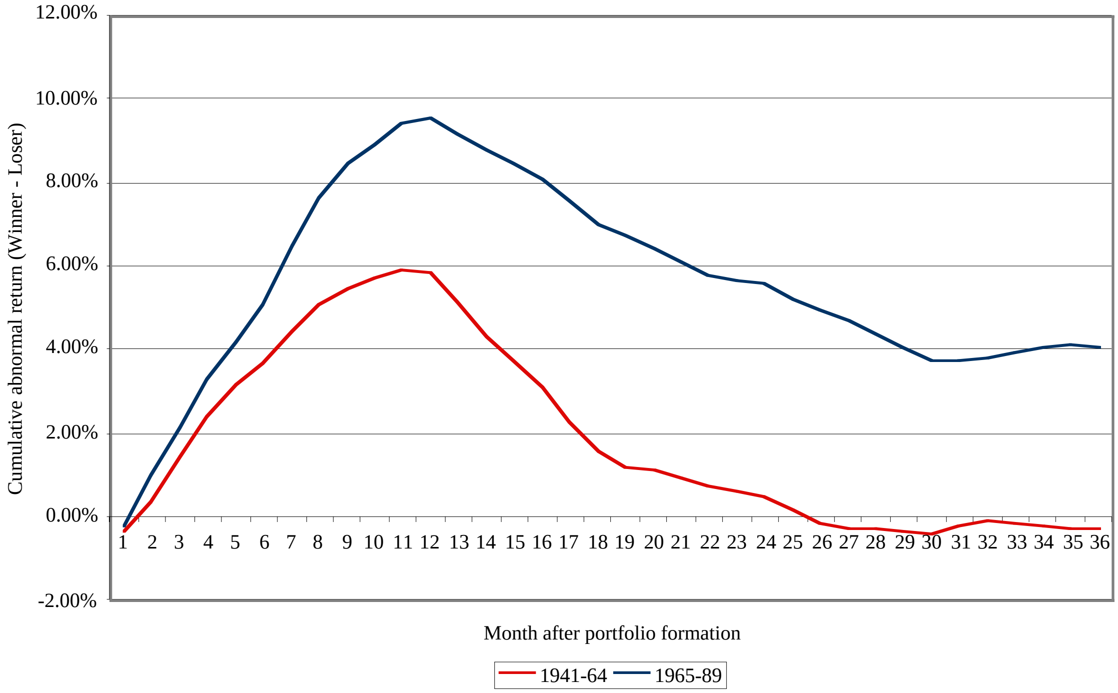

Source: Jegadeesh and Titman

There are two interesting findings in this graph. The first is that the winner portfolio actually outperforms the loser portfolio in the first 12 months. The second is that while loser stocks start gaining ground on winning stocks after 12 monthly, it took them 28 months in the 1941-64 time period to get ahead of them and the loser portfolio does not start outperforming the winner portfolio even with a 36-month time horizon in the 1965-89 time period. The payoff to buying losing companies may depend very heavily on whether you have to capacity to hold these stocks for long time periods.

b. Playing the Expectations Game

A more sophisticated version of contrarian investing is to play the expectations game. If you are right about markets overreacting to recent events, expectations will be set too high for stocks that have been performing well and too low for stocks that have been doing badly. If you can isolate these companies, you can buy the latter and sell the former. In this section, we will consider a couple of ways in which you can invest on expectations.

Bad companies can be good investments

Any investment strategy that is based upon buying well-run, good companies and expecting the growth in earnings in these companies to carry prices higher is dangerous, since it ignores the possibility that the current price of the company already reflects the quality of the management and the firm. If the current price is right (and the market is paying a premium for quality), the biggest danger is that the firm loses its luster over time, and that the premium paid will dissipate. If the market is exaggerating the value of the firm, this strategy can lead to poor returns even if the firm delivers its expected growth. It is only when markets under estimate the value of firm quality that this strategy stands a chance of making excess returns.

There is some evidence that well managed companies do not always make good investments. Tom Peters, in his widely read book on excellent companies a few years ago, outlined some of the qualities that he felt separated excellent companies from the rest of the market. Without contesting his standards, a study went through the perverse exercise of finding companies that failed on each of the criteria for excellence – a group of unexcellent companies and contrasting them with a group of excellent companies. Table 8.4 below provides summary statistics for both groups: [11]

Table 8.4: Excellent versus Unexcellent Companies – Financial Comparison

	Excellent companies	*Unexcellent companies*
Growth in assets	10.74%	4.77%
Growth in equity	9.37%	3.91%
Return on Capital	10.65%	1.68%
Return on Equity	12.92%	-15.96%
Net Margin	6.40%	1.35%

The excellent companies clearly are in much better financial shape and are more profitable than the unexcellent companies, but are they better investments? Figure 8.7 contrasts the returns would have made on these companies versus the excellent ones.

[11] See "Excellence revisited" by Michelle Clayman, Financial Analysts Journal, May/June 1994, pg 61-66.

Figure 8.7: Excellent versus Unexcellent Companies

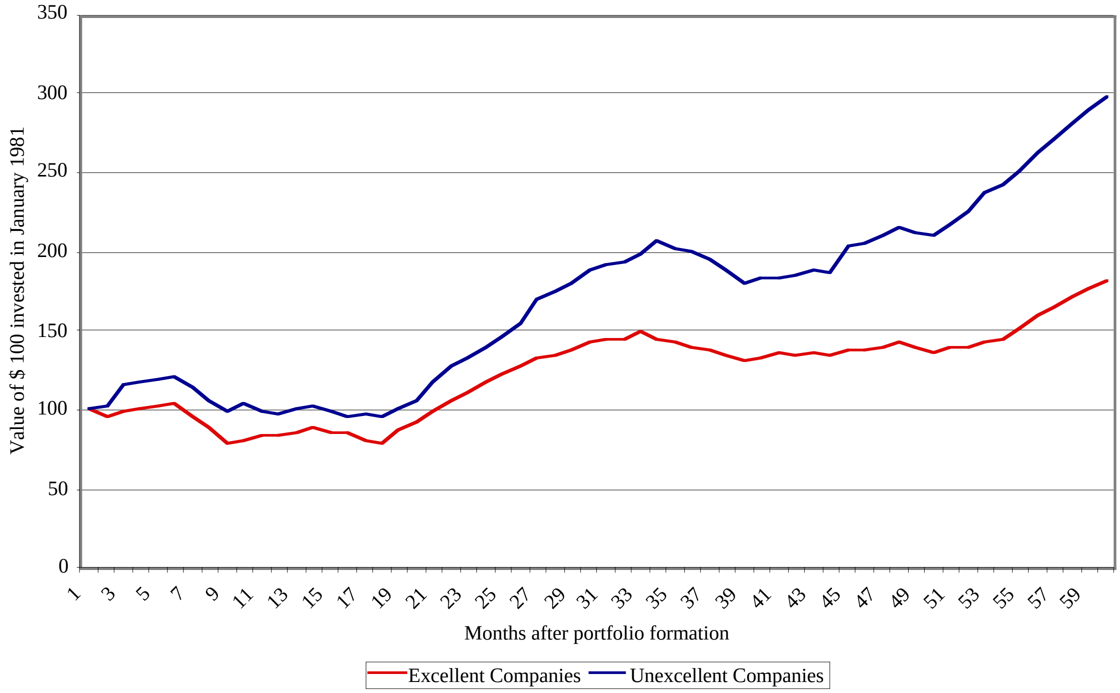

Source: Clayman

The excellent companies may be in better shape financially but the unexcellent companies would have been much better investments at least over the time period considered (1981-1985). An investment of $ 100 in unexcellent companies in 1981 would have grown to $ 298 by 1986, whereas $ 100 invested in excellent companies would have grown to only $ 182. While this study did not control for risk, it does present some evidence that good companies are not necessarily good investments, whereas bad companies can sometimes be excellent investments.

The second study used a more conventional measure of company quality. Standard and Poor's, the ratings agency, assigns quality ratings to stocks that resemble its bond ratings. Thus, an A rated stock, according to S&P, is a higher quality investment than a B+ rated stock, and the ratings are based upon financial measures (such as profitability ratios and financial leverage). Figure 8.8 summarizes the returns earned by stocks in different ratings classes, and as with the previous study, the lowest rated stocks had the highest returns and the highest rated stocks had the lowest returns.

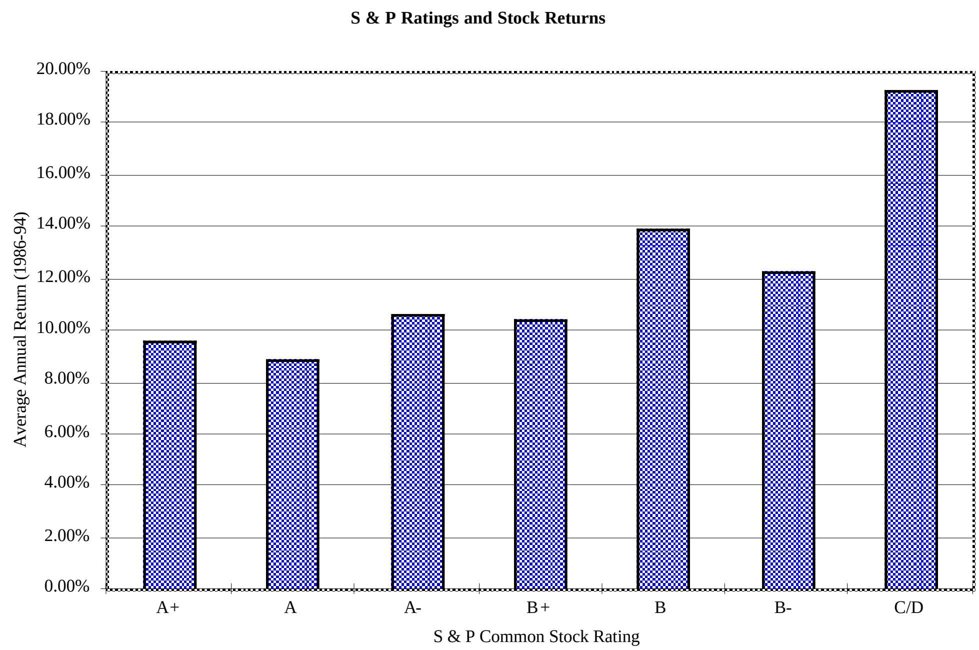

Again, the study is not definitive because it may well reflect the differences in risk across these companies, but it indicates that investors who bought the highest ranked stocks, expecting to earn higher returns, would have been sorely disappointed.

One version, perhaps an extreme one, of contrarian investing is vulture investing. In vulture investing, you buy the equity and bonds of companies that are in bankruptcy and bet either on a restructuring or a recovery. This is a high-risk strategy where your hope that a few big winners offset the many losers in your portfolio.

Caveats

As with the previous strategy of buying losers, a strategy of buying companies that rank low on financial criteria is likely to require a long time horizon and expose you to more risk, both from financial default and volatility. In addition, though, the following factors should be kept in mind while putting together a portfolio of “bad” companies.

The first is that not all companies that are poor performers are badly managed. Many are in sectors that are in long-term decline and have no turn-around in sight. It is entirely likely that these companies will continue to be poor performers in the future. Your odds of success are usually higher, if you buy a poorly performing company in a sector, where other companies are performing well. In other words, you are more likely to get the upside if there is potential for improvement.

Even if companies have potential for improvement in their sectors, part of the reason for the poor performance of the companies may be poor management. If the management of the company is entrenched, either because the managers hold a significant portion of the equity – at least the voting shares – or because of anti-takeover amendments in place, there may be little chance of improved performance in the future. You may have a better chance of succeeding at your portfolio, if you direct your investments to poorly managed firms, where there is a high (or at least reasonable) chance of removing incumbent management. You would, for instance, avoid poorly managed companies with unequal voting rights (voting and non-voting shares), substantial holdings by incumbent managers or anti-takeover amendments in place.

Finally, risk averse investors who wait for the absolute bottom before they will invest often fail at this strategy because timing it is just about impossible. You will have to accept the fact that bad companies will sometimes (or often) become worse before they become better, and that this may create some short-term damage to your portfolio.

Determinants of Success

The caveats presented in the section above suggest that success from buying losers or bad companies is not guaranteed and may prove illusive. In particular, you need the following –

a. Long Time Horizon: To succeed by buying these companies, you need to have the capacity to hold the stocks for several years. This is necessary not only because these stocks require long time periods to recover, but also to allow you to spread the high transactions costs associated with these strategies over more time. Note that having a long time horizon as a portfolio manager may not suffice if your clients can put pressure on you to liquidate holdings at earlier points. Consequently, you either need clients who think like you do and agree with you, or clients that have made enough money with you in the past that their greed overwhelms any trepidation they might have in your portfolio

b. Diversify: Since poor stock price performance is often precipitated or accompanied by operating and financial problems, it is very likely that quite a few of the companies in the loser portfolio will cease to exist. If you are not diversified, your overall returns will be extremely volatile as a result of a few stocks that lose all of their value. Consequently, you will need to spread your bets across a large number of stocks in a large number of sectors. One variation that may accomplish this is to buy the worst performing stock in each sector, rather than the worst performing stocks in the entire market.

c. Personal qualities: This strategy is not for investors who are easily swayed or stressed by bad news about their investments or by the views of others (analysts, market watchers and friends). Almost by definition, you will read little that is good about the firms in your portfolio. Instead, there will be bad news about potential default, management turmoil and failed strategies at the companies you own. In fact, there might be long periods after you buy the stock, where the price continues to go down further, as other investors give up on its future. Many investors who embark on this strategy find themselves bailing out of their investments early, unable to hold on to these stocks in the face of the drumbeat of negative information. In other words, you need both the self-confidence to stand your ground as others bail out and a stomach for short-term volatility (especially the downside variety) to succeed with this strategy.

Activist Value Investing

One of the more frustrating aspects of passive contrarian investing is that you, as an investor, do not control your destiny. Thus, you could invest in a poorly managed company, expecting management to change, but it may never happen, leaving you with an investment that wilts over time. In activist value investing, you acquire a large stake in an undervalued or poorly managed company, and then use your position as a large stockholder to push for changes that will release this value. In other words, you act as the catalyst for change, and enrich yourself in the process.

Strategies and Evidence

The strategies used by you as an activist value investor will be diverse, and will reflect why the firm is undervalued in the first place. With a conglomerate or multi-business firm that sells for less than the sum of its parts, you may push for divestitures or spin offs of the parts. When investing in a firm that is being far too conservative in its use of debt, you may push for a recapitalization (where the firm borrows money and buys back stock). Investing in a firm that could be worth more to another firm because of synergy, you may push for it to become the target of a hostile acquisition. When a company's value is weighted down because it is perceived as having too much cash, you may demand higher dividends or stock buybacks. In each of these scenarios, you may have to confront incumbent managers who are reluctant to make these changes. In fact, if your concerns are broadly about management competence, you may even push for a change in the top management of the firm.

Breaking up is hard to do

There are cases where large firms that operate in multiple businesses are penalized by the market, either because they are too complex to value or because of a perceived lack of efficiency that comes from being unfocused. In these cases, you could argue that pushing the firm to break up may create value for the component parts. In this section, we will first consider the overall evidence on how the market values multi-business firms, and then consider ways in which you may be able to release value at these firms.

The Conglomerate Discount

For the last few decades, strategists have gone back and forth on whether becoming a conglomerate creates or destroys value. In the 1960s and through much of the 1970s, the view was that conglomerates created valued, relative to their individual pieces, because you could pool of the strengths of the pieces to create a more powerful firm. A hidden subtext to many of these arguments was the premise that conglomerates were somehow less risky and more valuable than their individual components, because they were able to diversify away risk. Financial theorists pointed out that the fallacy in this argument by noting that individual investors could have accomplished the same diversification at far lower cost. Later, the argument shifted to one of superior management transferring its skills to poorly managed firms in different businesses, and creating often unnamed synergies.

Empiricists have approached this question from a different perspective. They have looked at the question of whether conglomerates trade at a premium or discount to their parts. To make this judgment, they valued the pieces of a conglomerate, using the typical multiple at which independent firms in the business trade at. Thus, you could break GE down into nine parts, and value each part based upon the enterprise value to EBITDA or PE ratio that other firms in the business trade at. You can then add up the values of the parts and compare it to the value of the conglomerate. In this comparison, the evidence[12] seems to indicate that conglomerates trade at significant discounts (ranging from 5 to 10%, depending upon the study) to their piecewise values. While one can contest the magnitude of these discounts on estimation grounds – it is difficult to estimate the true earnings of GE Capital, given allocations and other pooled costs – it is clear that some multi-business firms would be worth more as individual businesses.

So what can an activist investor who buys stock in such a company do to claim this surplus value? The most drastic step, in terms of separation from the parent company and existing management, is a divestiture of the individual pieces. There are less drastic alternatives as well, such as spin offs and split offs of independent businesses, that may

[12] See Berger, Philip G., and Eli Ofek, 1995, Diversification's effect on firm value, *Journal of Financial Economics* 37, 39–65. and Lang, Larry H.P., and René M. Stulz, 1994, Tobin's *q*, corporate diversification, and firm performance, *Journal of Political Economy* 102, 1248–1280.

accomplish the separation while preserving some of the benefits generated by having a linkage.

Divestitures

In a divestiture, a firm sells assets or a division to the highest bidder. On the sale, it receives cash that is either reinvested in new assets or returned to stockholders as dividends or stock buybacks. It is the most drastic of the actions described in this section, since the divested assets will belong to a new buyer and any connections with the parent company will be severed.

Process and Effect on Value

A divestiture can be initiated either by the divesting firm or by an interested buyer. In the first case, the divesting firm will offer assets for sale and invite potential bids. If the assets have substantial value, it will use the services of an investment banker in seeking out bidders. In the second case, the process starts with an interested buyer approaching the firm and offering to buy a division or assets. While this buyer cannot force the divestiture, it can elicit interest if it offers a high enough price. The final price will then be determined by negotiations between the two sides.

How does a divestiture affect a firm's value? To answer, you would need to compare the price received on the divestiture to the present value of the expected cash flows that the firm would have received from the divested assets. There are three possible scenarios:

1. If the divestiture value is equal to the present value of the expected cash flows, the divestitures will have no effect on the divesting firm's value.
2. If the divestiture value is greater than the present value of the expected cash flows, the value of the divesting firm will increase on the divestiture.
3. If the divestiture value is less than the present value of the expected cash flows, the value of the firm will decrease on the divestiture.

The divesting firm receives cash in return for the assets and can choose to retain the cash and invest it in marketable securities, invest the cash in other assets or new investments, or return the cash to stockholders in the form of dividends or stock buybacks. This action, in turn, can have a secondary effect on value.

Reasons for Divestitures

Why would a firm sell assets or a division? There are at least three reasons. The first is that the divested assets may have a higher value to the buyer of these assets. For assets to have a higher value, they have to either generate higher cash flows for the buyers or result in lower risk (leading to a lower discount rate). The higher cash flows can occur because the buyer is more efficient at utilizing the assets, or because the buyer finds synergies with its existing businesses. The lower discount rate may reflect the fact that the owners of the buying firm are more diversified that the owners of the firm selling the assets. In either case, both sides can gain from the divestiture and share in the increased value.

The second reason for divestitures is less value-driven and more a result of the immediate cash flow needs of the divesting firm. Firms that find themselves unable to meet their current operating or financial expenses may have to sell assets to raise cash. For instance, many leveraged acquisitions in the 1980s were followed by divestitures of assets. The cash generated from these divestitures was used to retire and service debt.

The third reason for divestitures relates to the assets not sold by the firm, rather than the divested assets. In some cases, a firm may find the cash flows and values of its core businesses affected by the fact that it has diversified into unrelated businesses. This lack of focus can be remedied by selling assets or businesses that are peripheral to the main business of a firm.

Market Reaction to Divestitures

A number of empirical questions are worth asking about divestitures. What types of firms are most likely to divest assets? What happens to the stock price when assets are divested? What effect do divestitures have on the operating performance of the divesting firm? Let us look at the evidence on each of these questions.

There are three scenarios in which firms divest assets. In the first, the firms are forced by the government to divest because of anti-trust laws. The second occurs when financially distressed firms need the cash to meet their financial obligations. In the third scenario, divestitures are part of a major restructuring effort, designed to return a firm to its core businesses. In some cases this process is initiated by the existing management, and in some cases by an acquirer. One study[13] looked at firms that were targets of hostile acquisitions, and noted that there were substantial asset divestitures in 60% of them; more than half the assets of the firms were divested in these cases. The divestitures were of units that were distinct from the rest of the firm's business and often had been acquired as part of an earlier diversification effort.

In a study in 1984, Linn and Rozeff examined the price reaction to announcements of divestitures by firms and reported an average excess return of 1.45% for 77 divestitures between 1977 and 1982. They also noted an interesting contrast between firms that announce the sale price and motive for the divestiture at the time of the divestiture, and those that do not: in general, markets react much more positively to the first group than to the second, as shown in Table 8.5.

---

[13] See Bhide (1989).

Table 8.5: Market Reaction to Divestiture Announcements

Price Announced	*Motive Announced*	
	Yes	No
Yes	3.92%	2.30%
No	0.70%	0.37%

It appears that financial markets view firms that are evasive about the reasons for and the proceeds from divestitures with skepticism. This finding was confirmed by Klein in 1986 , when she noted that that the excess returns are positive only for those divestitures where the price is announced at the same time as the divestiture. She extended the study and concluded that the magnitude of the excess return is a function of the size of the divestiture. For example, when the divestiture is less than 10% of the equity of the firm, there is no significant price effect, whereas if it exceeds 50%, the stock price increases by more than 8%.

Studies that have looked at the performance of parent firms after divestitures report improvements in a number of operating measures: operating margins and returns on capital increase, and stock prices tend to outperform the rest of the sector. In summary, firms that have lost focus often are most likely to diversify, markets respond positively to these divestitures if information is provided at the time of the divestiture and operating performance tends to improve after divestitures.

Spin Offs, Split Offs and Split Ups

In a spin off, a firm separates out assets or a division and creates new shares with claims on this portion of the business. Existing stockholders in the firm receive these shares in proportion to their original holdings. They can choose to retain these shares or sell them in the market. In a split up, which can be considered an expanded version of a spin off, the firm splits into different business lines, distributes shares in these business lines to the original stockholders in proportion to their original ownership in the firm, and then ceases to exist. A split off is similar to a spin off, insofar as it creates new shares in the undervalued business line. In this case, however, the existing stockholders are given the option to exchange their parent company stock for these new shares, which changes the proportional ownership in the new structure.

Process and Follow-up

Spin offs, split offs and split ups require far more procedural steps than a typical divestiture. Miles and Woolridge (1999) lay out the following steps in a typical spin off; they are similar for a split off or split up.

The process begins when the firm announces its intention to spin off a subsidiary or division. The market reaction to a spin off usually occurs on this announcement. Once the announcement has been made, the firm approaches the Internal Revenue Service or obtains a professional tax opinion on the tax status of the spin off. While the tax code in the United States treats a spin off as a dividend, the spin off is tax exempt if the firm fulfils the following requirements:

1. Both the parent and the subsidiary have been in active operations for at least 5 years prior to the spin off distribution date.
2. The parent company had control of the subsidiary before the spin off and gives up this control after the spin off. In general, the spun off shares have to represent at least 80% of the outstanding value of the unit, and the parent company must not be able to maintain effective control with the remaining shares. In other words, the subsidiary has to become independent of the parent company.
3. There must be a business reason for the spin off, and the objective cannot be purely distribution of profits. Legitimate business reasons are usually broadly defined to include giving managers a stake in ownership of the unit, complying with anti-trust laws and enhancing access to capital markets.

After obtaining a legal opinion, the firm will have to file Form 10 with the SEC. This form, which resembles the prospectus in an initial public offering, contains information about the unit being spun off and supporting financial statements. If the spin off is a large portion of the firm (as a percent of firm value) or if the corporate charter requires it, the firm must obtain stockholder approval for the action.

The firm will then either apply for stock exchange listing of the shares in the spun off unit or arrange for over-the-counter trading. Often, institutional investors will begin trading these units before they are actually issued; such trading is said to occur on a “when issued” basis. Thus, by the time the distribution of shares to existing stockholders occurs, the shares already have been priced in the market. Shareholders are then free to hold on to the shares or sell them in the market. The steps in the process are summarized in figure 8.9.

Figure 8.9: Steps in a Spin Off

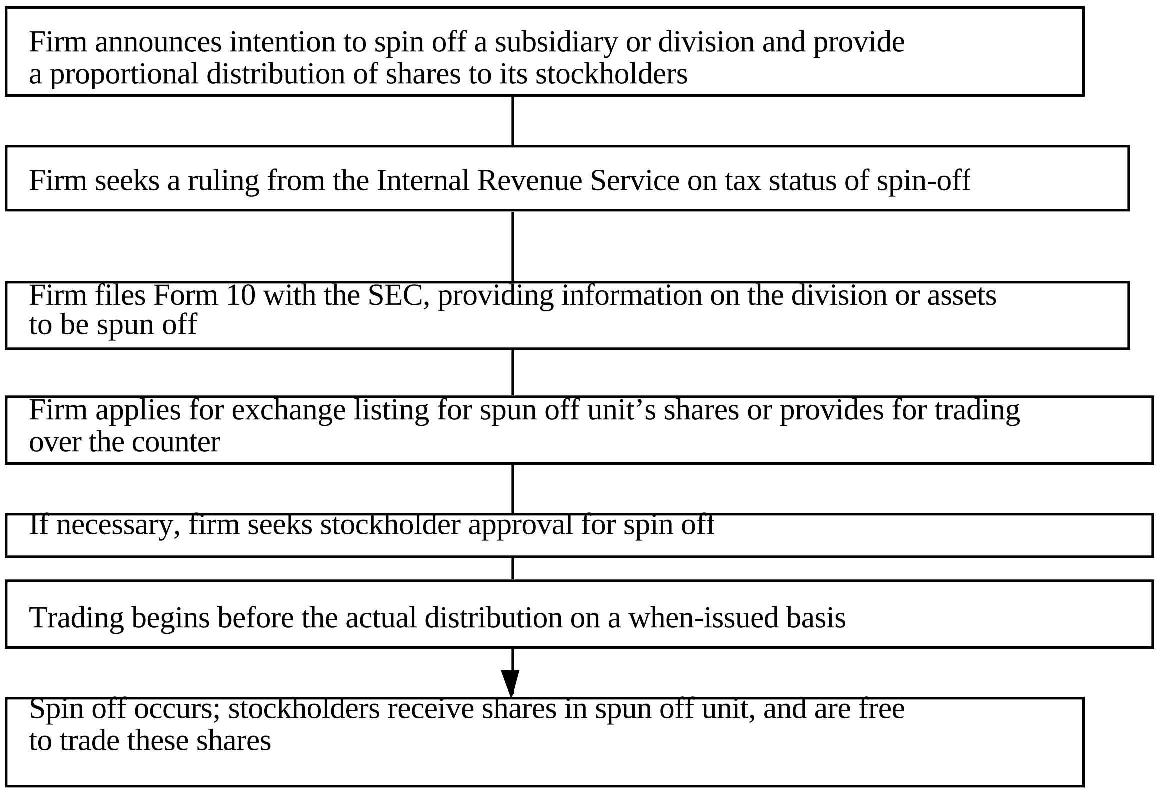

Reasons for Spin Offs

There are two primary differences between a divestiture and a spin off. The first is that there is often no cash generated for the parent firm in a spin off. The second is that the division being spun off usually becomes an independent entity, often with existing management in place. As a consequence, the first two reasons given for divestitures – a buyer who generates higher value from the assets than the divesting firm and the need to meet cash flow requirements – do not apply to spin offs. Improving the focus of the firm and returning to core businesses, which we offered as reasons for divestitures, can be arguments for spin offs as well. There are four other reasons:

- A spin off can be an effective way of creating value when subsidiaries or divisions are less efficient than they could be and the fault lies with the parent company, rather than the subsidiaries. For instance, Miles and Woolridge consider the case of Cyprus Minerals, a firm that was a mining subsidiary of Amoco in the early 1980s. Cyprus was never profitable as an Amoco subsidiary. In 1985, it was spun off after losing $ 95 million in the prior year. Cyprus cut overhead expenses by 30% and became profitable within six months of the spin off. Since the management of Cyprus remained the same after the spin off, the losses prior to it can be attributed to the failures of Amoco's management. When a firm has multiple divisions, and the sum of the divisional values is less than what the parent company is valued at, we have a strong argument for a split off, with each division becoming an independent unit.

- The second advantage of a spin off or split off, relative to a divestiture, is that it might allow the stockholders in the parent firm to save on taxes. If spin offs and split offs meet the tax tests described in the last section, they can save stockholders significant amounts in capital gains taxes. In 1992, for instance, Marriott spun off its hotel management business into a separate entity called Marriott International; the parent company retained the real estate assets and changed its name to Host Marriott. The entire transaction was structured to pass the tax test, and stockholders in Marriott were not taxed on any of the profits from the transaction.
- The third reason for a spin off or split off occurs when problems faced by one portion of the business affect the earnings and valuation of other parts of the business. As an example, consider the pressure brought to bear on the tobacco firms, such as Philip Morris and RJR Nabisco, to spin off their food businesses, because of investor perception that the lawsuits faced by the tobacco businesses weighed down the values of their food businesses as well.
- Finally, spin offs and split offs can also create value when a parent company is unable to invest or manage its subsidiary businesses optimally because of regulatory constraints. For instance, AT&T, as a regulated telecommunications firm, found itself constrained in decision making in its research and computer divisions. In 1995, AT&T spun off both divisions: the research division (Bell Labs) was renamed Lucent Technologies and its computer division reverted back to its original name of NCR.

Why would a firm use a split up instead of spin off or split off? By giving existing stockholders an option to exchange their parent company stock for stock in the split up unit, the firm can get a higher value for the assets of the unit. This is so because those stockholders who value the unit the most will be most likely to exchange their stock. The approach makes sense when there is wide disagreement between stockholders on how much the unit is worth.

Market Reactions to Spin Offs

Two issues have been examined by researchers who have looked at spin offs. The first relates to the stock price reaction to the announcement of spin offs. In general, these studies find that the parent company's stock price increases on the announcement of a spin off. A study by Schipper and Smith in 1983 examined 93 firms that announced spin offs between 1963 and 1981 and reported an average excess return of 2.84% in the two days surrounding the announcement. Similar results were reported by Hite and Owens in 1983 and by Miles and Rosenfeld in the same year. Further, there is evidence that the excess returns increase with the magnitude of the spun off entity. Schipper and Smith also find evidence that the excess returns are greater for firms in which the spin off is motivated by tax and regulatory concerns.

The second set of studies look at the performance of both the spun-off units, and the parent companies, after the spin off. These studies, which are extensively documented in Miles and Woolridge, can be summarized as follows:

- Cusatis, Miles and Woolridge report that both the spun off units and the parent companies report positive excess returns in the 3 years after the announcement of the spin offs. Figure 8.10 reports the total returns and the returns adjusted for overall industry returns in the three years after the spin off.

Figure 8.10: Returns at Spin Offs and Parent Company

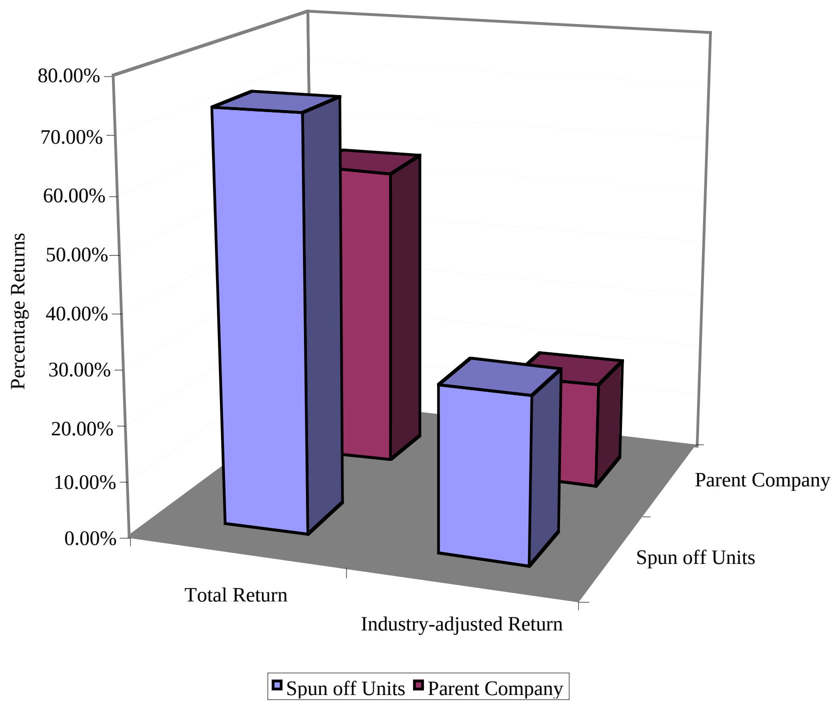

Source: Cusatis, Miles and Woolridge

Both groups are much more likely to be acquired, and the acquisition premiums explain the overall positive excess returns.

- There is a significant improvement in operating performance at the spun-off units in the 3 years after the spin off. Figure 8.11 reports on the change in revenues, operating income, total assets and capital expenditures at the spun off units in the three years after the spin off, before and after adjusting for the performance of the sector.

Figure 8.11: Operating Performance of Spun Off Units

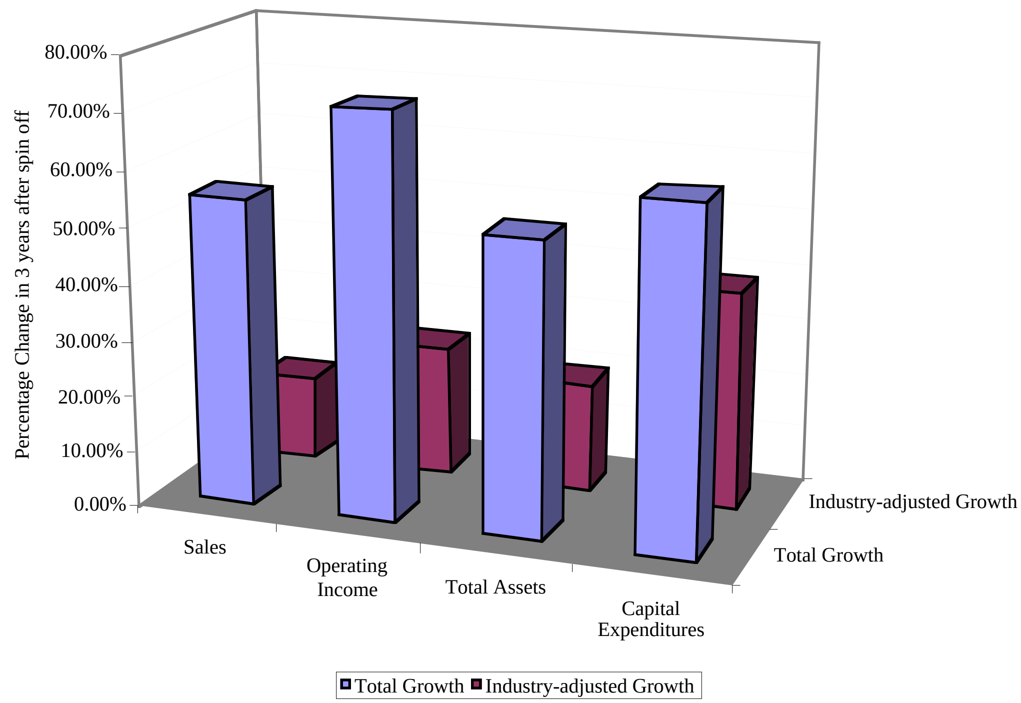

Source: Miles and Woolridge

Note that the spun off units grow faster than their competitors in terms of revenues and operating income; they also reinvest more in capital expenditures than other firms in the industry.

You can be too conservative

In corporate finance, there has long been a debate about whether firms can become more valuable as a result of changing the amount of debt that they carry on their books. There is one school of thought, attributed to Miller and Modigliani, that argues that value is independent of financial leverage, but only in a world without taxes and default risk. Another school of thought argues that in the presence of taxes and default risk, there is an optimal amount of debt that a firm can carry, and that value is maximized at that point. Finally, there is a school of thought argues that firms should not use debt, since it makes equity more risky, and that less debt is always better than more debt. We believe in the optimal debt ratio school and that firms can, in fact, be too conservative in their use of debt.

Are some firms underlevered?

What kinds of firms have too little debt or are underlevered? At an intuitive level, you would expect a firm with stable and large cashflows from operations and a high tax rate to gain substantial value from the use of debt. If such a firm chooses not to borrow money or has very little debt on its books, you could argue that it is in fact costing its stockholders

There is both anecdotal and empirical evidence that some firms are underlevered and that others are overlevered. You can come to this conclusion by comparing a firm to otherwise similar firms in the same business or by looking at the relationship between debt ratios and variables such as earnings variability and tax rates across the market. In 1984, Bradley, Jarrell, and Kim analyzed whether differences in debt ratios can be explained by some of the variables listed above. They noted that the debt ratio was lower for firms with more volatile operating income. Since these firms are also likely to face much higher likelihood of bankruptcy, this finding is consistent with the proposition that firms with high bankruptcy costs borrow less. They also looked at firms with high advertising and R&D expenses; lenders to these firms are likely to be much more concerned about recouping their debt if the firm gets into trouble, because the assets of these firms are intangible (brand names or patents) and difficult to liquidate. These firms, consistent with the theory, have much lower debt ratios. They also find that the there are a significant number of firms whose debt ratios are much lower and much higher than predicted by the crosssectional relationship.

So what if you were an activist investor in a firm with excess debt capacity and a conservative management? Left to themselves, the managers will not use the debt capacity. Investors can try to force them to borrow more and increase the proportion of capital that comes from debt – this process is called a recapitalization. At the limit, they may even use the firm’s debt capacity to borrow the money themselves and buy the entire company in a leveraged acquisition or buyout.

Recapitalization

In a recapitalization, a firm changes its financial mix of debt and equity, without substantially altering its investments or asset holdings. You can recapitalize in many ways. For instance, you could try to increase your debt ratio by borrowing money and paying a dividend or buying back stock. The first action increases debt and the second reduces equity. Alternatively, you can swap debt for equity, where equity investors in your firm are offered equivalent amounts (in market value terms) in debt. If you want to reduce your debt ratio, you would reverse these actions, raising equity and reducing debt.

The boom in debt for equity recapitalization occurred in the late 1980s. A stud that looked at these recapitalizations came to two conclusions. The first was that almost every one of them was triggered by the threat of a hostile takeover. In other words, it is external pressure that forces managers to increase financial leverage. The second was that the average stock price reaction to recapitalizations is very positive. On average, in the sample of 45 recapitalizations studied, the stock price increased by 21.9%. This finding is not restricted to just stock buybacks. A study of 52 offers to exchange debt for equity found that stock prices increased by 14%.

We might be overreaching when we conclude that this is definitive evidence that these firms were under levered. After all, the stock price reaction to a buyback or exchange offer may be explained by a much simpler story, say dilution - there are fewer shares outstanding after these actions. Notwithstanding this, the evidence seems to indicate that firms that issue debt are often treated favorably by markets.

Leveraged Acquisitions

Another phenomenon of the late 1980s was the leveraged buyout. Here, a group of investors raise debt against the assets of a publicly traded firm, preferably one with stable earnings and marketable assets, and use the debt to acquire the outstanding shares in the firm. If they succeed in their endeavor, the firm becomes a private company, and the debt is partly or substantially paid down with the firm's cashflows or from asset sales over time. Once the firm has been nursed back to health and efficiency, it is taken public again, reaping (at least if all goes according to plan) substantial payoffs to the equity investors in the deal.

Studies of leveraged acquisitions suggest that they do, on average, deliver significant returns to their investors. However, some of the leveraged buyouts done towards the end of the 1980s failed spectacularly, highlighting again that leverage is a two-edged sword, elevating returns in good times and reducing them in bad times.

You have lousy managers

Both conglomerate discounts and under leverage are manifestations of a larger problem, which is that managers do not always put stockholder interests first. While you can fashion specific solutions to both of these problems, they may not be sufficient in a firm where the source of the problem is poor management. For such firms, the only long term solution to value generation is a new management team.

Changing Top Management

If you are an activist investor in a firm with incompetent management, how would you go about instituting change? Needless to say, you will not have the cooperation of the existing management, who you have labeled as not up to the job. If you are able to harness enough stockholders to your cause, though, you may be able to increase the pressure on the top management to step down. While some may view the loss of top management in a company to be bad news, it really depends upon the market's perception of the management. The overall empirical evidence suggests that changes in management are generally viewed as good news. In figure 8.12, for instance, we examine how stocks react when a firm's CEO is replaced.

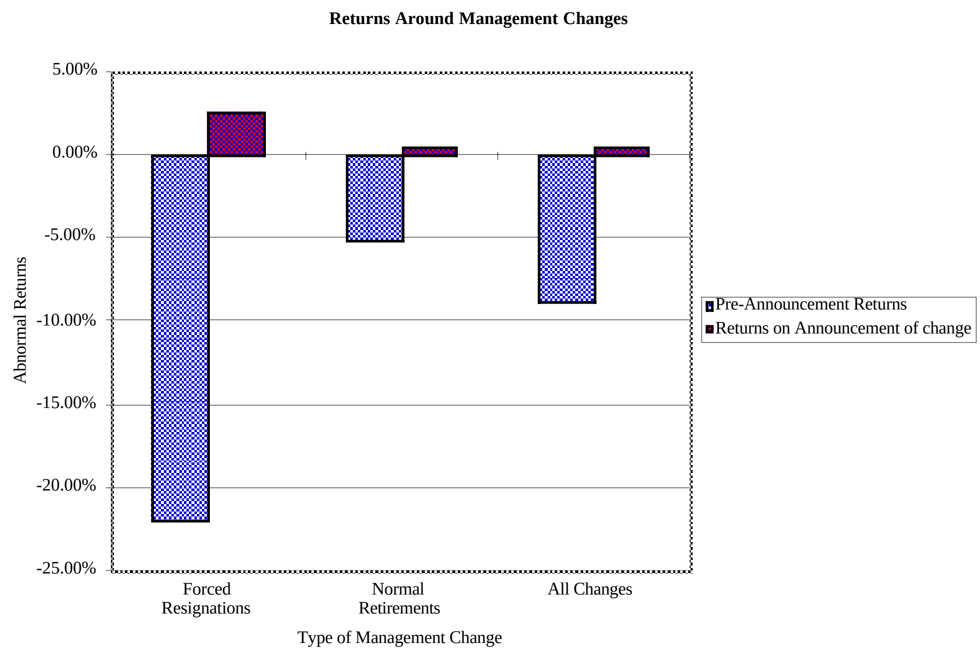

The price goes up, on average, when top management is changed. However, the impact of management changes is greatest when the change is forced. Management is more likely to be forced out in the aftermath of negative returns, and stock prices increase after the change is announced.

Hostile Acquisitions

If you cannot get top management to leave the firm, you can actively seek out hostile acquirers for the firm. If others share your jaundiced view of the management of the firm, you may very well succeed. There is evidence that indicates that badly managed firms are much more likely to be targets of acquisitions than well managed firms. Figure 8.13 summarizes key differences between target firms in friendly and hostile takeovers.

Target Characteristics - Hostile vs. Friendly Takeovers

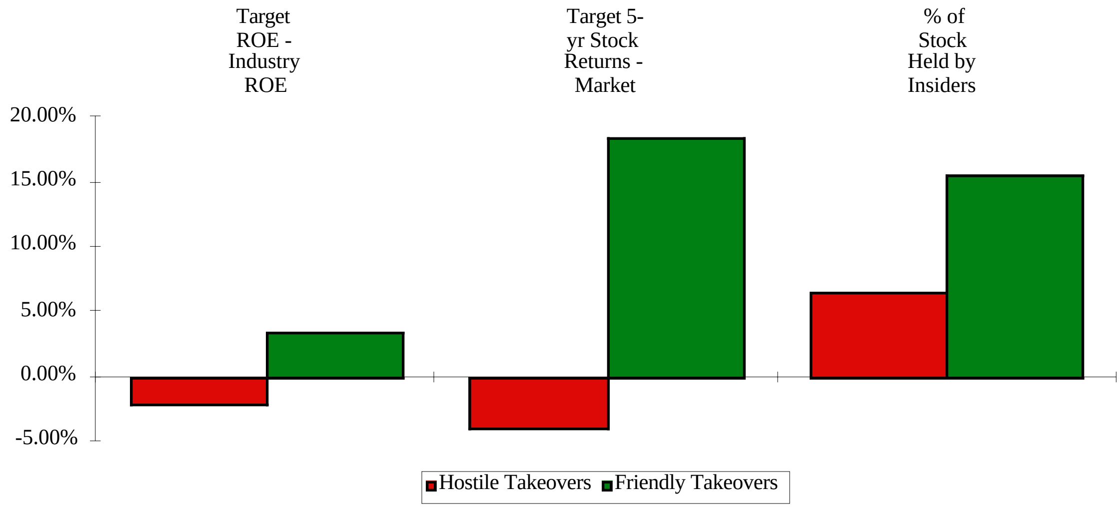

Source: Bhide

Note that target firms in hostile takeovers generally have much lower returns on equity (relative to their peer group), done worse for their stockholders and have less insider holdings than target firms in friendly takeovers. Needless to say, the payoff to being the stockholder of a firm that is the target of a hostile takeover is huge.

Empirical Evidence on Activist Investing

The overall evidence on whether activist investing works is mixed. While there are individual activist investors who have earned high returns by getting corporate managers to bend to their wishes, studies indicate that managers are both stubborn and resilient. For instance, studies[14] that have examined proxy fights find that there is little or no stock price reaction to proxy proposals by activist investors. This suggests that markets are not optimistic about changes occurring as a result of these proposals. However, a study by Wahal indicates that the price reaction to proxy fights is more positive when you look at only the sub-sample of companies that were targeted for poor stock price performance.

A study by Caton, Goh and Donaldson looked at companies on the Focus List – a list of poorly performing companies targeted by the Council of Institutional Investors. On average, these companies report higher earnings and stock returns after they are put of the list. However, when the sample of 138 companies was broken up into companies that traded at a market value less than replacement cost (Tobin's Q< 1) and at a market value greater than the replacement cost (Tobin's Q > 1), the improvement in earnings and stock prices was only in the latter group. Summarizing the evidence, we would suggest that shareholder

---

[14] See Karpoff, Malatesta and Walkling (1996).

activism has a chance of succeeding at firms whose stock prices have done badly and where there is potential for improved performance. It is unlikely to yield results when it is focused on firms with positive stock price performance or where management is not at fault for poor performance.

Determinants of Success

Activist value investors have an advantage over passive value investors since they can provide the catalysts for value creation. So, what is it that stops all of us from being activist value investors? When we consider some of the pre-requisits for being a successful value investor, we can also see why there are so few successful ones.

- This power of activist value investing usually comes from having the capital to buy significant stakes in poorly managed firms and using these large stockholder positions to induce management to change their behavior. Managers are unlikely to listen to small stockholders, no matter how persuasive their case may be.
- In addition to capital, though, activist value investors need to be willing to spend substantial time fighting to make themselves heard and in pushing for change. This investment in time and resources implies that an activist value investor has to pick relatively few fights and be willing to invest substantially in each fight.
- Activist value investing, by its very nature, requires a thorough understanding of target firms, since you have to know where each of these firms is failing and how you would fix these problems. Not surprisingly, activist value investors tend to choose a sector that they know really well and take positions in firms within that sector. It is clearly not a strategy that will lead to a well diversified portfolio.
- Finally, activist value investing is not for the faint hearted. Incumbent managers are unlikely to roll over and give in to your demands, no matter how reasonable you may thing them to be. They will fight, and sometimes fight dirty, to win. You have to be prepared to counter and be the target for abuse. At the same time, you have to be adept at forming coalitions with other investors in the firm since you will need their help to get managers to do your bidding.

If you consider all these requirements for success, it should come as no surprise that most conventional mutual funds steer away from activist value investing. Even though they might have the capital to be activist investors, they do not have the stomach or the will to go up against incumbent managers. The most successful activist value investors have either been individuals, like Michael Price, or small focused mutual funds, like the Lens Fund. As a small individual investor, you can try to ride their coattails, and hope that they succeed, but it is unlikely that you could succeed at activist value investing.

Michael Price: Activist Investing

In the 1990s, Michael Price acquired a reputation for buying stock in what were perceived as poorly managed companies and pushing for change. In the process, he enriched shareholders at Mutual Shares, the fund that he ran. One example was his investment in Chase Manhattan in the mid-80s, where after he acquired the shares, he pushed the firm to merge with Chemical. He argued that the latter's management would shake up the moribund culture at Chase and make it a more profitable firm. While Chase's management initially fought the merger, they ultimately succumbed to his pressure and the subsequent merger generated substantial returns for Mutual Shares.

Price served his apprenticeship in Max Heine, a German Jew who fled Austria and became a contrarian value investor who became co-manager of Mutual Shares. Heine looked for cheap assets that were out of favor. He bought railroad bonds for cents on the dollar in the 1970s and made his money back several times over. As Price paraphrases it, Heine taught him to "stay away from the crowd and buy thing at a big discount". Like Heine, Price prefers less visible stocks that are underpriced, though unlike Heine, he has been willing to take large positions in high profile firms like Dow Jones and Sunbeam and push for change. Price also does not have much faith in equity research, which he believes is designed to enrich Wall Street and not investors.

Conclusion

Value investing comes in many stripes. First, there are the screeners, who we view as the direct descendants of the Ben Graham school of investing. They look for stocks that trade at low multiples of earnings, book value or revenues, and argue that these stocks can earn excess returns over long periods. It is not clear whether these excess returns are truly abnormal returns, rewards for having a long time horizon or just the appropriate rewards for risk that we have not adequately measured. Second, there are contrarian value investors, who take positions in companies that have done badly in terms of stock prices and/or have acquired reputations as poorly managed or run companies. They are playing the expectations game, arguing that it is far easier for firms such as these to beat market expectations than firms that are viewed as successful firms. Finally, there are activist investors who take positions in undervalued and/or badly managed companies and by virtue of their holdings are able to force changes in corporate policy or management that unlock this value.

What, if anything, ties all of these different strands of value investing together? In all of its forms, the common theme of value investing is that firms that are out of favor with the market, either because of their own performance or because the sector that they are in is in trouble, can be good investments.

Lessons for investors

To be a value investor, you should have

- A long time horizon: While the empirical evidence is strongly supportive of the long-term success of value investing, the key word is long term. If you have a time horizon that is less than 2 or 3 years, you may never see the promised rewards to value investing.
- Be willing to bear risk: Contrary to popular opinion, value investing strategies can entail a great deal of risk. Firms that look cheap on a price to earnings or price to book basis can be exposed to both earnings volatility and default risk.

In addition to these, to be a contrarian value investor, you need

- A tolerance for bad news: As a contrarian investor who buys stocks that are down and out, you should be ready for more bad news to come out about these stocks. In other words, things will often get worse before they get better.

In addition to all of the above, to be an activist investor, you have to

- Be willing to fight: Incumbent managers in companies that you are trying to change will seldom give in without a fight.

CHAPTER 9

THE ALLURE OF GROWTH: SMALL CAP AND GROWTH INVESTING

There is a widespread belief that while value investing is for the risk averse, growth investing is the investment philosophy of those who like to take risk. Though there is nothing wrong with seeking out risk, taking on risk for the sake of doing so is foolhardy. Growth clearly has value, but the real issue is whether you can buy it at a reasonable price. In this chapter, we will examine the basis of growth investing and dispense with the notion that all growth investors are risk seekers. As with value investing, we will look at the various strands of growth investing and examine what you would need to succeed with each.

Who is a growth investor?

Many services define a growth investor as one who buys stocks that trade at high multiples of earnings. Though this may be a convenient way to categorize investors, it is not an accurate one. In fact, it leaves us with the misleading picture of growth investors as being uninterested in the value of what they are buying. While this may be true for some growth investors, does anyone really believe that Peter Lynch, who built Fidelity Magellan by focusing on growth companies, cares less about value than Warren Buffett does?

We will define growth investors as those who buy companies whose growth potential is being undervalued by the market. With our categorization, note that growth investors care just as much about value as value investors do. What then, you might wonder, is the distinction between growth and value investors? In our view, the key difference lies in where the focus for finding value lies. As we argued in the last chapter, value investors believe that you are more likely to find under valuation of assets in place and tend to invest in mature firms with substantial existing assets, albeit underperforming ones. Growth investors believe that they are more likely to find bargains in growth investments.

In the sections that follow, we will consider the different strands of growth investing. We will begin by looking at passive growth investing strategies, where we focus on investing in stocks that passes a specific screen - for instance, PE ratios that are less than expected growth rates in earnings per share. We will then consider active growth investing strategies, where investors not only take large positions in growth companies, but also actively involve themselves in the management of these companies. It is in this category that we consider venture capital and private equity investing.

Passive Growth Investing

In passive growth investing, as in passive value investing, we use screens to find stocks that are under valued by the market. The simplest version of passive growth investing is investing in small growth companies, with small defined in terms of market capitalization. Next, we look at investing in initial public offerings, with the intent of capturing any excess returns associated with the stock going up after the offering. Finally, we consider more conventional growth investing strategies, by first looking at a strategy of buying companies with high growth, then evaluate a strategy of buying high PE stocks and finally a more nuanced strategy of buying growth stocks, but only at a reasonable price.

Small Cap Investing

One of the most widely used passive growth strategies is the strategy of investing in small companies, with small defined in terms of market capitalization. While you could construct a value oriented, small cap portfolio, most small cap portfolios tend to be tilted towards growth companies, and we believe that this category fits better in this chapter. We will begin by reviewing the empirical evidence on small cap investing, and then look at the requirements for success at this strategy.

The Small Cap Effect

Studies have consistently found that smaller firms (in terms of market value of equity) earn higher returns than larger firms of equivalent risk, where risk is defined in terms of the market beta. Figure 9.1 summarizes annual returns for stocks in ten market value classes, for the period from 1927 to 2001.[1] The portfolios were reconstructed at the end of each year, based upon the market values of stock at that point in time, and held for the subsequent year.

---

[1] These annual returns were obtained from the annual returns data set maintained by Ken French and Gene Fama on market value classes.

Figure 9.1: Annual Returns by Market Value Class - 1927 - 2001

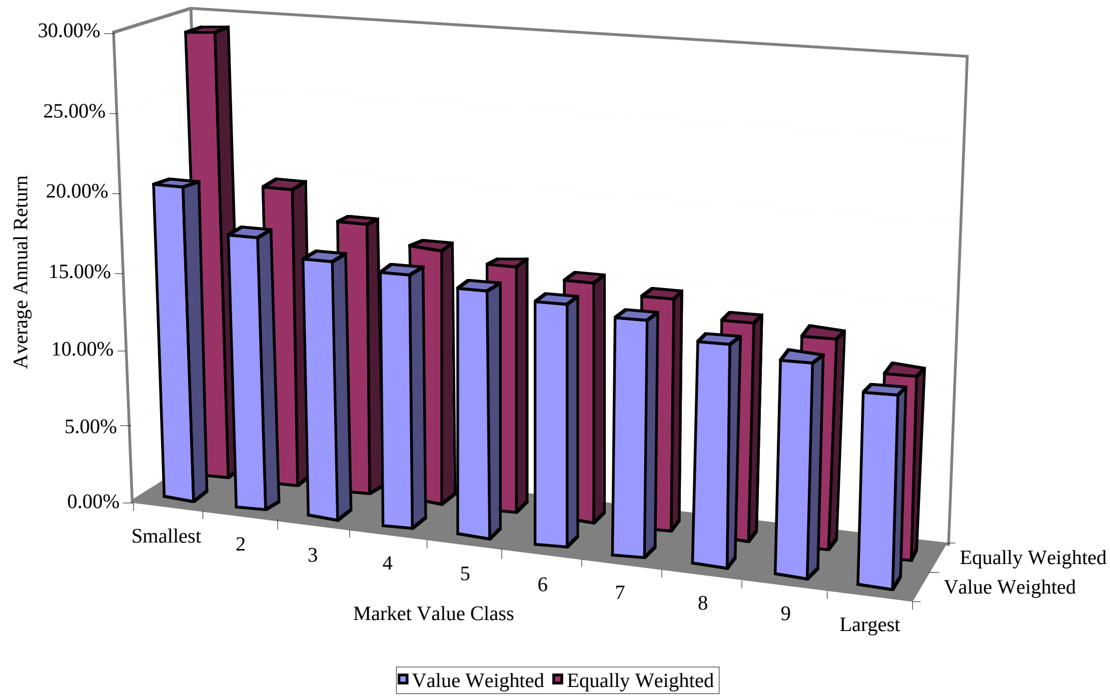

Source: Raw data from French

If we look at value weighted portfolios, the smallest stocks earned an annual return of about 20% over the period as contrasted with the largest stocks which earned an annual return of 11.74%. If we use an equally weighted portfolio, the small firm premium is much larger, an indication that the premium is being earned by the smallest stocks. In other words, to capture the small cap premium, you would have to invest in the very smallest companies in the market. Nevertheless, these results are impressive and provide a rationale for the number of portfolio managers who focus on buying small cap stocks. Before we conclude that small cap investing is the way to go, though, we do have to consider some of the details of the small stock premium.

Small Cap Cycles

On average, have small cap stocks outperformed large cap stocks over this period? Absolutely, but, success from this strategy is by no means guaranteed in every time period. While small cap stocks have done better than large cap stocks in more periods than not, there have been extended periods where small cap stocks have underperformed large cap stocks. Figure 9.2 graphs the premium earned by small cap stocks over large cap stocks from 1927 to 2001.

Figure 9.2: Small Firm Premium over time- 1927 -2001

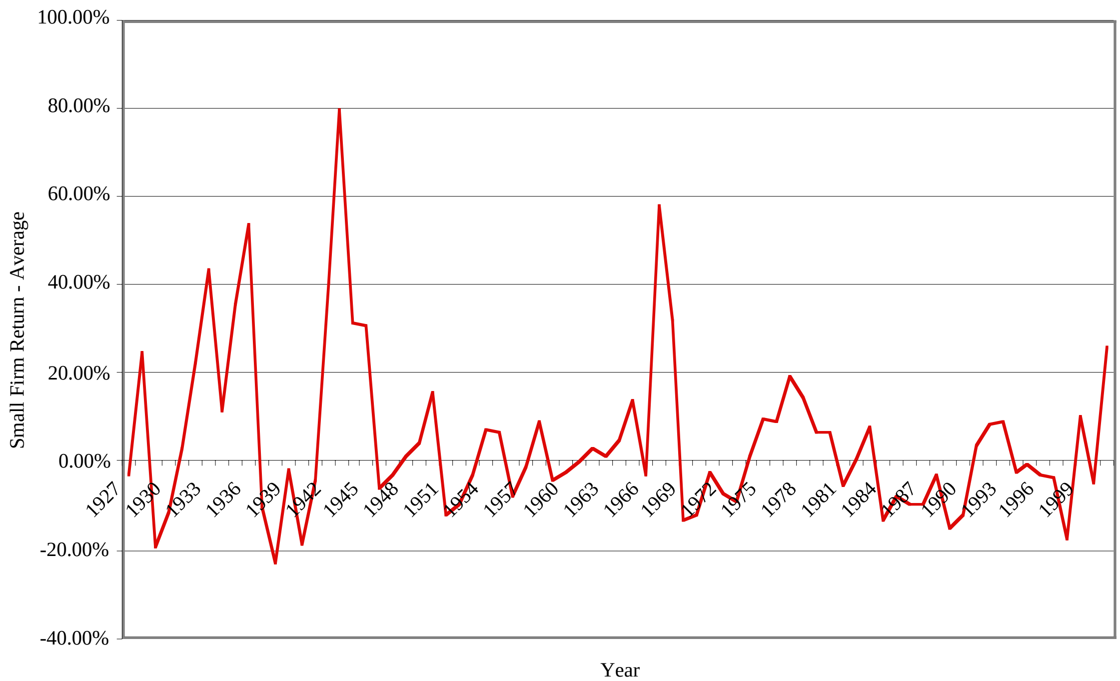

Source: Raw data from French

Note that the premium is negative in a significant number of years – small stocks earned lower returns than large stocks in those years. In fact, during the 1980s, large market cap stocks outperformed small cap stocks by a significant amount, creating a debate about whether this was a long term shift in the small stock premium or just a temporary dip. On the one side, Jeremy Siegel notes that the small stock premium can be almost entirely attributed to the performance of small stocks in the late 1970s. Since this was a decade with high inflation, could the small stock premium have something to do with inflation? On the other side are small cap portfolio managers, arguing that the events of the 1980s were an aberration and that the small stock premium would return. On cue, the small stock premium returned in the 1990s, as can be seen in figure 9.3 below:

Small Cap Effect over Time

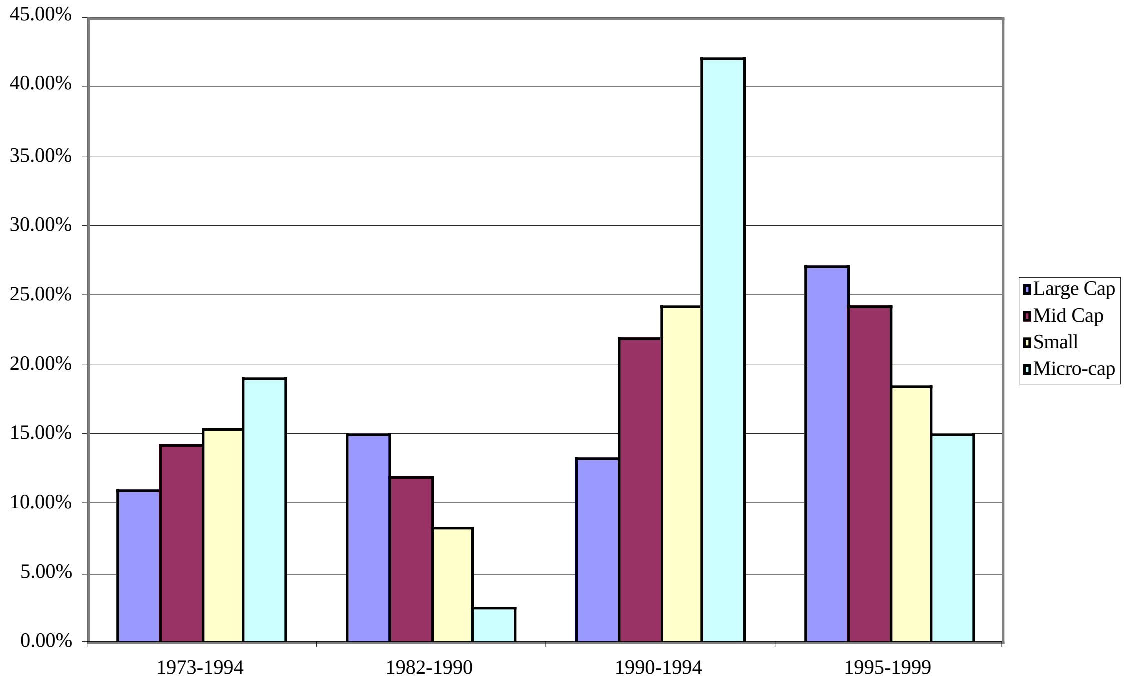

Source: Pradhuman (1998)

Pradhuman takes a close look at the small cap premium in his book on the topic.[2] He notes that small cap stocks tend to do much better than large cap stocks when the yield curve is downward sloping and inflation is high, which may explain why the premium was high in the 1970s. He also finds that the small cap premium tends to be larger when default spreads on corporate bonds narrow. In summary, there is a return premium for small cap stocks but it is a volatile one. While the premium clearly exists over long time periods, it also disappears over extended periods.

Deconstructing the Small Cap Effect

A number of studies have tried to take a closer look at the small cap effect to see where the premium comes from. The following are some of the conclusions:

- The small cap effect is greatest in the micro-cap companies, i.e., the really small companies). In fact, many of these companies have market capitalizations of $250 million or lower. All too often these are also companies that have low priced and illiquid stocks, not followed by equity research analysts.

---

[2] The book titled “Small Cap Dynamics” is one of the most detailed looks at the phenomenon.

- A significant proportion of the small cap premium is earned in January. Figure 9.4 presents the contrast between small cap and large cap companies in January and for the rest of the year between 1935 and 1986:

Figure 9.4: The Small Firm Effect in January

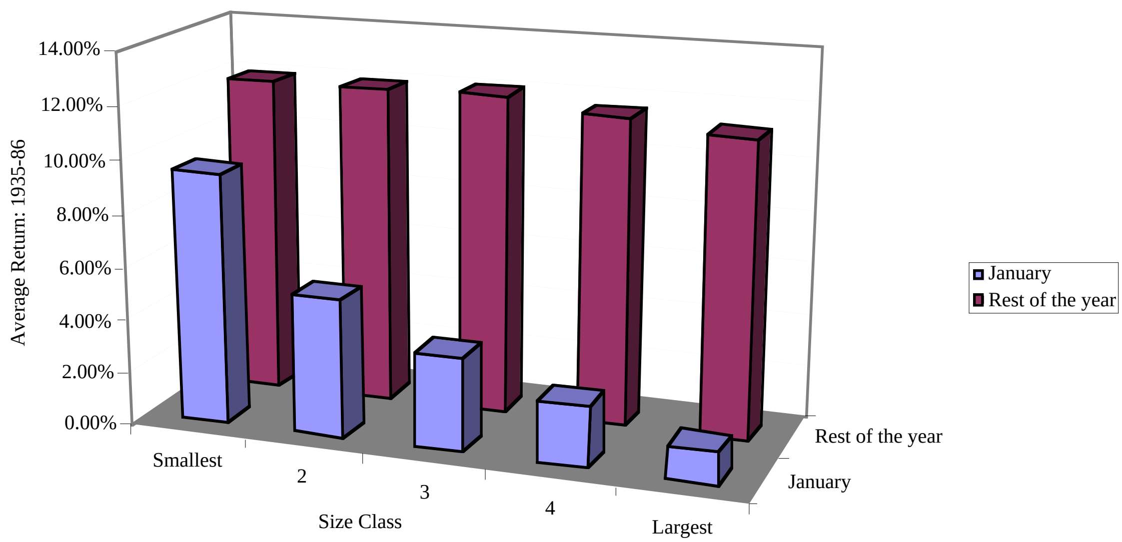

Source: Raw data from French

In fact, you cannot reject the hypothesis that there is no small cap premium from February to December. Many of the other temporal anomalies that we noted in chapter 7 such as the weekend effect also seem to be greater for small cap companies.

- There is evidence of a small firm premium in markets outside the United States. Studies find small cap premiums of about 7% from 1955 to 1984 in the United Kingdom,[3] 8.8% in France and a much smaller size effect in Germany[4] and a premium of 5.1% for Japanese stocks between 1971 and 1988.[5]

Explanations for the Small Stock Premium

The persistence of the small stock premium has led many to argue that what looks like a premium in empirical studies comes the failure to allow for transactions costs and

---

[3] See Dimson and Marsh,

[4] Updated numbers are reported by Fama and French.

[5] Chan, Hamao and Lakonishok

measure risk correctly in firms. There is truth in these arguments, though it is unclear whether the small stock premium would disappear even if they were considered.

Transactions Costs

The transactions costs of investing in small stocks are significantly higher than the transactions cots of investing in larger stocks, and the premiums are estimated prior to these costs. In chapter 5, for instance, we looked at the bid-ask spread as a percent of the stock price and noted that it tended to be higher for smaller companies. In addition the price impact from trading is also higher for small cap stocks because they are less liquid. Can the difference in transactions costs overwhelm the small cap premium? The answer has to depend upon your time horizon. With short time horizons, the transactions costs can wipe out any perceived excess returns associated with small cap companies, With longer time horizons, though, you can spread the costs over your holding period and the excess returns may persist.

In a telling illustration of the difficulties associated with replicating the small firm premiums that are observed in the studies in real time, we compare the returns on a hypothetical small firm portfolio (CRSP Small Stocks) with the actual returns on a small firm mutual fund (DFA Small Stock Fund), which passively invests in the same small stocks in figure 9.5:

Figure 9.5: Returns on CRSP Small Stocks versus DFA Small Stock Fund

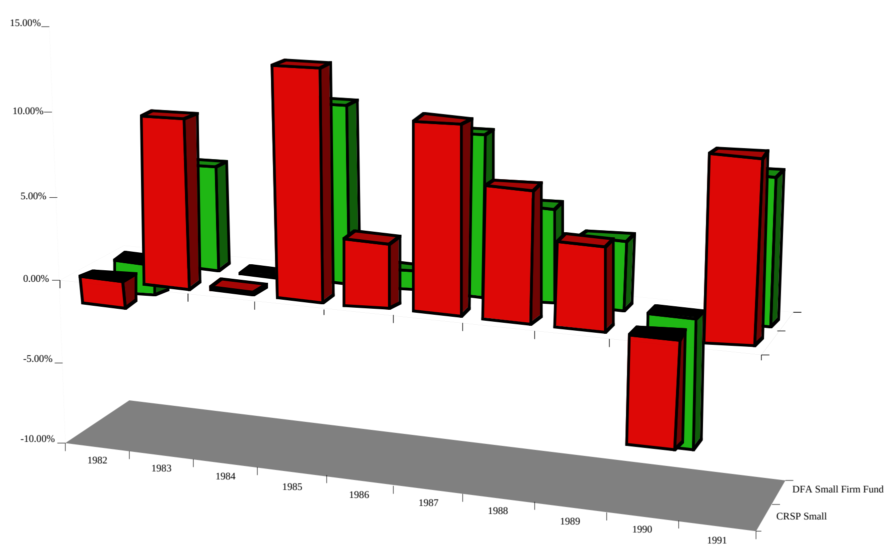

Note that the returns on the DFA fund consistently lag the returns on the hypothetical portfolio by about 2%, reflecting the transactions and execution costs faced by the fund.

Failure to consider liquidity and estimation risk

Many of the studies that uncover a small cap premium measure the risk of stocks using a market beta and the capital asset pricing model. It is entirely possible that the capital asset pricing model is not the right model for risk, and betas under estimate the true risk of small stocks. Thus, the small firm premium may really reflect the failure of the market beta to capture risk. The additional risk associated with small stocks may come from several sources. First, the estimation risk associated with estimates of beta for small firms is much greater than the estimation risk associated with beta estimates for larger firms, partly because of the fact that small companies tend to change more over time and partly because of their short histories. The small firm premium may be a reward for this additional estimation risk.[6] Second, there may be much greater liquidity risk associated with investing in small companies. This risk (which is also partially responsible for the higher transactions costs noted in the previous section) is not captured in betas.

While the argument that liquidity and estimation risk can be significant problems for small cap stocks seems unexceptional, there is one problem with it. Note that portfolios of small cap stocks do not carry the same risk as individual stocks and that estimation risk, in particular, should be diversifiable. Estimation risk will lead you to under estimate the risk (or betas) of some small companies and over estimate the risk (or betas) of other small companies. The beta of a portfolio of such companies should still be predictable, because the estimation errors should average out. With illiquidity, the diversification argument is tougher to make, since it manifests itself as a higher cost (bid-ask spread or price impact) for all small stocks. Thus, the illiquidity risk will show up as higher transactions costs in a small-cap portfolio and will increase as trading in the portfolio increases.

Information Risk

When investing in publicly traded companies, we tend to rely not only on the financial reports filed by the company but also on the opinions of analysts following the company. We expect these analysts, rightly or wrongly, to collect information about the firm and reveal this information in their reports. With a large and widely held firm, it is not uncommon to see 25 or 30 analysts following the firm and substantial external information on the firm. Many small cap firms are followed by one or two analysts and many are not followed by any, as you can see in figure 9.6.

---

[6] The problem with this argument is that it does not allow for the fact that estimation risk cuts both ways – some betas will be underestimated and some will be overestimated – and should be diversifiable.

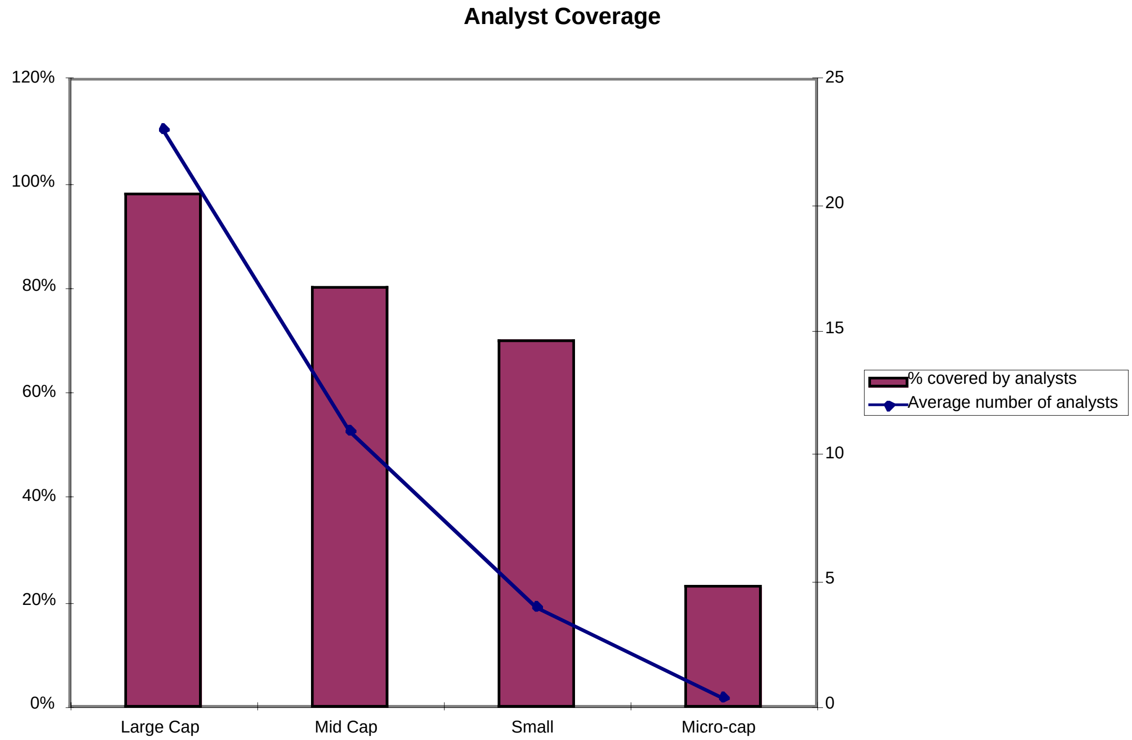

With some small-cap firms, you may find that the only source of information is the firm itself. While the firm may follow all of the regulatory requirements, the information revealed is unlikely to be unbiased, and it is entirely possible that bad news about the firm's operations may be withheld. Since you cannot diversify away this risk, you may demand a premium when investing in these companies.

Determinants of Success at Small-cap Investing

Let us concede, notwithstanding the period in the 1980s where the premium waned, that small cap stocks earn a premium over large cap stocks, when we adjust for risk using conventional measures like beta. Given the discussion in the last section about potential explanations for this premium, what do you need to do to succeed at small cap investing?

- The first and most critical factor seems to be a long time horizon, given the ups and downs of small cap premium. In figure 9.7, we examine the percent of time a small cap investor would have outstripped a large cap investor with different time horizons. Note that the number is close to 50% for time horizons up to five years, no different from a random strategy. Beyond 5 years though, small cap investing wins decisively.

Figure 9.7: Time Horizon and the Small Firm Premium

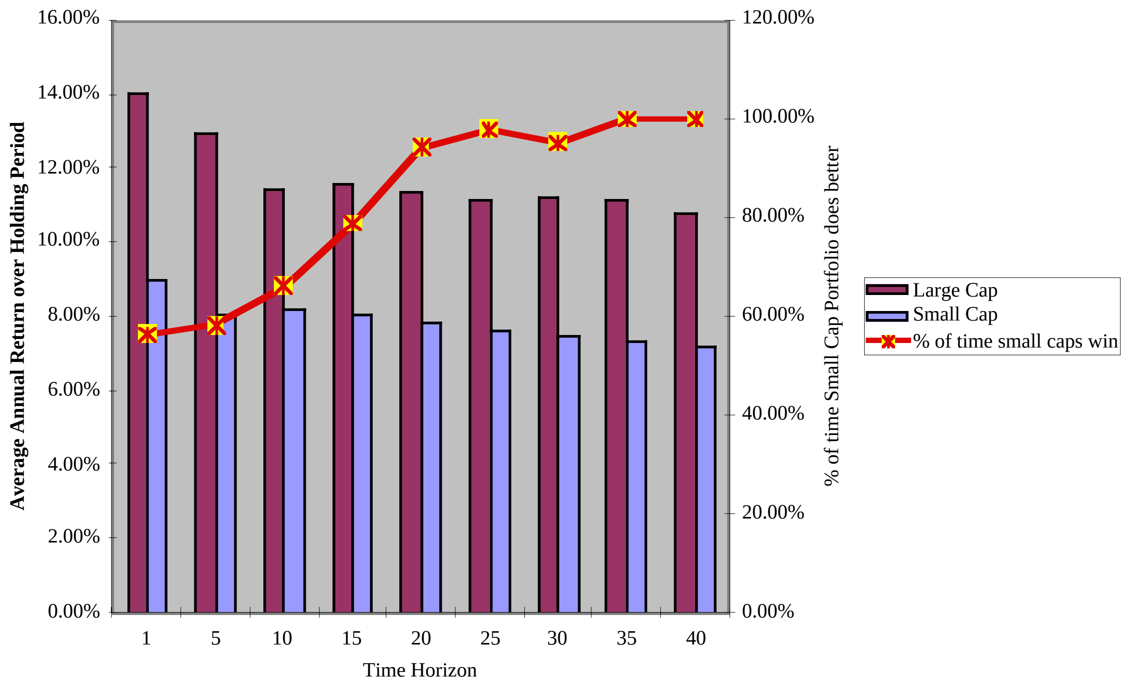

A long time horizon will also go a long way towards reducing the bite taken out of returns by transactions costs.

- The importance of discipline and diversification become even greater, if you are a small cap investor. Since small cap stocks tend to be concentrated in a few sectors, you will need a much larger portfolio to be diversified with small cap stocks.[7] In addition, diversification should also reduce the impact of estimation risk and some information risk.
- When investing in small cap stocks, the responsibility for due diligence will often fall on your shoulders as an investor, since there are often no analysts following the company. You may have to go beyond the financial statements and scour other sources (local newspapers, the firm’s customers and competitors) to find relevant information about the company.

If you combine the need for more stocks in your portfolio with additional research on each, you can see that small cap investing is likely to be more time and resource intensive than most other investment strategies. If you are willing to expend these

[7] The conventional rule of thumb for being diversified (where you diversify away 95% of the firm-specific risk) with large cap stocks is about 25 stocks. With small cap stocks, you would need to hold more stocks. How many more? It will depend upon your strategy, but you should consider holding at least 40-50 stocks.

resources and have a long time horizon, you may well be able to claim a large portion of the small cap stock premium going forward.

Small Cap Value Investing

While we have considered small cap investing as a strand of growth investing, you can be a small-cap value investor, if you focus on small companies that trade low PE or low PBV ratios – the conventional measures of value companies. Investors who do this hope to combine the excess returns that have been uncovered for buying stocks that trade at low multiples of earnings and book value with the excess returns associated with small cap investing.

Pradhuman, in his book on small cap investing, contrasts a strategy of buying small cap value stocks with small cap growth stocks and presents several results. First, the excess return on a small cap, value strategy is less than the sum of the excess return on a value strategy and the excess return on a small cap strategy. In other words, there is some leakage in returns from both strategies when you combine them. Second, the difference in returns between value and growth small-cap stocks mirrors the difference in returns between value and growth large-cap stocks, but the cycles are exaggerated. In other words, when value stocks outperform (underperform) growth stocks across the market, small-cap value stocks outperform (underperform) small-cap growth stocks by an even larger magnitude. Third, the excess returns in the last two decades on a small-cap, value strategy seem to be more driven by the value component than by the small-cap component. [8]

Initial Public Offerings

In initial public offerings, private firms make the transition to being publicly traded firms by offering their shares to the public. In contrast with equity issues by companies that are already publicly traded, where there is already a market price for the stock that acts as an anchor, an initial public offering has to be priced by an investment banker based upon perceptions of demand and supply. There are some investors who believe that they can exploit both the uncertainty in the process and the biases brought to the pricing by investment bankers to make excess returns.

The Process of an Initial Public Offering

When a private firm becomes publicly traded, the primary benefit it gains is increased access to financial markets and to capital for projects. This access to new capital is

[8] We came to this conclusion by regressing excess returns on stocks against market capitalization and price to book ratio. The latter explained far more of the differences in excess returns than the former.

a significant gain for high growth businesses, with large and lucrative investment opportunities. A secondary benefit is that the owners of the private firm are able to cash in on their success by attaching a market value to their holdings. These benefits have to be weighed against the potential costs of being publicly traded. The most significant of these costs is the loss of control that may ensue from being a publicly traded firm. Other costs associated with being a publicly traded firm are the information disclosure requirements and the legal requirements[9]. Assuming that the benefits outweigh the costs, there are four steps involved in an initial public offering.

I. Choosing an Investment Banker

Once the decision to go public has been made, a firm generally cannot approach financial markets on its own. This is so because it is largely unknown to investors and does not have the expertise to go public without help. Therefore, a firm has to pick intermediaries to facilitate the transaction. These intermediaries are usually investment bankers, who provide several services. First, they help the firm meet the requirements of the Securities and Exchange Commission (SEC) in preparing and filing the necessary registration statements needed for the public offering. Second, they provide the credibility a small and unknown private firm may need to induce investors to buy its stock. Third, they provide their advice on the valuation of the company and the pricing of the new issue. Fourth, they absorb some of the risk in the issue by guaranteeing an offer price on the issue; this guarantee is called an underwriting guarantee. Finally, they help sell the issue by assembling a group called an underwriting syndicate, who try to place the stock with its clients. The underwriting syndicate is organized by one investment bank, called the lead investment bank. Private firms tend to pick investment bankers based upon reputation and expertise, rather than price. A good reputation provides the credibility and the comfort level needed for investors to buy the stock of the firm; expertise applies not only to the pricing of the issue and the process of going public but also to other financing decisions that might be made in the aftermath of a public issue. The investment banking agreement is then negotiated, rather than opened up for competition.

II. Valuing the Company and Setting Issue Details

Once the firm chooses an investment banker to take it public, the next step is to estimate a value for the firm. This valuation is generally done by the lead investment bank, with substantial information provided by the issuing firm. The value is sometimes estimated

---

[9] The costs are two fold. One is the cost of producing and publicizing the information itself. The other is the loss of control over how much and when to reveal information about the firm to others.

using discounted cash flow models, similar to those described in chapter 5. More often, though, the value is estimated by using a multiple, like a price earnings ratio, and by looking at the pricing of comparable firms that are already publicly traded. Whichever approach is used, the absence of substantial historical information, in conjunction with the fact that these are small companies with high growth prospects, makes the estimation of value an uncertain one at best.

The other decision the firm has to make relates to the size of the initial issue and the use of the proceeds. In most cases, only a portion of the firm's stock is offered at the initial public offering; this reduces the risk on the under pricing and enables the owners to test the market before they try to sell more stock. In most cases, the firm uses the proceeds from the initial stock issue to finance new investments.

The next step in this process is to set the value per share for the issuer. To do so, the equity in the firm is divided by the number of shares, which is determined by the price range the issuer would like to have on the issue. If the equity in the firm is valued at $ 50 million, for example, the number of shares would be set at 5 million to get a target price range of $10, or at 1 million shares to get a target price range of $ 50 per share.

The final step in this process is to set the offering price per share. Most investment banks set the offering price below the estimated value per share for two reasons. First, it reduces the bank's risk exposure, since it ensures that the shares will be bought by investors at the offering price. (If the offering price is set too high and the investment bank is unable to sell all of the shares being offered, it has to use its own funds to buy the shares at the offering price.) Second, investors and investment banks view it as a good sign if the stock increases in price in the immediate aftermath of the initial issue. For the clients of the investment banker who get the shares at the offering price, there is an immediate payoff; for the issuing company, the ground has been prepared for future issues.

In setting the offering price, investment bankers have the advantage of first checking investor demand. This process, which is called building the book, involves polling institutional investors prior to pricing an offering, to gauge the extent of the demand for an issue. It is also at this stage in the process that the investment banker and issuing firm will present information to prospective investors in a series of presentations called road shows. In this process, if the demand seems very strong, the offering price will be increased; in contrast, if the demand seems weak, the offering price will be lowered. In some cases, a firm will withdraw[10] an initial public offering at this stage, if investors are not enthusiastic about it.

III. SEC Requirements

In order to make a public offering the United States, a firm has to meet several requirements. First, it has to file a registration statement and prospectus with the SEC, providing information about the firm's financial history, its forecasts for the future and how it plans for the funds it raises from the initial public offering. The prospectus provides information about the riskiness and prospects of the firm for prospective investors in its stock. The SEC reviews this information and either approves the registration or sends out a deficiency memorandum asking for more information. While the registration is being reviewed, the firm may not sell any securities, though it can issue a preliminary prospectus, titled a red herring, for informational purposes only.

Once the registration has been approved by the SEC, the firm can place a tombstone advertisement in newspapers and other publications. This ad contains details of the issue, the name of the lead investment banker, and the names of other investment bankers involved in the issue. The order in which the investment bankers are listed is significant. At the top is the lead investment banker and the co-managers of the issue, followed by the major bracket investment bankers. The categorization is based both upon reputation and national focus. Then comes the mezzanine bracket, which includes smaller investment banks that operate nationally, and at the bottom are the regional investment bankers involved with the issue. Figure 9.8 shows a typical tombstone advertisement for an initial public offering.

[10] One study of initial public offerings between 1979 and 1982 found that 29% of firms terminated their initial public offerings at this stage in the process.

Figure 9.8: Tombstone Advertisement

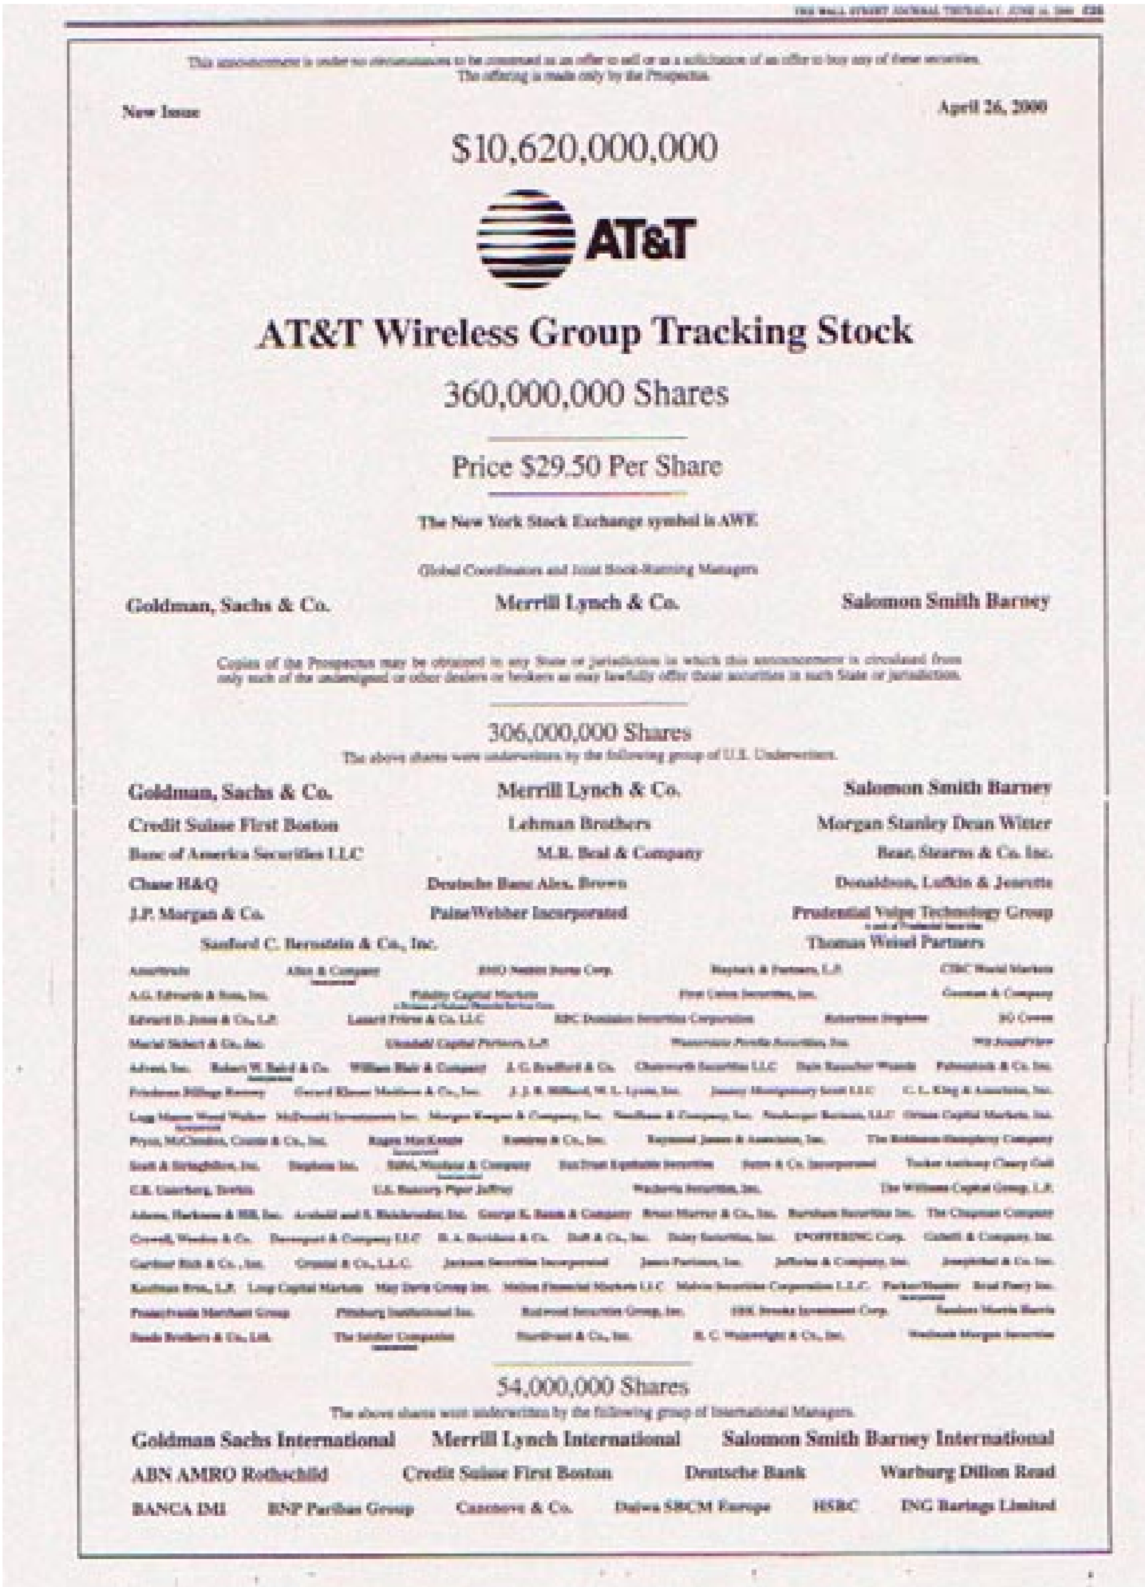

New Issue

April 26, 2000

$10,620,000,000

AT&T

AT&T Wireless Group Tracking Stock

360,000,000 Shares

Price $29.50 Per Share

The New York Stock Exchange symbol is AWE

Global Coordinators and Joint Book-Running Managers

Goldman, Sachs & Co. | Merrill Lynch & Co. | Salomon Smith Barney

306,000,000 Shares

The above shares were underwritten by the following group of U.S. Underwriters.

Goldman, Sachs & Co. | Merrill Lynch & Co. | Salomon Smith Barney

Credit Suisse First Boston | Lehman Brothers | Morgan Stanley Dean Witter

Banc of America Securities LLC | M.R. Beal & Company | Bear, Stearns & Co. Inc.

Chase H&Q | Deutsche Banc Alex. Brown | Donaldson, Lufkin & Jenrette

J.P. Morgan & Co. | PaineWebber Incorporated | Prudential Volpe Technology Group

Sanford C. Bernstein & Co., Inc. | Thomas Weisel Partners

54,000,000 Shares

The above shares were underwritten by the following group of International Managers.

Goldman Sachs International | Merrill Lynch International | Salomon Smith Barney International

ABN AMRO Rothschild | Credit Suisse First Boston | Deutsche Bank | Warburg Dillon Read

BANCA IMI | BNP Paribas Group | Cazenove & Co. | Daiwa SBCM Europe | HSBC | ING Barings Limited

IV. The Issue

Once the offering price has been set and the tombstone advertisement published, the die is cast. If the offering price has indeed been set below the true value, the demand will exceed the offering, and the investment banker will have to choose a rationing mechanism to allocate the shares. On the offering date — the first date the shares can be traded — there will generally be a spurt in the market price. If the offering price has been set too high, as is sometimes the case, the investment bankers will have to discount the offering to sell it and make up the difference to the issuer, because of the underwriting agreement.

IPO Pricing and Investment Strategies

How well do investment bankers price initial public offerings? One way to measure this is to compare the price when the stock first starts trading to the offering price. While precise estimates vary from year to year, the average initial public offering seems to be under priced by 10-15%. The under pricing also seems to be greater for smaller public offerings. One study[11] estimates the under pricing as a function of the issue proceeds for 1767 iPOs between 1990 and 1994, and the results are presented in figure 9.9 below:

Figure 9.9: Average Initial Return and Issue Size

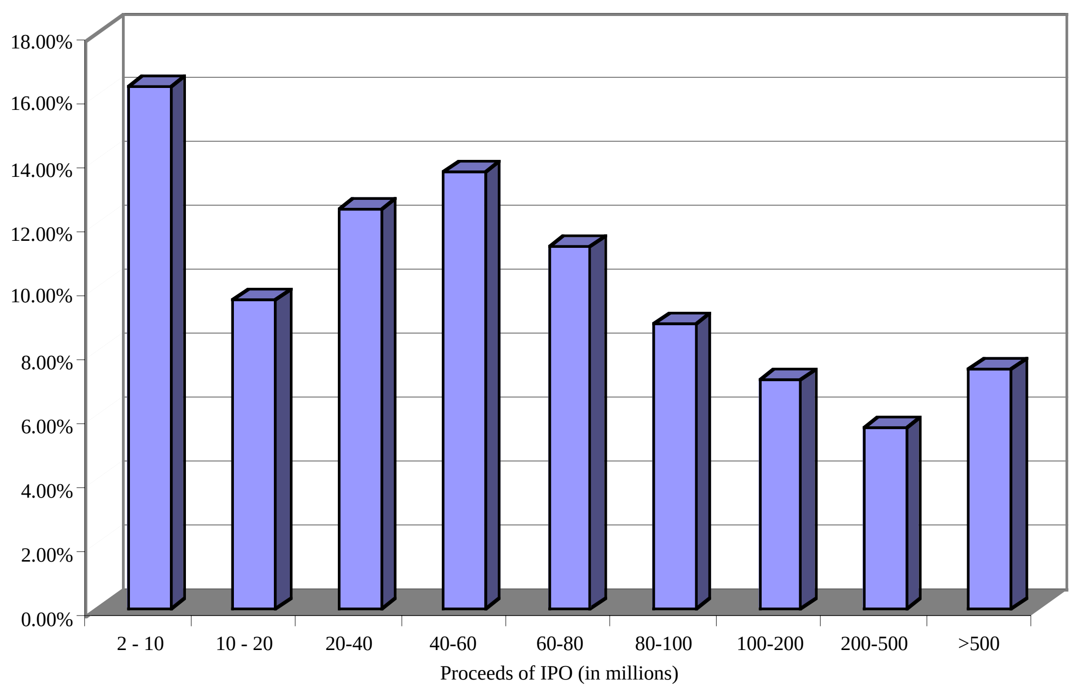

Source: Lee, Lockhead, Ritter and Zhao

[11] See Lee, Lockhead, Ritter and Zhao (1996)

The smaller the issue, the greater the underpricing – the smallest offerings often are underpriced by more than 17% but the underpricing is much smaller for the larger issues.

Some of the studies have broken down initial public offerings on other dimensions to examine the reasons for the underpricing. In 1998, Ritter provided a comprehensive summary of both they hypotheses on why the under pricing occurs and the empirical evidence on it. We summarize a few of his findings:

- The average initial return is 15.8% across a sample of 13,308 initial public offerings. However, about 15% of all initial public offerings are over priced. In other words, the stock price drops from the initial offering price on the date of the offering. Thus, investing in IPOs is by no means a riskless or guaranteed strategy, even if you were guaranteed an allotment in every one at the offering price.
- Initial public offerings where the offering price is revised upwards prior to the offering are more likely to be under priced than initial public offerings where the offering price is revised downwards. Table 9.1 below contrasts the initial returns and the percent of offerings that were under priced for both classes from 1991 to 1996.

Table 9.1: Average Initial Return – Offering Price Revision

Offering price	*Number of IPOs*	*Average initial return*	*% of offerings underpriced*
Revised down	708	3.54%	53%
Revised up	642	30.22%	95%

While the evidence that initial public offerings go up on the offering date is strong, it is not clear that these stocks are good investments in the years after. Loughran and Ritter tracked returns on 5821 IPOs in the five years after the offerings and contrasted them with returns in figure 9.10.

Figure 9.10: Post Issue Returns - IPOs versus Non IPOs

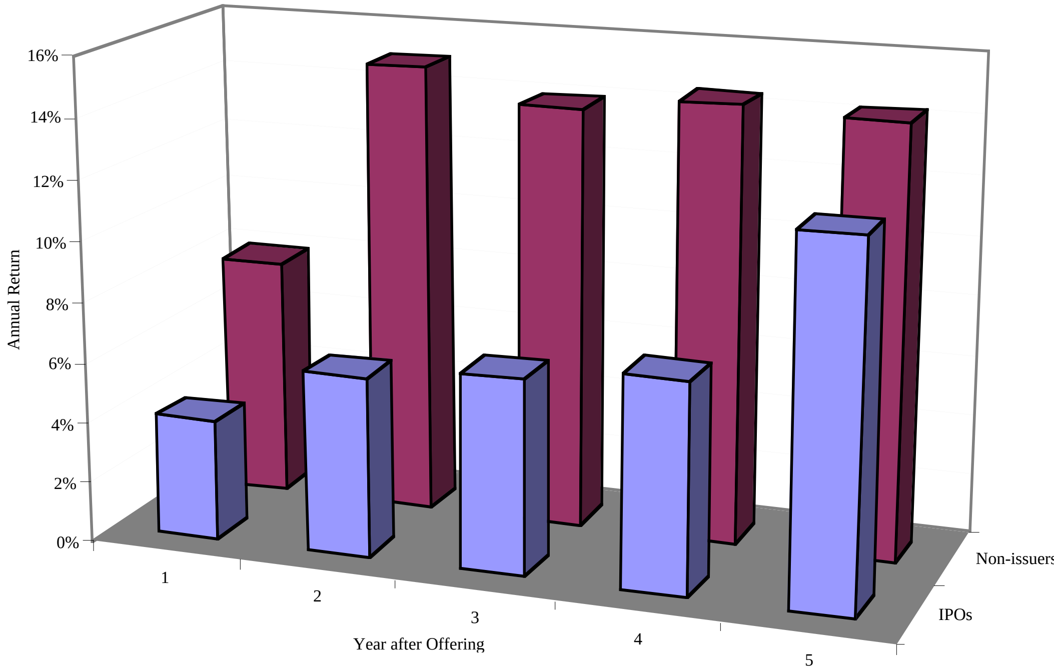

Note that the IPO firms consistently under perform the non-issuing firms and that the under performance is greatest in the first few years after the offering. While this phenomenon is less pronounced for larger initial public offerings, it still persists.

The Allotment Process

If initial public offerings, on average, are under priced, an obvious investment strategy is to subscribe a large number of initial public offerings and to construct a portfolio based upon allotments of these offerings. There is, however, a catch in the allotment process that may prevent this portfolio from earning the excess returns from the under pricing. When investors subscribe to initial public offerings, the number of shares that they are allotted will depend upon whether and by how much the offering is under priced. If it is significantly under priced, you will get only a fraction of the shares that you requested. On the other hand, if the offering is correctly priced or over priced, you will get all of the shares that you requested. Thus, your portfolio will be underweighted in under priced initial public offering and overweighted in overpriced offerings.

Is there a way in which you can win this allotment game? There are two. The first is to be the beneficiary of a biased allotment system, where the investment bank gives you more than your share of your requested shares in under priced offerings. While this is illegal in the United States[12], it is legal in many other countries in the world. The second and more legitimate way is to develop an analytical system that allows you to separate under priced from over priced offerings, using public information contained in the prospectus and other SEC filings. You would then request shares in only those offerings that you identified as under priced. If you are reasonably accurate, you should end up with a portfolio that more closely resembles (or even beats) the hypothetical portfolios created across all initial public offerings.

The IPO Cycle

Initial public offerings ebb and flow with the overall market. There are periods where the market is flooded with initial public offerings and periods where there are almost no offerings. Contrast, for instance, the salad days of the late 1990s, when firms went public at an extraordinary pace and 2001, when the number slowed to a trickle. In addition, the initial public offerings during any period tend to share a common sector focus. For instance, the bulk of the initial public offerings during 1999 were of young technology and telecomm firms. This does create two problems for investment strategies that focus exclusively on initial public offerings. The first is that your portfolio will not be diversified in periods of plenty, and will be over weighted in whichever sector is in favor at that point in time. The second is that there will be extended periods where you will find nothing to invest in, because there are few or no initial public offerings.

Ritter (1998) provides a summary of the number of offerings made each year from 1960 to 1996 and the average initial returns on those offerings. His results are summarized in figure 9.11:

[12] Notwithstanding restrictions on this practice, investment banks in the 1990s used allotments in initial public offerings as a lead into other business with clients. Thus, large portfolio managers often were given more than their fair share of initial public offerings that were in demand.

Figure 9.11: Number of IPOs and Average Initial Return

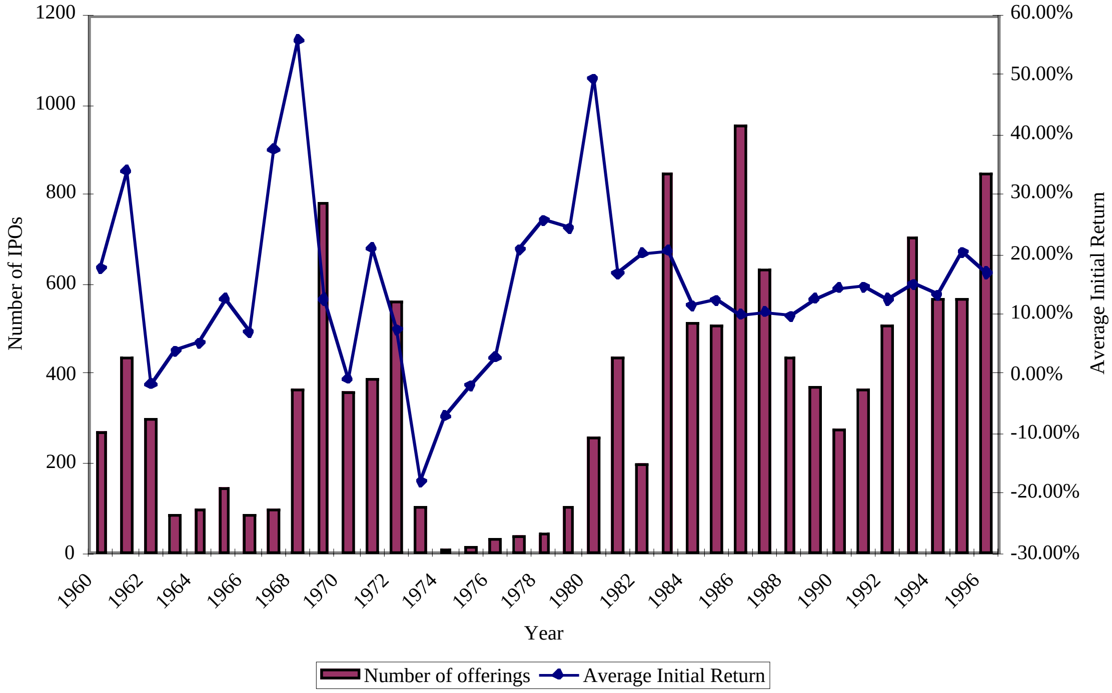

Note that the number of offerings drops to almost zero in the early 1970s and the returns to IPOs drops as well. A portfolio manager who focused only on initial public offerings would have gone out of business in that period.

Determinants of Success

A strategy of investing in initial public offerings makes more sense as an ancillary strategy rather than a primary strategy, partly because the sector concentration of initial public offerings during hot periods and partly because of the absence of offerings during cold periods. Assuming that it is used as an ancillary strategy, you would need to do the following to succeed:

- Have the valuation skills to value companies with limited information and considerable uncertainty about the future, so as to be able to identify the companies that are under or over priced.
- Since this is a short term strategy, often involving getting the shares at the offering price and flipping the shares on the offering date, you will have to gauge the market mood and demand for each offering, in addition to assessing its value. In other words, a shift in market mood can leave you with a large allotment of over-priced shares in an initial public offering.

- Play the allotment game well, asking for more shares than you want in companies which you view as severely under priced and fewer or no shares in firms that are overpriced or that are priced closer to fair value.

In recent years, investment banks have used and misused the allotment process to reward selected clients. In periods when demand for initial public offerings is high, they have also been able to punish investors who sell immediately by withholding or rationing future allotments. If you are required to hold these stocks for the long term to qualify for the initial offering, you may very well find that the under performance of these stocks in the post-offering period (see figure 9.10) can very quickly be decimated by poor returns in subsequent periods.

Growth Screens

If you were a portfolio manager whose choices come from a very large universe of stocks, your most effective way of building a portfolio may be to screen stocks and pick those that pass specific screens. In other words, you do for growth stocks what Ben Graham did for value stocks. In this section, we consider three screening strategies – a strategy of buying stocks with high expected growth rates in earnings, the high flyer strategy, where you pick stocks with high PE ratios and the growth at a reasonable price (GARP) strategy, where you pick growth stocks that trade at low prices, given their expected growth.

High Earnings Growth Strategy

The strategy that follows most logically for most growth investors is to buy stocks with high growth rates in earnings. You can look at past growth in earnings as a predictor of future growth and buy companies with high historical earnings growth rates or you can look for companies where analysts are predicting high expected earnings growth.

Historical Growth

Is the growth rate in the past a good indicator of growth in the future? Not necessarily. Past growth rates are useful in forecasting future growth, but there are two problems.

- The first is that they have considerable noise associated with them and are noisy predictors of future growth. In a 1960 study of the relationship between past growth rates and future growth rates, Little coined the term "Higgledy Piggledy Growth" because he found little evidence that firms that grew fast in one period continued to grow fast in the next period. In the process of running a series of correlations between growth rates in earnings in consecutive periods of different length, he frequently found negative correlations between growth rates in the two periods and the average correlation across the two periods was close to zero (0.02). If past growth in earnings is not a reliable indicator of future growth at average firms, it becomes even less so at smaller firms. The growth rates at smaller firms tend to be even more volatile than growth rates at other firms in the market. The correlation between growth rates in earnings in consecutive time periods (five-year, three-year and one-year) for firms in the United States, categorized by market value, is reported in Figure 9.12.

Figure 9.12: Correlations in Earnings Growth by Market Capitalization

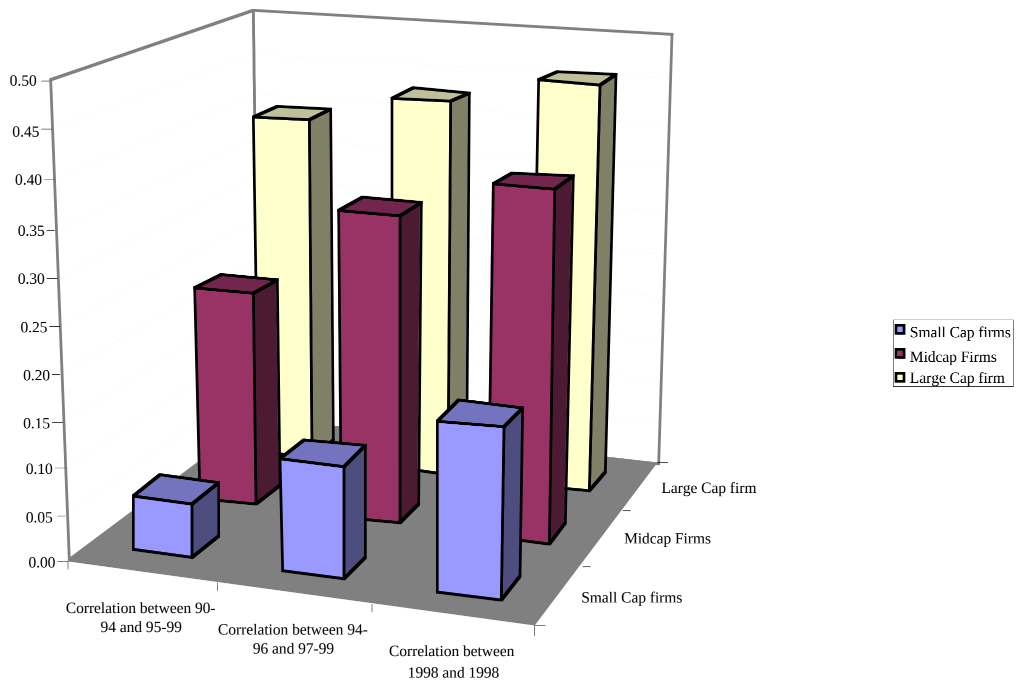

While the correlations tend to be higher across the board for one-year growth rates than for 3-year or 5-year growth rates in earnings, they are also consistently lower for smaller firms than they are for the rest of the market. This would suggest that you should be more cautious about using past growth, especially in earnings, for forecasting future growth at these firms.

- The second problem is that there is mean reversion in earnings growth rates. In other words, companies that are growing fast will see their growth rates decline towards the market average whereas below average growth companies will see their growth rates increase. This tendency is chronicled by Dreman and Lufkin when they track companies in the highest and lowest earnings growth classes for 5 years after the portfolios are formed. While the highest earnings growth companies have an average growth rate which is 20% higher than the average growth rate for the lowest earnings growth companies in the year the portfolio is formed, the difference is close to zero five years later.

- In general, revenue growth tends to be more persistent and predictable than earnings growth. This is because accounting choices have a far smaller effect on revenues than they do on earnings. Figure 9.13 compares the correlations in revenue and earnings growth over one-year, three-year and five-year periods at U.S. firms.

Figure 9.13: Correlation in Revenues and Earnings

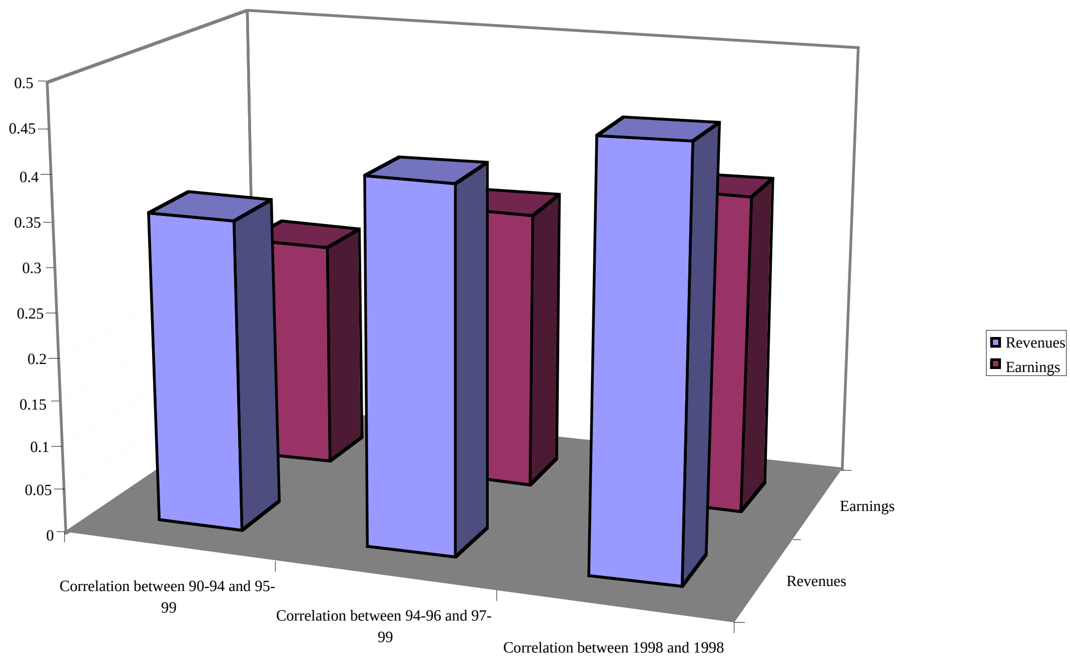

Revenue growth is consistently more correlated over time than earnings growth. The implication is that historical growth in revenues is a far more useful number when it comes to forecasting than historical growth in earnings.

There are some investors who believe that it is not earning growth per se that you should be looking at but momentum in growth. In other words, you want to invest in stocks whose earnings growth is accelerating. This is, in fact, a big component of what Value Line's acclaimed stock picking measures are based upon. While Value Line may have been successful with this strategy in its earlier years, much of what we have said about earnings growth probably also applied to earnings momentum.

In summary, past earnings growth is not a reliable indicator of future growth and investing in companies with high past growth does not yield significant returns. In fact, if there is mean reversion and you pay a large premium for companies with high growth, you will find yourself with a losing portfolio.

Expected Earnings Growth

Value is ultimately driven by future growth and not past growth. It seems reasonable, therefore, that you would be better served investing in stocks where expected growth is high rather than historical growth. Here, you do run into a practical problem. In a market as large as the United States, you cannot estimate expected growth for each firm in the market. Instead, you have to rely on analyst estimates of expected growth. That information, though, is freely accessible now to most investors and you could buy stocks with high expected growth rates in earnings. But will such a strategy generate excess returns?

Stocks with highest expected growth: Take a look at 50 stocks with the highest expected growth in earnings per share.

Consider what you would need for this strategy to be successful. First, analysts have to be proficient at forecasting long term earnings growth. Second, the market price should not already reflect or overprice this growth. If it does, your portfolio of high growth companies will not generate excess returns. On both conditions, the evidence works against the strategy. When it comes to forecasting growth, analysts have a tendency to overestimate growth and the forecast errors are high for long-term forecasts. In fact, some studies find that time series model match or even outperform analysts when it comes to long term growth. As for pricing growth, markets historically have been more likely to over price growth than under price it, especially during periods of high earnings growth for the market.

High PE Strategy

The easiest growth strategy, albeit the riskiest, is to buy the stocks with the highest PE ratios on the market, on the assumption that these are growth companies where the growth will deliver the excess returns in the future.

The Overall Evidence

We should begin by noting that the overall evidence on buying high PE ratio stocks is grim. As we noted in chapter 8, when looking at the value stocks, buying low PE ratio stocks seems to outperform high PE ratio stocks by significant margins. Figure 9.14 presents the difference in annual returns from buying low PE stock and high PE stock portfolio from 1952 to 2001.

Figure 9.14: PE Ratios and Stock Returns - 1952 -2001

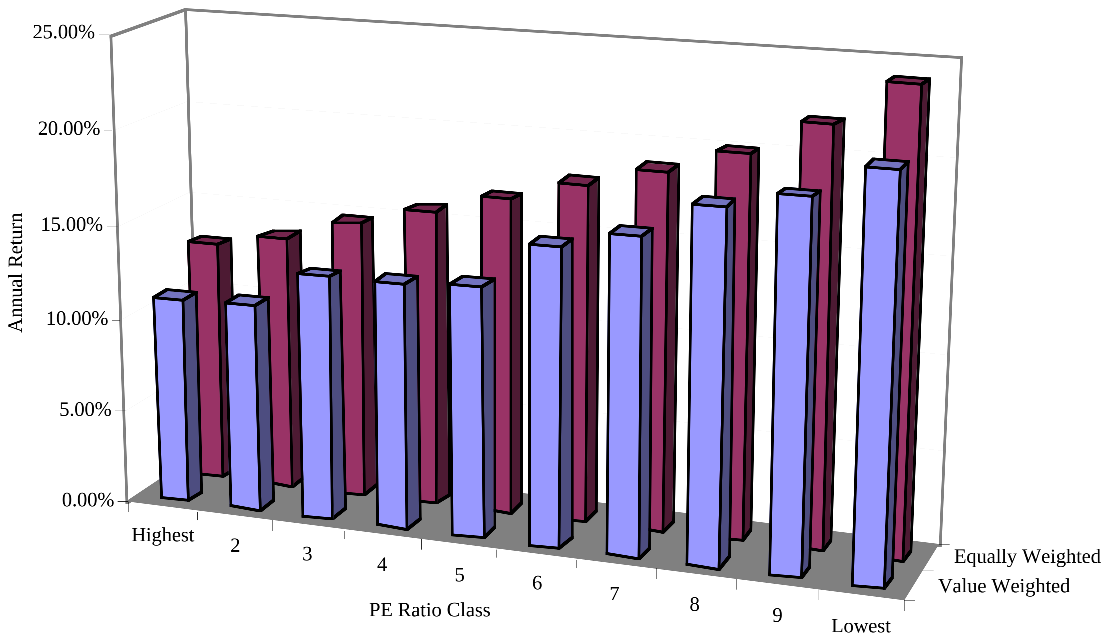

On both an equally-weighted and a value-weighted basis, high PE stocks have under performed low PE ratio stocks. In fact, it is this consistent under performance of high PE stocks that has led to the value investing bias that we often see in both academic and practitioner research.

The Growth Investors' Case

Given this sorry performance, what you might wonder attracts investors to this strategy? The answer lies in cycles. There have been extended time periods where high PE stocks seem to outperform low PE stocks. For instance, growth investing seems to do much better when the earnings growth in the market is low and value investing tends to do much better when earnings growth is high. In figure 9.15, we have graphed the difference between a low PE and a high PE portfolio and the growth in earnings in each period:

Figure 9.15: Relative Performance of Growth and Value versus Earnings Growth

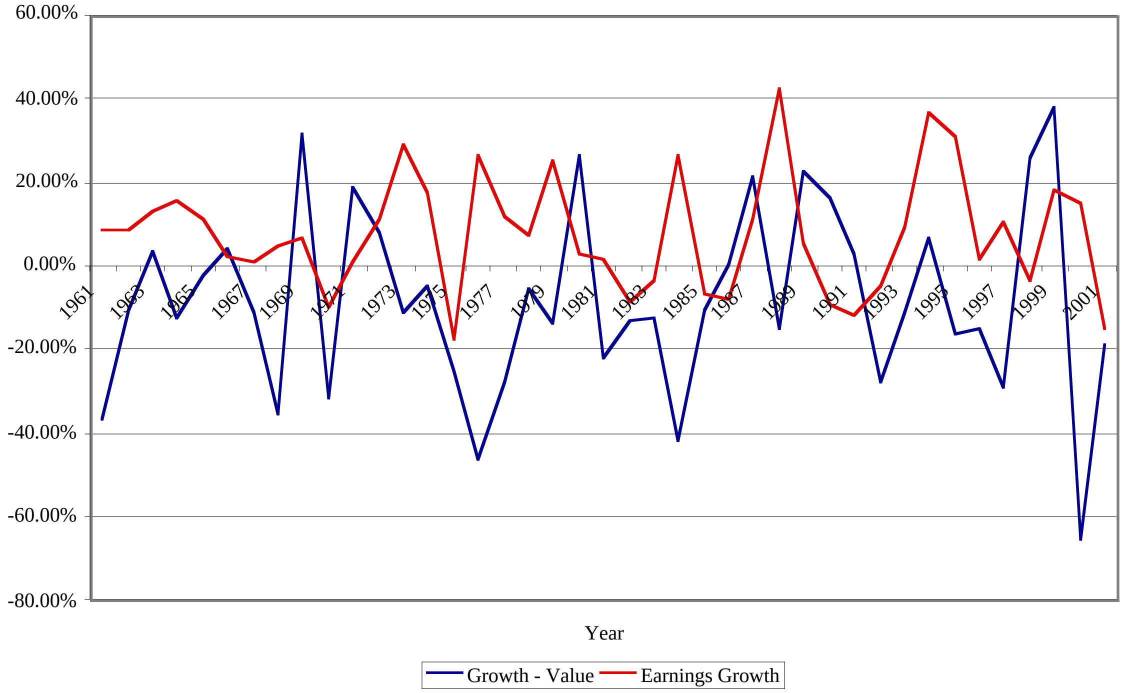

We measure the performance of growth versus value by looking at the difference between the returns earned on a portfolio of stocks in the top decile in terms of PE (growth) and a portfolio of stocks in the lowest decile (value). Thus, a positive value indicates that high PE stocks outperformed low PE stocks in that year. Growth investing does best in years when earnings growth is low. This may be due to the fact that growth stocks are more desirable in periods when earnings growth is low, because they are scarce. By the same token, when all companies are reporting high earnings growth, investors seem to be unwilling to pay a premium for growth.

Growth investing also seems to do much better when the yield curve is flat or downward sloping and value investing does much better when is much more upward sloping. Figure 9.16 presents the relationship between the slope of the yield curve and the performance of growth investing.

Figure 9.16: Relative Performance of Growth Stocks versus Yield Curve

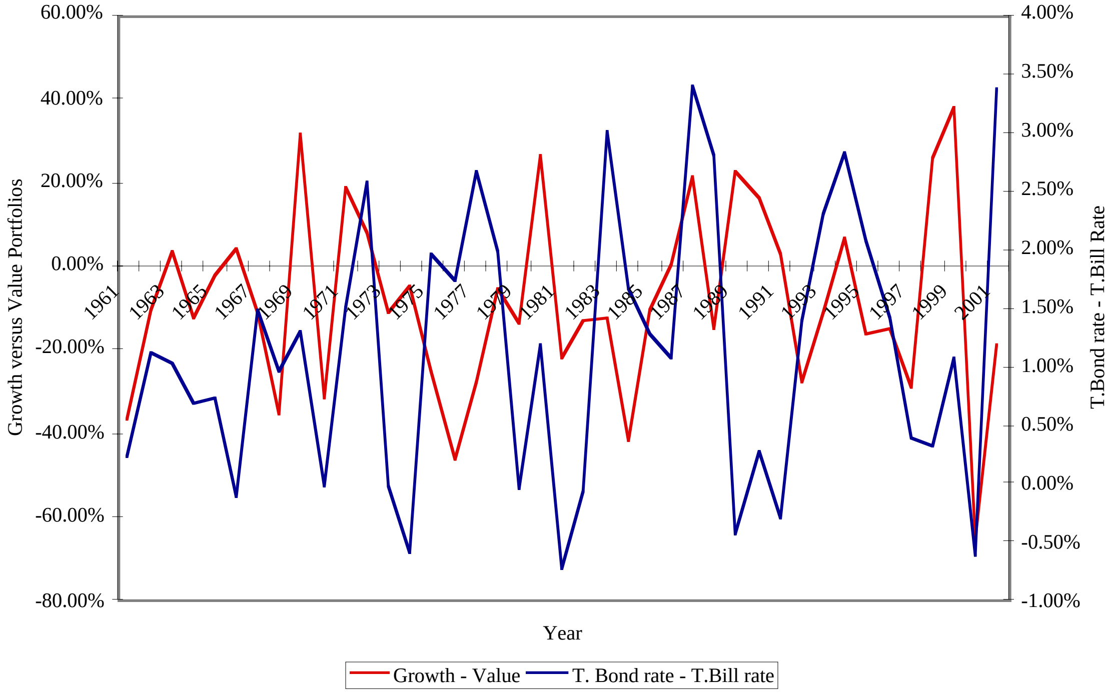

The most interesting evidence on growth investing, however, lies in the percent of active money managers who beat their respective indices. When measured against their respective indices, active growth investors seem to beat growth indices more often than active value investors beat value indices. In his paper on mutual funds in 1995, Malkiel provides additional evidence on this phenomenon. He notes that between 1981 and 1995, the average actively managed value fund outperformed the average actively managed growth fund by only 16 basis points a year, while the value index outperformed a growth index by 47 basis points a year. He attributes the 32 basis point difference to the contribution of active growth managers, relative to value managers. We will look at this evidence in more detail in chapter 13.

Peter Lynch: Finding value in growth stocks

If Warren Buffett is the icon for value investors, Peter Lynch occupies a similar position for growth investors. His reputation was made during his stewardship of Fidelity Magellan, a small high growth fund that he took over in 1977 and made into the largest equity mutual fund in the world over the next decade. The reason for its growth was its performance. An investment of $ 10,000 in the Magellan fund would have grown 20 fold over the next ten years. During that period, Lynch also helped dispel the notion that growth investors were incurable optimists who bought stocks on promises. He introduced the rigors of value investing to growth investing, and he described much of what he did in his books on investing and his articles for Worth, a financial magazine. One of these articles includes Lynch's maxims on finding good investments:

1. Pay attention to facts and not forecasts.
2. Before you invest, check the balance sheet to see if the company is financially sound.
3. Don't buy options, and don't invest on margin. With options, time works against you, and if you're on margin, a drop in the market can wipe you out.
4. When several insiders are buying the company's stock at the same time, it's a positive.
4. Average investors should be able to monitor five to ten companies at a time, but nobody is forcing you to own any of them.
5. Be patient. Stocks often make their greatest gains in the third or fourth year that you own them. A few took ten years.
6. Enter early -- but not too early. Think of investing in growth companies in terms of baseball. Try to join the game in the third inning because a company has proved itself by then. If you buy before the lineup is announced, you're taking an unnecessary risk. If you buy in the late innings, you may be too late.
7. Don't buy "cheap" stocks just because they're cheap. Buy them because the fundamentals are improving.
8. Buy small companies after they've had a chance to prove they can make a profit.
9. Long shots usually backfire or become "no shots."
10. Investigate ten companies and you're likely to find one with bright prospects that aren't reflected in the price.

Worth Magazine, 1996.

GARP Strategies

There are many growth investors who would blanch at the strategy of buying high PE stocks. Their mission, they would argue, is to buy high growth stocks where growth is undervalued. To find these stocks, they have developed a number of strategies where you consider both expected growth and the current pricing of the stock. We will consider two of these strategies in this section – buying stocks with a PE less than the expected growth rate or buying stocks with a low ratio of PE to growth (called a PEG ratio).

PE less than Growth Rate

The simplest GARP strategy is to buy stocks that trade at a PE ratio less than the expected growth rate. Thus, a stock that has a PE ratio of 12 and an expected growth rate of 8% would be viewed as undervalued, whereas a stock with a PE of 40 and an expected growth rate of 50% would be viewed as undervalued. While this strategy clearly has the benefit of simplicity, it can be dangerous for several reasons.

- *Interest rate effect*: Since growth creates earnings in the future, the value of growth is a present value. The value created by any given growth rate will be greater when interest rates are low (which makes the present values higher) than when interest rates are high. Thus, the stock with a PE of 40 and an expected growth rate of 50% when interest rates are 7% may find itself with a PE of 60 if interest rates drop to 5% but growth remains unchanged. It is not surprising, therefore, that portfolio managers who use this strategy not only find far more attractive stocks when interest rates are high but also find many emerging market stocks (where interest rates tend to be higher) bargains. The effect on interest rats on the relationship between PE and growth can be best illustrated by looking at the percent of firms that trade at less than their expected growth rate as a function of the treasury bond rate. In 1981, when treasury bond rates hit 12%, more than 65% of firms traded at PE ratios less than the expected growth rate. In 1991, when rates had dropped to about 8%, the percent of stocks trading at less than the expected growth rate also dropped to about 45%. By the end of the nineties, with the treasury bond rate dropping to 5%, the percent of stocks that traded at less than the expected growth rate had dropped to about 25%.
- *Growth Rate Estimates*: When this strategy is used for a large number of stocks, you have no choice but to use the growth rate estimates of others. In some cases, the consensus growth rates estimated by all analysts following a firm are obtained from a data service and used. When you do this, you have to wonder both about the differences in the quality of the growth estimates across different analysts and the comparability. Given that these estimated growth rates are at most for five years, you may penalize companies that have expected growth for much longer periods by focusing just on the 5-year rate.

It is also possible that in low interest rate scenarios, very few stocks pass this screen and that you will end up with little to invest in.

PEG Ratios

An alternative approach that seems to offer more flexibility than just comparing the PE ratio to expected growth rates is to look at the ratio of the PE ratio to expected growth. This ratio is called the PEG ratio and is widely used by analysts and portfolio managers following growth companies.

Defining the PEG Ratio

The PEG ratio is defined to be the price earnings ratio divided by the expected growth rate in earnings per share:

$$\text{PEG ratio} = \frac{\text{PE ratio}}{\text{Expected Growth Rate}}$$

For instance, a firm with a PE ratio of 40 and a growth rate of 50% is estimated to have a PEG ratio of 0.80. There are some who argue that only stocks with PEG ratios less than one are desirable, but this strategy is equivalent to the strategy of comparing the PE to the expected growth rate.

Stocks with lowest PEG ratios: Take a look at the 50 stocks with the lowest PEG ratios.

Consistency requires the growth rate used in this estimate be the growth rate in earnings per share. Given the many definitions of the PE ratio, which one should you use to estimate the PEG ratio? The answer depends upon the base on which the expected growth rate is computed. If the expected growth rate in earnings per share is based upon earnings in the most recent year (current earnings), the PE ratio that should be used is the current PE ratio. If it based upon trailing earnings, the PE ratio used should be the trailing PE ratio. The forward PE ratio should generally not be used in this computation, since it may result in a double counting of growth.[13] Building upon the theme of uniformity, the PEG ratio should be estimated using the same growth estimates for all firms in the sample. You should not, for instance, use 5-year growth rates for some firms and 1-year growth rates for others. One way of ensuring uniformity is to use the same source for earnings growth estimates for all the firms in the group. For instance, both I/B/E/S and Zacks provide consensus estimates from analysts of earnings per share growth over the next 5 years for most U.S. firms. Many analysts who use PEG ratios, though, prefer to use short-term growth rates in earnings to compute them.

Using the PEG Ratio

How do analysts use PEG ratios? A stock with a low PEG ratio is considered cheap, because you are paying less for the growth. It is viewed as a growth neutral measure that can be used to compare stocks with different expected growth rates. In a study concluded in 1998, Morgan Stanley found that a strategy of buying stocks with low PEG ratios yielded returns that were significantly higher than what you would have made on the S&P 500. They came to this conclusion by looking at the 1000 largest stocks on the U.S. and

---

[13] If the forward earnings are high because of high growth in the next year, and this high growth results in a high growth rate for the next 5 years, you will understate your PEG ratio.

Canadian exchanges each year from January 1986 through March 1998, and categorizing them into deciles based upon the PEG ratio. They found that the 100 stocks with the lowest PEG ratio earned an annual return of 18.7% during the period, much higher than the market return of about 16.8% over the period. While no mention was made of risk adjustment, it was argued that the difference was larger than could be justified by the risk adjustment.

We updated this study to examine how this strategy would have done from 1991 to 2001, creating five portfolios at the end of each year based upon the PEG ratio and examining the returns in the following year. Figure 9.17 summarizes the average annual returns on PEG ratios classes in the 1991-1996 and 1997-2001 time periods.

Figure 9.17: PEG Ratios and Annual Returns

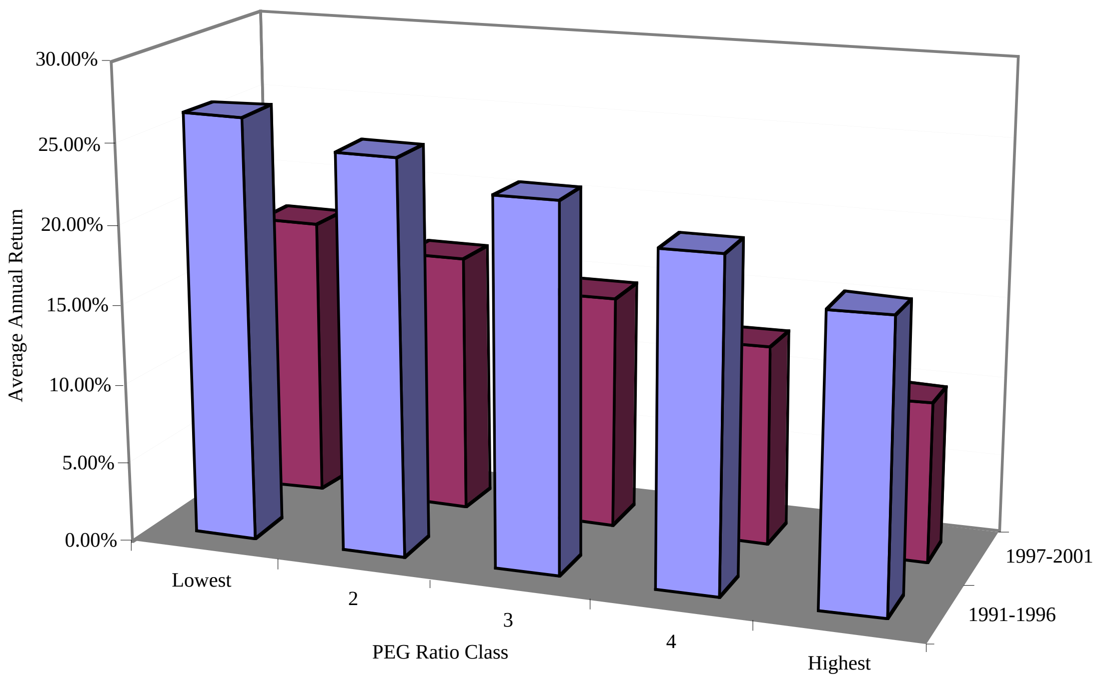

A strategy of investing in low PEG ratio stocks would have generated an average return about 3% higher than the average returns on a high PEG ratio portfolio, before adjusting for risk, during both time periods.

Potential Problems

If, on average, stocks with low PEG ratios outperform other stocks, why should we not adopt this as a screening strategy? There are two potential problems with PEG ratios that may lead us to misidentify riskier stocks with higher growth rates as undervalued.

The first and most obvious problem is that the PEG ratio is obtained by dividing the PE ratio by the expected growth rate and the uncertainty about that expected growth rate is not factored into the number. Intuitively, you would expect, riskier stocks for any given growth rate, to have lower PE ratios. Thus, a stock that looks cheap on a PEG ratio basis may be, in fact, correctly or even over valued. The relationship between risk and growth can be illustrated in two ways. The first is by computing the PE ratio for a hypothetical firm, holding growth and cashflows constant, but varying the risk.[14] In figure 9.18 below, for instance, we vary the beta of a stock with an expected growth rate of 25% for five years and 8% forever therafter, and compute the PEG ratio:

Figure 9.18: PEG Ratios and Beta: Firm with 25% growth for next 5 years, 8% thereafter

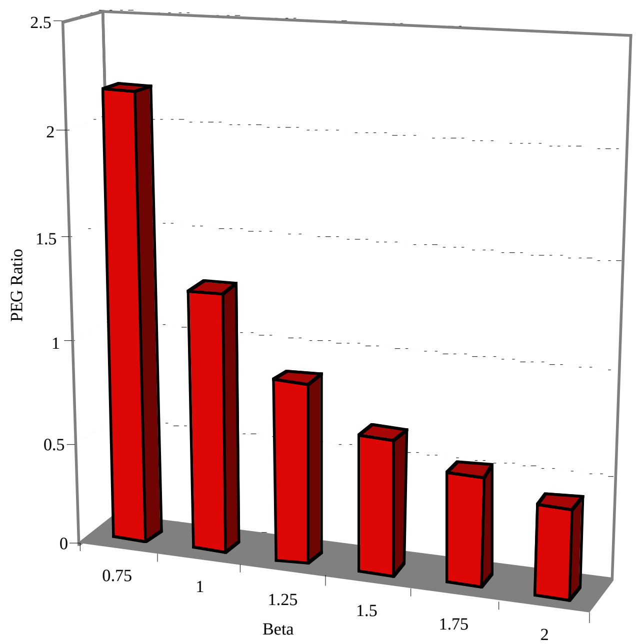

Note that the PEG ratio for the firm with a beta of 0.75 is almost four times higher than the PEG ratio for the same firm (with the same growth rate) with a beta of 2.00. You can also see the relationship between risk and PEG ratios by computing the average PEG ratios for all stocks listed in the United States and categorizing them based upon their riskiness.

---

[14] To do this, you first have to compute the PE ratio based upon fundamentals and then divide by the expected growth rate. A more detailed exposition is provided in my book on Investment Valuation, but the PEG ratio in a two-stage dividend discount model can be written as

$$\text{PEG} = \frac{(\text{Payout Ratio})(1+g)\left(1-\dfrac{(1+g)^n}{(1+k_{e,hg})^n}\right)}{g(k_{e,hg} - g)} + \frac{(\text{Payout Ratio}_n)(1+g)^n(1+g_n)}{g(k_{e,st} - g_n)(1+k_{e,hg})^n}$$

Figure 9.19 classifies all firms in the United States into six risk classes[15] and computes the average PE ratios and PEG ratios for firms in each class in January 2002.

Figure 9.19: PE and PEG ratios by Risk Class- January 2002

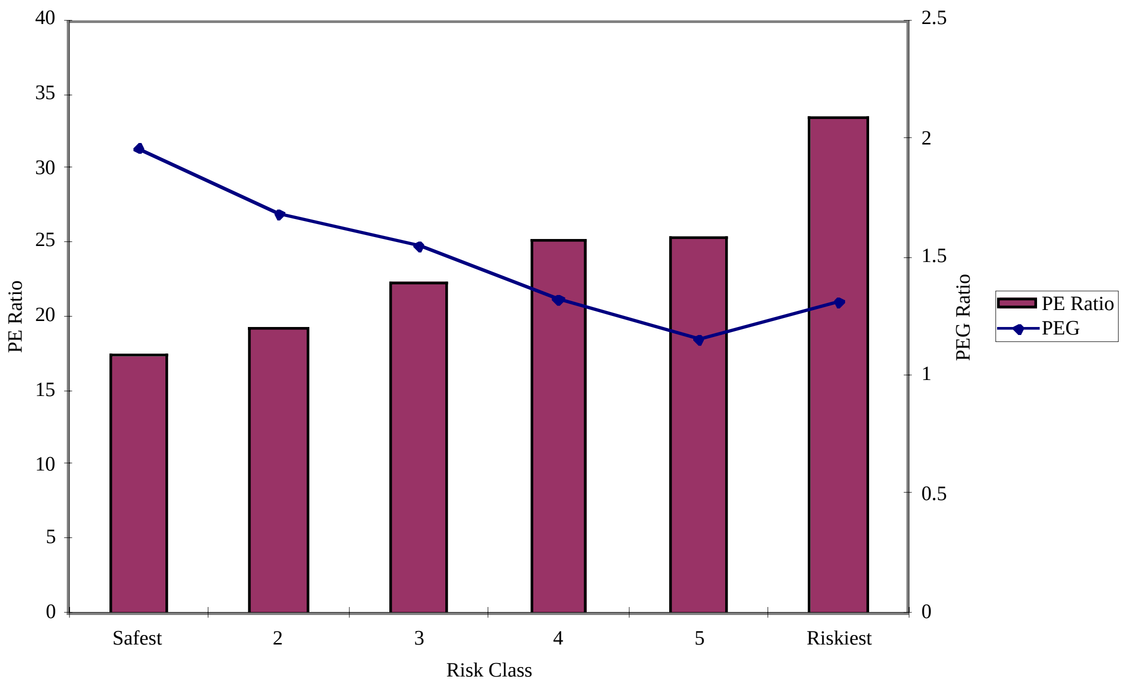

While the highest risk firms have higher PE ratios than the safer firms, on average, they also have lower PEG ratios. Thus, a portfolio of the stocks with the lowest PEG ratios will tend to include a large number of high risk stocks.

> *Stocks with low PEG ratios and below average risk:* Take a look at the 50 stocks with the lowest PEG ratios and below-average risk.

The second potential problem with PEG ratios is less obvious but just as dangerous. When we use PEG ratios, we make the implicit assumption that as growth doubles, the PE ratios doubles, and if it is halved, the PE ratio will be halved as well. In other words, we assume a linear relationship between PE and expected growth and this clearly is not correct. To see why, consider what should happen to the PE as expected growth drops to zero. If you have a firm that has a dollar in earnings that it pays out in dividends and you expect to get this dollar in dividend in perpetuity, you would still be willing to pay a price for its stock. In other words, your PE does not go to zero. On the other side, you will find that PE ratios

---

[15] This categorization was based upon stock price standard deviation, but we did try alternate measures such as beta and obtained similar results.

increase as you increase the expected growth rate but at a decreasing rate. In other words, your PE ratio will change much more dramatically when your expected growth rate goes from 3 to 4% than when it goes from 23 to 24%. Again, the effect on PEG ratio of varying the growth rate can be shown in one of two ways. Using the same process that we used to examine the relationship between PEG ratios and risk, we can estimate the PEG ratio for a hypothetical firm in figure 9.20 as you change the expected growth rate during its high growth phase.

Figure 9.20: PEG Ratios and Expected Growth Rate over next 5 years

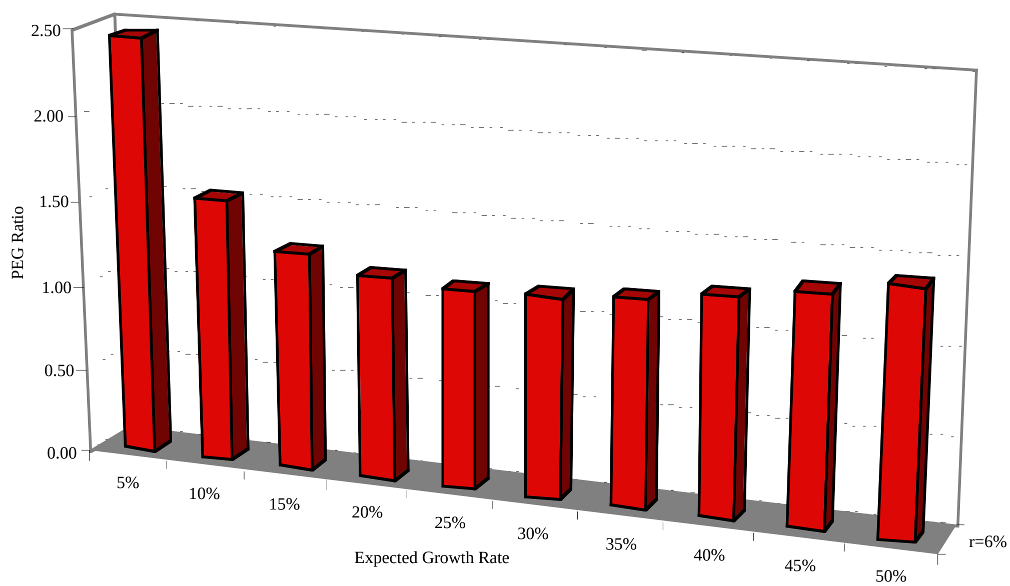

The PEG ratio is highest when the expected growth rate is low, but is lower at higher expected growth rates. Clearly, the problem is greatest when you are comparing high growth firms to low growth firms, since PEG ratios will be understated for the former and overstated for the latter. It is less of an issue if you are comparing PEG ratios across firms with high growth rates, since the effect is muted.

In short, picking stocks based upon low PEG ratios can leave you with a portfolio of stocks with high risk and high growth that are not undervalued. Can you correct for these errors? You can adjust for risk by either considering it as a separate factor (you pick stocks with low PEG ratios and low risk) or modifying the PEG ratio. Morgan Stanly, for example, aware of the potential bias towards risk in the PEG ratio modified it to include the dividend yield in the denominator to create a new ratio called the PEGY ratio:

$$PEGY = \frac{PE}{(\text{Expected Growth Rate} + \text{Dividend Yield})}$$

Thus, a firm with a PE ratio of 12, an expected growth rate of 5% and a dividend yield of 4% would have a PEGY ratio of 1.33 (12/(5+4)). It is much more difficult to adjust for the linearity assumption. While you can use a modified version of the growth rate in the denominator, the measure loses its intuitive appeal when this is done.[16]

Determinants of Success at Passive Growth Investing

The overall empirical evidence on the efficacy of screens is much less favorable for growth screens than it is for value screens. While there are cycles during which growth screens like low PEG ratios and high PE ratios may yield excess returns, they are trumped over longer periods by value screens such as price to book value ratios. From our perspective, there are three key determinants of success at this strategy:

- Since growth is the key dimension of value in these companies, obtaining better estimates of expected growth should improve your odds of success. If you are a growth investor following a fairly small set of companies, you may try to estimate growth yourself. If you can estimate growth more precisely than the overall market, you should get a payoff. If this is not a feasible option because you do not have the resources to estimate expected growth rates for the hundreds of firms that you follow, you should compare the different sources that you have for this input to see which one has the best track record. For instance, you may find that Value Line estimates of growth are better than the consensus estimates of growth from all analysts or that estimates of growth from a sub-set of analysts (say the top 5) do better than estimates that look at all analysts.
- If your underlying strategy is sound, a long time horizon increases your chances of earning excess returns. In other words, if you conclude after careful analysis that buying stocks that have PE ratios less than the expected growth rate would have yielded high returns over the last two decades, you will be more likely to replicate these results if you have a 5-year horizon than with a 1-year horizon.
- Finally, there are extended cycles where the growth screens work exceptionally well and other cycles where they are counter productive. If you can time these cycles, you could augment your returns substantially. Since many of these cycles are related to how the overall market is doing, this boils down to your market timing ability.

---

[16] For example, using the natural log of the growth rate in the denominator of the PEG ratio seems to make the relationship more linear, but the PEG ratio is no longer intuitive.

Being successful at these strategies will require not only long time horizons but also the capacity to be right on market cycles.

Estimating Growth from Fundamentals

If obtaining better estimates of growth is key to successful growth investing, you may want to consider breaking your dependence on estimates of growth made by equity research analysts. As we will see in the coming chapters, analysts often do not estimate long term growth and even when they do, they provide biased and erroneous estimates. One alternative that may yield better and more robust estimates is to link the expected growth to fundamental aspects of how a firm is run. In fact, the expected growth rate in earnings for a firm comes from two sources - its willingness to reinvest its earnings back into new projects and assets and its capacity to earn high returns on these investments. The growth rate in earnings for a firm in the long term should be a product of the proportion of its earnings that are reinvested back in the business and the return on this investment. For equity earnings, it can be computed as follows;

Growth rate in earnings per share = (1 – Dividends/ Earnings) (Return on Equity)

Consider, for instance, a company like Microsoft, which pays no dividends and earns about 25% on its equity. Its expected growth rate, if it can sustain these numbers, will be 25%. In contrast, Procter and Gamble, which pays about 50% of its earnings as dividends and earns about 16% on its equity will have an expected growth rate of 8%.

When computing growth in operating earnings, you will have to modify the equation to make it consistent:

Growth rate in operating earnings = [(Capital Expenditures – Depreciation + Change in Working capital)/ EBIT (1-t)]* Return on Capital

Thus, Cisco which reinvested 110% of its after-tax operating income and earned a return on capital of 35% in 1999 was able to post a growth rate of 38.5% in that year.

Activist Growth Investing

In activist growth investing, you not only take a position in a growth business but you also play an active role in making it successful. Since most growth businesses start off as small and privately owned, the most common forms of activist growth investing involve taking positions in these businesses before they go public and in nurturing them towards eventual public offerings and large profits. In this section, we will consider venture capital and private equity investing as examples of activist growth investing.

Description

In venture capital investing, you provide equity financing to small and often risky businesses in return for a share of the ownership of the firm. The size of your ownership share will depend upon two factors. First, at the minimum, you will demand an ownership share based upon how much capital you contribute to the firm, relative to total firm value. For instance, if you provide $ 2 million and the estimated value of the firm is $10 million, you will expect to own at least 20% of the firm. Second, if the business can raise the funds from other sources, its bargaining position will be stronger, and it may be able to reduce your share down to a small premium over the minimum specified above. If a business has no other options available to raise the equity financing, however, its bargaining position is considerably weaker, and the owner of the business will have to give up a disproportionate share of the ownership to get the required funding. In general, the capacity to raise funds from alternative sources or to go public will increase with the size of the firm and decrease with the uncertainty about its future prospects. Thus, smaller and riskier businesses are more likely to seek venture capital and are also more likely to be asked to give up a greater share of the value of the firm when receiving the venture capital.

The Market for Private Equity and Venture Capital

Until a few decades ago, venture capital was provided by a relatively small number of individuals. They tended to specialize in a sector, invest in relatively few firms and take an active role in the operations of these firms. In recent decades, though, as the market for venture capital has increased, you have seen three categories emerge.

The first are venture capital funds that trace their lineage back to the 1950s. One of the first was American Research and Development that provided seed money for the founding of Digital Equipment. During the 1960s and 1970s, these funds multiplied and helped start and expand companies such as Intel and Apple that were then taken public. The second are leveraged buyout funds that developed during the 1980s, using substantial amounts of debt to take over publicly traded firms and make them private firms. The publicity they generated – positive as well as negative – in the form of personalities, books and movies helped shaped the public's view of all acquisitions for a generation.[17] More recently, we have seen the growth of private equity funds that pool the wealth of individual investors and invest in private firms that show promise. This has allowed investors to invest in private businesses without either giving up diversification or taking an active role in

[17] Movies like Wall Street and Other People's Money and books like Barbarians at the Gate were based upon raiders who did leveraged buyouts for a living.

managing these firms. Pension funds and institutional investors, attracted by the high returns earned by investments in private firms, have also set aside portions of their overall portfolios to invest in private equity.

Most private equity funds are structured as private limited partnerships, where the managers of the fund are the general partners and the investors in the fund – both individual and institutional – are limited partners. The general partners hold on to the power on when and where to invest, and are generously compensated, with annual compensation ranging from 1.5% to 2.5% of the total capital invested. Partnerships typically last from 10 to 12 years and limited partners have to agree to make capital commitments for periods of 5 to 7 years.

The Process of Venture Capital investing

Venture capital can prove useful at different stages of a private firm's existence. *Seed-money venture capital*, for instance, is provided to start-up firms that want to test a concept or develop a new product, while *start-up venture capital* allows firms that have established products and concepts to develop and market them. Additional rounds of venture capital allow private firms that have more established products and markets to expand. There are five steps associated with how venture capital gets to be provided to firms, and how venture capitalists ultimately profit from these investments:

- *Provoke equity investor's interest*: There are hundreds of small firms interested in raising finance from private equity investors, and relatively few venture captialists and private equity investors. Given this imbalance, the first step that a private firm wanting to raise private equity has to take is to get private equity investors interested in investing in it. There are a number of factors that help the private firm, at this stage. One is the *type of business* that the private firm is in, and how attractive this business is to private equity investors. In the late 1980s and early 1990s, for instance, firms in bio-technology were the favored targets for private equity investors. By the late 1990s, the focus had shifted to internet and technology stocks.
- The second factor is the track record of the top manager or managers of the firm. Top managers, who have a track record of converting private businesses into publicly traded firms, have an easier time raising private equity capital. For instance, Jim Clark, who founded Netscape Communications and Silicon Graphics, both successful publicly traded firms, was able to raise private equity for Healtheon, the venture he founded after leaving Netscape, because of his past track record.
- *Valuation and Return Assessment*: Once private equity investors become interested in investing in a firm, the value of the private firm has to be assessed by looking at

both its current and expected prospects. While venture capitalists sometimes use discounted cash flow models to value firms, they are much more likely to value private businesses using what is called the *venture capital method*. Here, the earnings of the private firm are forecast in a future year, when the company can be expected to go public. These earnings, in conjunction with a price-earnings multiple, estimated by looking at publicly traded firms in the same business, is used to assess the value of the firm at the time of the initial public offering; this is called the exit or terminal value.

For instance, assume that a small private software firm is expected to have an initial public offering in 3 years, and that the net income in three years for the firm is expected to be $ 4 million. If the price-earnings ratio of publicly traded software firms is 25, this would yield an estimated exit value of $ 100 million. This value is discounted back to the present at what venture capitalists call a *target rate of return*, which measures what venture capitalists believe is a justifiable return, given the risk that they are exposed to. This target rate of return is usually set at a much higher level[18] than the traditional cost of equity for the firm.

Discounted Terminal Value = Estimated exit value /(1+ Target return)^n

In this example, if the venture capitalist requires a target return on 30% on his or her investment, the discounted terminal value for the firm would be

Discounted Terminal value for InfoSoft = $ 100 million/1.30^3 = $ 45.52 million

- *Structuring the Deal*: In structuring the deal to bring private equity into the firm, the private equity investor and the firm have to negotiate two factors. First, the private equity investor has to determine what proportion of the value of the firm he or she will demand, in return for the private equity investment. The owners of the firm, on the other hand, have to determine how much of the firm they are willing to give up in return for the same capital. In these assessments, the amount of new capital being brought into the firm has to be measured against the estimated firm value. In the software firm example described above, assuming that the venture capitalist is considering investing $ 12 million, he or she would want to own at least 26.36% of the firm.[19]

  Ownership proportion = Capital provided/ Estimated Value

---

[18] By 1999, for instance, the target rate of return for private equity investors was in excess of 30%.

[19] Private equity investors draw a distinction between what a firm will be worth without their capital infusion (pre-money) and what it will be worth with the infusion (post-money). Optimally, they would like their share of the firm to be based upon the pre—money valuation, which will be lower.

= \$ 12/ \$ 45.52 = 26.36%

Second, the private equity investor will impose constraints on the managers of the firm in which the investment is being made. This is to ensure that the private equity investors are protected and that they have a say in how the firm is run.

- *Post-deal Management*: Once the private equity investment has been made in a firm, the private equity investor will often take an active role in the management of the firm. Private equity investors and venture capitalists bring not only a wealth of management experience to the process, but also contacts that can be used to raise more capital and get fresh business for the firm.
- *Exit*: Private equity investors and venture capitalists invest in private businesses because they are interested in earning a high return on these investments. How will these returns be manifested? There are three ways in which a private equity investor can profit from an investment in a business. The first and usually the most lucrative alternative is an initial public offering made by the private firm. While venture capitalists do not usually liquidate their investments at the time of the initial public offering, they can sell at least a portion of their holdings once they are traded[20]. The second alternative is to sell the private business to another firm; the acquiring firm might have strategic or financial reasons for the acquisition. The third alternative is to withdraw cash flows from the firm and liquidate the firm over time. This strategy would not be appropriate for a high growth firm, but it may make sense if investments made by the firm no longer earn excess returns.

The Payoff to Venture Capital and Private Equity Investing

Note that the act of seeking and receiving venture capital is voluntary, and both sides enter into the relationship with the hope of gaining from it. The business gains access to funds that would not have been available otherwise; these funds in turn might enable the firm to bridge the gap until it can become a publicly traded firm. The venture capitalist might contribute management and organizational skills to the venture and provide the credibility needed for the business to raise more financing. The venture capitalist also might provide the know-how needed for the firm to eventually make a public offering of its equity. The venture capitalist gains as well. If the venture capitalist picks the right businesses to fund and provides good management skills and advice, there can be large returns on the initial

---

[20] Black and Gilson (1998) argue that one of the reasons why venture capital is much more active in the U.S. than in Japan or Germany is because the option to go public is much more easily exercised in the U.S.

investment. While the venture capitalist may reap returns from the private business itself, the largest payoff occurs if and when the business goes public and the venture capitalist is able to convert his or her stake into cash at the market price.

How well do venture capital and private equity investors do, relative to the market? There is clearly anecdotal evidence that some private equity investors do very well on individual deals and over time. There are also periods of time when private equity investing collectively offers extraordinary returns. During the 1990s, for instance, venture capital funds earned an average return of 29.5%, compared to the S&P 500's annual return of 15.1%, but there are three potential problems with this comparison. The first is that the appropriate comparison would really be to the NASDAQ, which boomed during the 1990s and contained companies much like those in a venture capital portfolio – young technology firms .The second and related point is that these returns (both on the venture capital funds and the NASDAQ) are before we adjust for the substantial risk associated with the types of companies in their portfolios. The third is that the returns on the venture capital funds themselves are suspect because they are based upon assessments of value (often made by the venture capitalists) of non-traded investments. In fact, many of these venture capital funds were forced to confront both the risk and self-assessment issues in 2000 and 2001 as many of their investments, especially in new technology businesses, were written down to true value. From September 2000 to 2001, for instance, venture capital funds lost 32% of their value, private equity funds lost 21% and buyout funds lost 16% of their value.

When we look at the longer period returns on private equity investing over the last two decades what emerges is the sobering evidence that venture capital does yield high returns but not of the magnitude that some investors expect. Venture Economics, a data service that tracks the returns on private equity investments reported the following short term and long term returns on private equity investments as of September 2001:

Figure 1. Venture Economics' US Private Equity Performance Index (PEPI)

Returns as of September 30, 2001

Fund Type	*1 Yr*	*3 Yr*	*5 Yr*	*10 Yr*	*20 Yr*
Early/Seed Venture Capital	-36.3	81	53.9	33	21.5
Balanced Venture Capital	-30.9	45.9	33.2	24	16.2
Later Stage Venture Capital	-25.9	27.8	22.2	24.5	17
All Venture Capital	-32.4	53.9	37.9	27.4	18.2
All Buyouts	-16.1	2.9	8.1	12.7	15.6
Mezzanine	3.9	10	10.1	11.8	11.3
All Private Equity	-21.4	16.5	17.9	18.8	16.9

S&P 500					

On average, private equity and venture capital funds have outperformed the S&P 500 but the difference is surprisingly small. Between 1991 and 2001, for instance, all private equity funds earned an annual average return only 3.2% higher than the S&P 500 over the same period. Given the high risk associated with these investments, that does not seem like a significant excess return.

There is one final point worth making about private equity and venture capital investments. The average returns reported above are pushed up by the presence of a few investments that make very high returns. Most private equity and venture capital investments fail, and the median (rather than the average) return indicates this propensity. Consider, for instance, the glory years of 1997 through 1999. The conventional wisdom is that private equity investments did well in those years. In 1999, the weighted-average internal rate of return on private equity investments was 119%, but the median return in that year was 2.9%. The median trailed the average badly in 1997 and 1998, as well.

Determinants of Success at Activist Growth Investing

While venture capital and private equity investing, in general, is not a recipe for riskfree high returns, there are some venture capital and private equity investors who succeed and earn extraordinary returns. What set them apart and how can you partake in their success? The keys seem to be the following:

- Pick your companies (and managers) well: Most small private businesses do not succeed, either because the products or services they offer do not find a ready audience or because of poor management. Good venture capitalists seem to have the capacity to find the combination of ideas and management that make success more likely.
- Diversify: The rate of failure is high among private equity investments, making it critical that you spread your bets. The earlier the stage of financing – seed money, for example – the more important it is that you diversify.
- Support and supplement management: Venture capitalists are also management consultants and strategic advisors to the firms that they invest in. If they do this job well, they can help the managers of these firms convert ideas into commercial success.
- Protect your investment as the firm grows: As the firm grows and attracts new investment, you as the venture capitalist will have to protect your share of the business from the demands of those who bring in fresh capital.

- Know when to get out: Having a good exit strategy seems to be as critical as having a good entrance strategy. Know how and when to get out of an investment is critical to protecting your returns.

As a successful venture capitalist, you will still find yourself holding not only a risky portfolio but a relatively undiversified one, with large stakes in a number of small and volatile business. In short, activist growth investing is best suited for investors who have substantial capital, long time horizons and are willing to take risk.

Conclusion

If value investors bet on the market getting it wrong when pricing assets in place, growth investors place their bets on mis-assessments of the value of growth. While some categorize growth investors based upon their willingness to buy high PE stocks, that characterization does not capture the diversity of growth investors. In this chapter, we began by looking at investing in small cap stocks and initial public offerings as growth investing strategies. We then considered a variety of growth screens used by investors to find undervalued growth, ranging from high PE ratios to low PEG ratios. While the empirical evidence is not as supportive of growth screens as it is for value screens, investors who are disciplined, have long time horizons and are good at gauging market cycles can earn significant excess returns.

In the last part of the chapter, we examined venture capital and private equity investing and categorized them as activist growth investing strategies, since they require taking large positions in young growth businesses and then taking an active role in making them succeed. While there are some venture capital and private equity investors who earn huge returns, the overall returns to private equity investing reflect only a modest premium over investing in publicly traded stocks. A large appetite for risk and a long time horizon are pre-requisits for success.

Lessons for Investors

To be a growth investor, you need to

- Make more precise estimates of growth and price it well: The success of growth investing ultimately rests on your capacity to forecast growth and to price it right. If you are better at these roles than the market, you improve your odds of success.
- Catch the growth cycles when they occur: Growth investing has historically done best when earnings growth in the market is low and investors are pessimistic about the future.

To be an activist growth investor, you need to

- Accept skewed returns: Private equity and venture capital investing may offer a few investors spectacular returns, but the average returns to all investors in these categories are low (relative to investing in publicly traded stocks).
- Invest in the right businesses: To succeed at private equity investing, you have to pick the right businesses to make the investments in, diversify your bets and have a well devised exit strategy.

CHAPTER 10

INFORMATION PAYS: TRADING ON NEWS

Information affects stock prices. This undeniable fact is brought home every day as we watch financial markets react to announcements not only from individual firms but also from analysts following these firms and market mavens. When firms report better than anticipated earnings, stock prices go up, whereas announcements that contain bad news push the stock price down. Given this reality, any investor who is able to gain access to information prior to it reaching the market can buy or sell ahead of the information, depending upon whether it is good or bad news, and make a killing.

There is a catch, though. Information that has not yet been made available to markets is viewed as inside information, and trading on it can often be illegal. For any portfolio manager who wants to trade on information, there are three alternatives. One is to use the rumor mills that always exist in financial markets, screening the rumors for credibility (based upon the both the news in the rumors and the source of the rumors) and then trading on them. The other is to wait until the information reaches the market and then to trade on market reaction. Implicit in this approach is the assumption that markets react inappropriately to news items and that it is possible to take advantage of these mistakes. The third is to use private information and hope that you do not cross the line and break the law on insider trading.

In this chapter, we begin by considering the individuals who are most likely to have access to private information – insiders and analysts – and consider whether they use it to earn high returns. We then move on to consider information announcements made by firms – earnings and dividend announcements, acquisitions and investment announcements, for instance – and how markets react to that information and whether there is a possibility of profiting on information, by trading after the announcement. In the final part of the paper, we will consider the essential ingredients of a successful information-based trading strategy.

Information and Prices

No matter how markets are structured, the market price of an asset is an estimate of its value. Investors in the market make assessments of the price based upon their expectations for the future cash flows on the asset. They form these expectations using the information that is available to them and this information can arrive in different forms. It can be public information available in annual reports or filings with the SEC, or information available to one or a few investors. In this section, we will begin by drawing a distinction between private and public information, and then considering how an efficient market should react to information.

Private and Public Information

While the steps in the pricing process – receive information, process the information to form expectations and trade on the asset – may be the same for all investors, there are wide variations across investors in how much information they have available, and how they process the information. Some investors have access to more information than others. For instance, an equity research analyst whose job it is to evaluate a stock as an investment will have access to more information about the firm than a small investor making the same decision. These differences in information are compounded by the different ways in which investors use the information to form expectations. Some investors build complex quantitative models, converting the information into expected earnings and cash flows, and then value investments. Other investors use the same information to make comparisons across traded investments. The net effect is that, at any point in time, investors will disagree on how much an asset is worth. Those who think that it is worth more will be the buyers of the asset, and those who think it is worth less will sell the asset. The market price represents the price at which the market clears, i.e, where demand (buying) is equal to supply (selling).

Let us now consider the relationship between price and value. In chapter 4, we argued that the value of an asset is the present value of the expected cash flows over its lifetime. The price of that asset represents the product of a process in which investors use the information available on the asset to form expectations about the future. The price can and usually will deviate from the value for three reasons. First, the information available may be insufficient or incorrect; then expectations based upon this information will also be wrong. Second, investors may not do a good job of processing the information to arrive at expectations. Third, even if the information is correct and investors, on average, form expectations properly, there might still be investors who are willing to trade at prices that do not reflect these expectations. Thus, an investor who assesses the value of a stock to be $ 50 might still be willing to buy the stock for $ 60, because he or she believes that it can be sold to someone else for $ 75 later.

Information Efficiency: How stock prices react to news

One of the key questions we need an answer to before we examine whether to invest in new projects or how to finance them is whether markets are efficient. There are three ways of measuring or defining market efficiency. One is to look at how much and for how long prices deviate from true value. The second is to measure how quickly and completely prices adjust to reflect new information. The third is to measure whether some investors in markets consistently earn higher returns than others who are exposed to the same amount of risk. It is the last definition that we used in chapter 6.

If we define market efficiency in terms of how much the price of an asset deviates from a firm's true value, the smaller and less persistent the deviations are, the more efficient a market is. Market efficiency does not require that the market price be equal to true value at every point in time. All it requires is that errors in the market price be unbiased, i.e., prices can be greater than or less than true value, as long as these deviations are random. Another way of assessing market efficiency is to look at how quickly and how well markets react to new information. The value of an asset should increase when new information that affects any of the inputs into value – the cash flows, the growth or the risk – reaches the market. In an efficient market, the price of the asset will adjust instantaneously and, on average, correctly to the new information, as shown in figure 10.1.[1]

Figure 10.1: Price Adjustment in an Efficient Market

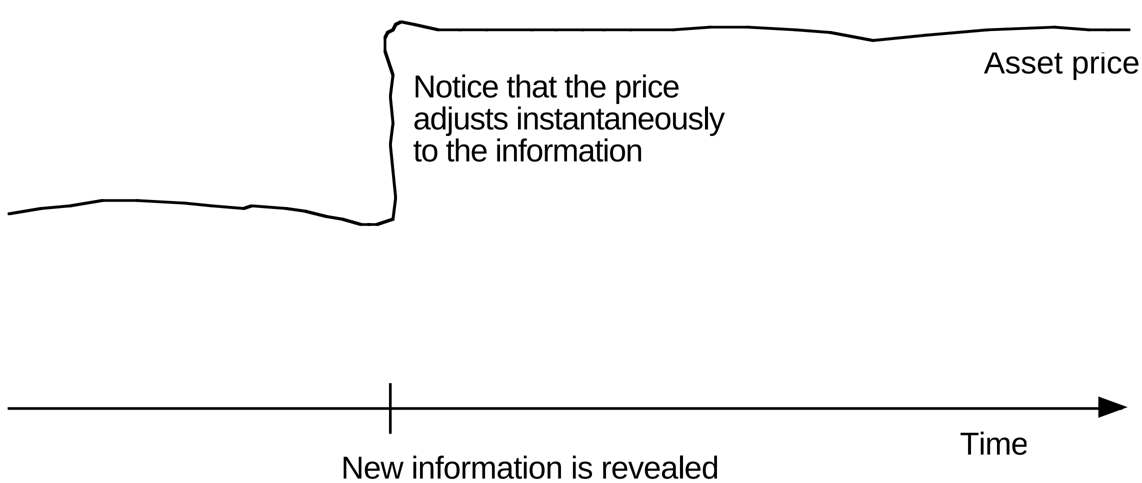

The adjustment will be slower if investors are slow in assessing the impact of the information on value. In figure 10.2, we show the price of an asset adjusting slowly to new information. The drift in prices that we observe after the information arrives is indicative of a slow learning market.

[1] Brown, Harlow and Tinic (1988) present a more sophisticated version of this test by allowing information to change both expectations about cash flows in the future and expectations of risk.

Figure 10.2 A Slow Learning Market

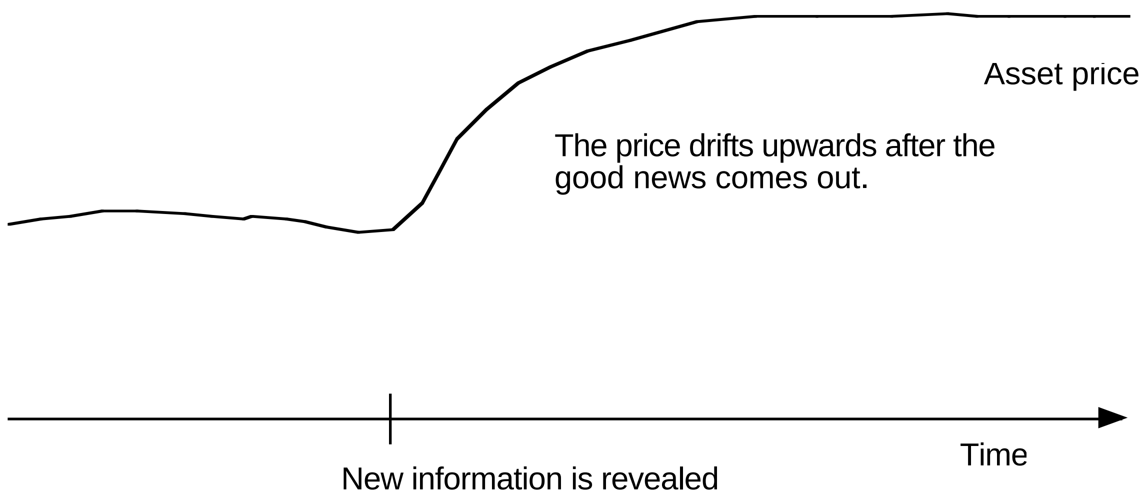

In contrast, the market could adjust instantaneously to the new information but overestimate the effect of the information on value. Then, the price of the asset will increase by more than it should, given the effect of the new positive information on value, or drop by more than it should, with negative information. Figure 10.3 shows the drift in prices in the opposite direction, after the initial reaction.

Figure 10.3: An Overreacting Market

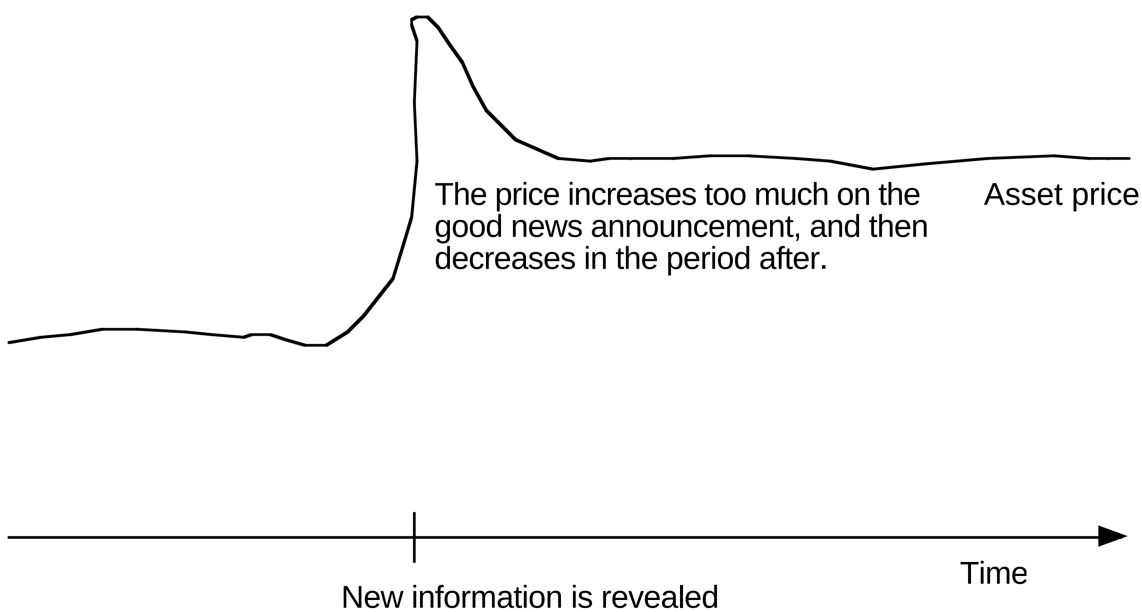

Trading on Private Information

Do investors who have information that no one else has access to, i.e. private information, able to use this information to profit? While the answer seems obviously yes, it is very difficult to test whether they do. The reason for this is that the regulatory authorities, at least in the United States, specifically forbid trading in advance of significant information releases. Thus, insiders who follow the law and register their trades with the SEC are not likely to be trading on specific information in the first place. Notwithstanding this selection bias, we will begin by looking at whether insider buying and selling operate as signals of future price movements, since insiders may still have access to general information about the firm that outsiders do not. We will then look at the more difficult question of whether those who trade illegally on private information make excess returns. While this may seem like an impossible test to run, we can at least draw inferences about this trading by looking at trading volume and price movements prior to major news announcements.

Insiders

The SEC defines an insider to be an officer or director of the firm or a major stockholder (holding more than 5% of the outstanding stock in the firm). Insiders are barred from trading in advance of specific information on the company and are required to file with the SEC when they buy or sell stock in the company. In this section, we will begin by looking at the relationship between insider trading and subsequent stock price changes, and then consider whether non-insiders can use information on insider trading to earn excess returns themselves.

Insider Trading and Stock Prices

If it is assumed, as seems reasonable, that insiders have better information about the company, and consequently better estimates of value, than other investors, the decisions by insiders to buy and sell stock should signal future movements in stock prices. Figure 10.4, derived from an early study of insider trading by Jaffe, examines excess returns on two groups of stock, classified on the basis of insider trades. The "buy group" includes stocks where insider buys exceeded sells by the biggest margin, and the "sell group" includes stocks where insider sells exceed buys by the biggest margin.

Stocks with most insider holdings: Take a look at the 50 stocks with the highest insider holdings as a percent of outstanding stock.

Figure 10.4: Cumulative Returns Following Insider Trading: Buy vs Sell Group

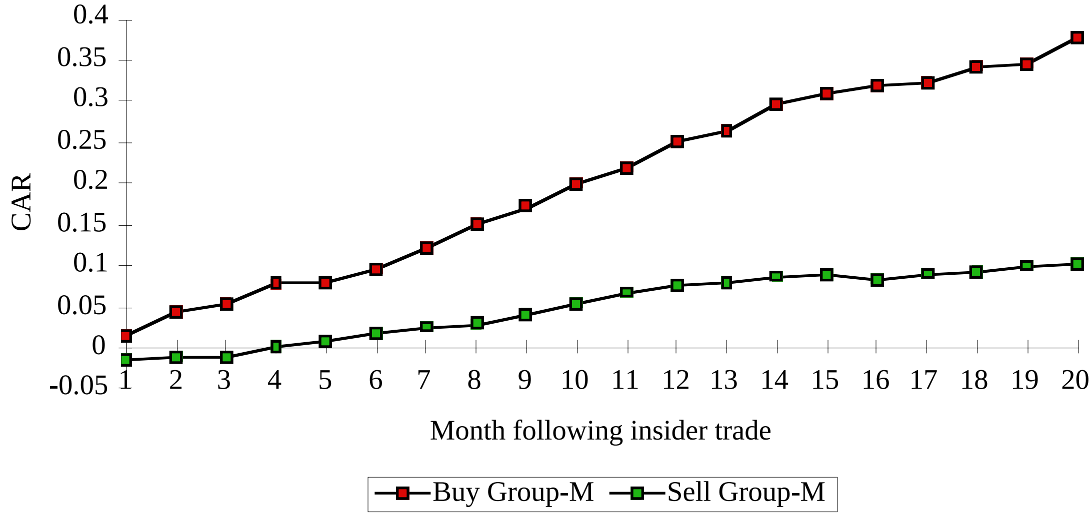

Source: Jaffe

Studies since support this finding[2], but it is worth noting that insider buying is a noisy signal – about 4 in 10 stocks where insiders are buying turn out to be poor investments, and even on average, the excess returns earned are not very large. In a study in 1998, Lakonishok and Lee take a closer look at the price movements around insider trading. They find that firms with substantial insider selling have stock returns of 14.4% over the subsequent 12 months, which is significantly lower than the 22.2% earned by firms with insider buying. However, they find that the link between insider trading and subsequent returns is greatest for small companies and that there is almost no relationship at larger firms.

Can you follow insiders?

If insider trading offers advance warning, albeit a noisy one, of future price movements, can we as outside investors use this information to make better investment decisions? In other words, when looking for stocks to buy should we consider the magnitude of insider buying and selling on the stock? To answer this question, we first have to recognize that since the SEC does not require an immediate filing of insider trades, investors will find out about insider trading on a stock with a lag, of a few weeks or even a few months. In fact, until recently, it was difficult for an investor to access the public filings

---

[2] See Finnerty (1976), Seyhun (1986) and Rozeff and Zaman (1988)

on insider trading. As these filings have been put online in recent years, this information on insider trading has become available to more and more investors.

A study of insider trading examined excess returns around both the date the insiders report to the SEC and the date that information becomes available to investors in the official summary. Figure 10.5 presents the contrast between the two event studies.

Figure 10.5: Abnormal Returns around Reporting Day/ Official Summary Availability Day

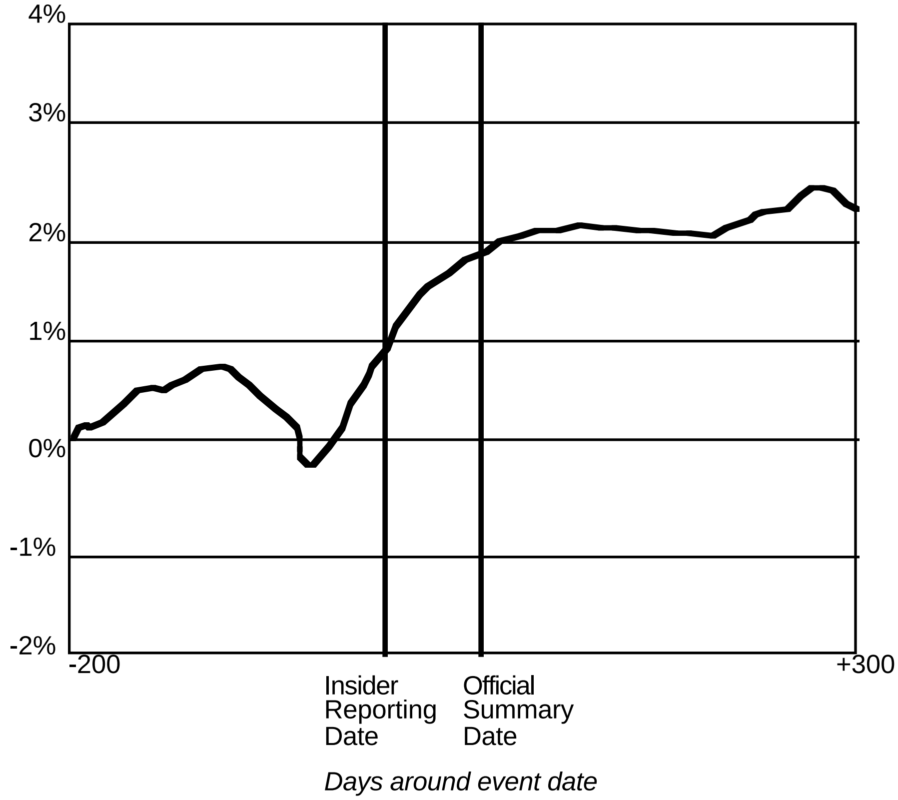

Given the opportunity to buy on the date the insider reports to the SEC, investors could have marginal excess returns (of about 1%), but these returns diminish and become statistically insignificant, if investors are forced to wait until the official summary date. If you control for transactions costs, there are no excess returns associated with the use of insider trading information.[3]

Does this mean that insider trading information is useless? It may be so if we focus on total insider buying and selling but there may be value added if we can break down insider trading into more detail. Consider the following propositions:

- Not all insiders have equal access to information. Top managers and members of the board should be privy to much more important information and thus their trades

[3] This is also the conclusion drawn by Seyhun (1986) and Rozeff and Zaman (1988).

should be more revealing. A study by Bettis, Vickrey and Vickery finds that investors who focus only on large trades made by top executives, rather than total insider trading may, in fact, be able to earn excess returns.

- As investment alternatives to trading on common stock have multiplied, insiders have also become more sophisticated about using these alternatives. As an outside investor, you may be able to add more value by tracking these alternative investments. For instance, Bettis, Bizjak and Lemmon find that insider trading in derivative securities (options specifically) to hedge their common stock positions increases immediately following price run-ups and prior to poor earnings announcements. In addition, they find that stock prices tend to go down after insiders take these hedging positions.

Illegal Insider Trading

None of the studies quoted above answer the question of whether insiders themselves make excess returns. The reporting process, as set up now by the SEC, is biased toward legal and less profitable trades, and away from illegal and more profitable trades. Though direct evidence cannot be easily offered for this proposition, insiders trading illegally on private information must make excess returns. To support this proposition, we can present three pieces of evidence.

- The first (and weakest) is anecdotal. When insiders are caught trading illegally, they almost invariably have made a killing on their investment. Clearly, some insiders made significant returns off their privileged positions. The reason that it has to be viewed as weak evidence, though, is because the SEC looks for large profits as one of the indicators of whether it will prosecute. In other words, an insider who trades illegally on information may be breaking the law but is less likely to be prosecuted for the act if he or she loses money.
- Almost all major news announcements made by firms are preceded by a price run-up (if it is good news) or a price drop (if it is bad news). Thus, you see that the stock price of a target firm starts drifts up before the actual takeover announcement, and that the stock price of a firm reporting disappointing earnings drops in the days prior to the earnings report. While this may indicate a very prescient market, it is much more likely that someone with access to the privileged information (either at the firm or the intermediaries helping the firm) is using the information to trade ahead of the news. In fact, the other indicator of insider trading is the surge in

trading volume in both the stock itself and derivatives prior to big news announcements.[4]

- In addition to having access to information, insiders are often in a position to time the release of relevant information to financial markets. Knowing as they do that they are not allowed to trade ahead of this information, insiders often adjust information disclosure to make it less likely that they will be targeted by the SEC. One study[5] find that insiders sell stock between 3 and 9 quarters before their firms report a break in consecutive earnings increases.[6] They also find, for instance, that insider selling increases at growth firms prior to periods of declining earnings.

Using Insider Trading in Investment Decisions

As the information on insider trades has become more accessible, it has also become less useful. In addition, the spurt in the use of options in management compensation schemes has introduced a substantial amount of noise in the reporting system, since a large proportion of insider trades now are associated with managers exercising options and then selling a portion of their stock holding for liquidity and diversification reasons. For information on insider trading to pay off, you need to look beyond the total insider trading numbers at the insiders themselves, focusing on large trades by top managers at smaller, less followed firms. Even then, you should not expect miracles, since you are using publicly available information.

We believe that the real payoff comes from tracking illegal insider trading by looking at trading volume and bid-ask spreads. The relationship between trading volume and private information may provide an intuitive rationale for the use of some of the volume measures described in chapter 7 as technical indicators.

Tracking Insider Trading

Tracking what legal insiders are doing has become both easier and more timely. You can look at the filings made by companies on the SEC website (http://www.sec.gov). The insider trading information is available in forms 3, 4 and 144s. Many of the more popular financial web sites such as Yahoo! Finance report on recent insider transactions on individual companies. If you are willing to pay more, you can subscribe to services that consolidate the information and provide it to you.

---

[4] It is for this reason that the SEC tracks trading volume. Sudden increases in volume often trigger investigations of insiders at firms.

[5] See Ke, Huddart and Petroni (2002)

[6] You generally face legal jeopardy when you sell in the quarter or two before the news announcement.

Analysts

Analysts clearly hold a privileged position in the market for information, operating at the nexus of private and public information. Using both types of information, analysts make earnings forecasts for the firms that they follow, and issue buy and sell recommendations to their clients, who trade on its basis. In this section, we will consider where there is valuable information in these forecasts and recommendations and whether incorporating them into investment decisions leads to higher returns.

Who do analysts follow?

The number of analysts tracking firms varies widely across firms. At one extreme are firms like GE, Cisco and Microsoft that are followed by dozens of analysts. At the other extreme, there are hundreds of firms that are not followed by any analysts. Figure 10.6 shows the divergence across firms in the United States, in terms of the number of analysts following them.

Figure 10.6 : Number of analysts estimating earnings per shar: U.S. firms in January 2001

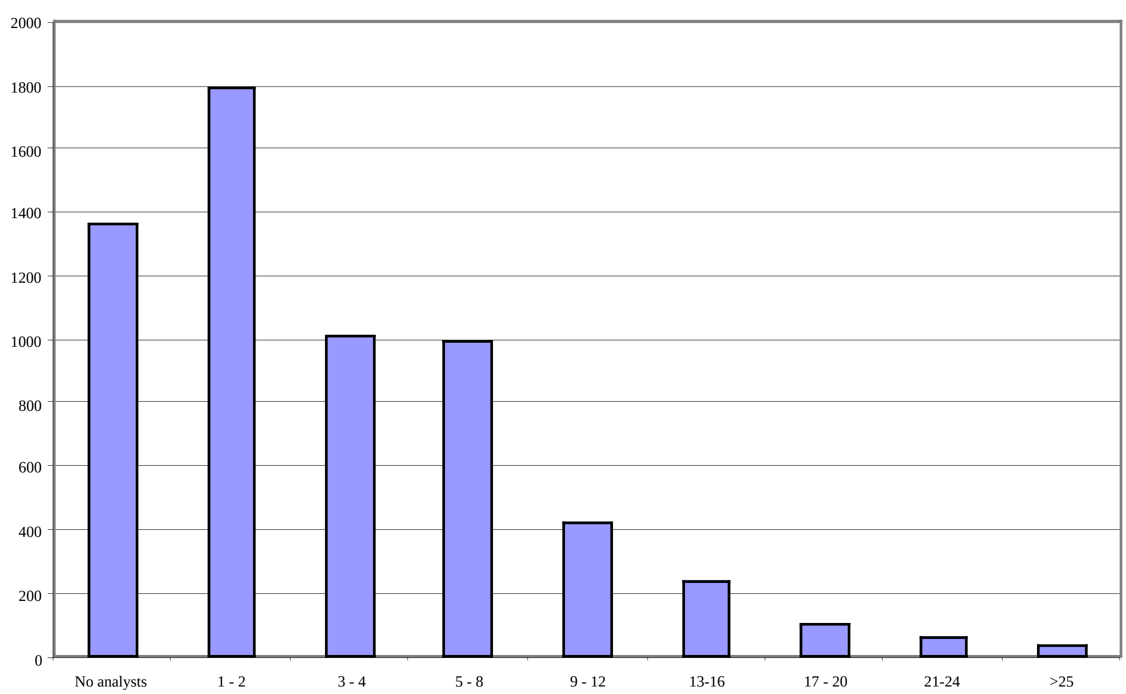

Source: Value Line

Why are some firms more heavily followed than others? These seem to be some of the determinants:

- *Market Capitalization*: The larger the market capitalization of a firm, the more likely it is to be followed by analysts.
- *Institutional Holding*: The greater the percent of a firm's stock that is held by institutions, the more likely it is to be followed by analysts. The open question, though, is whether analysts follow institutions or whether institutions follow analysts. Given that institutional investors are the biggest clients of equity research analysts, the causality probably runs both ways.
- *Trading Volume*: Analysts are more likely to follow liquid stocks. Here again, though, it is worth noting that the presence of analysts and buy (or sell) recommendations on a stock may play a role in increasing trading volume.

Sell Side and Buy Side Analysts: A Primer

There are thousands of financial analysts who try to value stocks and most of them toil anonymously. The analysts who receive the most attention are the sell-side analysts who work for investment banks. Their research is primarily for external consumption and their roles are complex. They interact with the firms they research and sell their research to portfolio managers and individual investors. Buy side analysts, on the other hand, work for money management companies like Fidelity. Their research is intended primarily for internal consumption and is designed to help portfolio managers pick better stocks.

Why does it matter? Sell side equity research may have a higher profile than buy-side research but it is also buffeted by far more conflicts of interest and bias. The fact that the investment banks that churn out the research do not have to invest in the stocks that they recommend should give pause to individual investors who intend to follow these recommendations. In addition, sell side analysts have to spend substantially more time selling than buy side analysts do.

I. Earnings Forecasts

Analysts spend a considerable amount of time and resources forecasting earnings per share both for the next quarter and for the next financial year. Presumably, this is where their access to company management and private information should generate an advantage. Thus, when analysts revise their forecasts upwards or downwards, they convey information to financial markets and prices should react. In this section, we examine how markets react to analyst forecast revisions and whether there is potential for us as investors to take advantage of this reaction.

The Information in Analyst Forecasts

There is a simple reason to believe that analyst forecasts of growth should be better than using historical growth rates. Analysts, in addition to using historical data, can avail themselves of other information that may be useful in predicting future growth.

1. Firm-specific information that has been made public since the last earnings report: Analysts can use information that has come out about the firm since the last earnings report, to make predictions about future growth. This information can sometimes lead to significant re-evaluation of the firm's expected earnings and cash flows.

2. Macro-economic information that may impact future growth *:* The expected growth rates of all firms are affected by economic news on GNP growth, interest rates and inflation. Analysts can update their projections of future growth as new information comes out about the overall economy and about changes in fiscal and monetary policy. Information, for instance, that shows the economy growing at a faster rate than forecast will result in analysts increasing their estimates of expected growth for cyclical firms.

3. Information revealed by competitors on future prospects*:* Analysts can also condition their growth estimates for a firm on information revealed by competitors on pricing policy and future growth. For instance, a negative earnings report by one telecommunications firm can lead to a reassessment of earnings for other telecommunication firms.

4. Private information about the firm: Analysts sometimes have access to private information about the firms they follow which may be relevant in forecasting future growth. This avoids answering the delicate question of when private information becomes illegal inside information. There is no doubt, however, that good private information can lead to significantly better estimates of future growth. In an attempt to restrict this type of information leakage, the SEC issued new regulations preventing firms from selectively revealing information to a few analysts or investors. Outside the United States, however, firms routinely convey private information to analysts following them.

5. Public information other than earnings: Models for forecasting earnings that depend entirely upon past earnings data may ignore other publicly available information that is useful in forecasting future earnings. It has been shown, for instance, that other financial variables such as earnings retention, profit margins and asset turnover are useful in predicting future growth. Analysts can incorporate information from these variables into their forecasts.

The Quality of Earnings Forecasts[7]

If firms are followed by a large number of analysts and these analysts are indeed better informed than the rest of the market, the forecasts of growth that emerge from analysts should be better than estimates based upon either historical growth or other publicly available information. But is this presumption justified? Are analyst forecasts of growth superior to other estimates?

The general consensus from studies that have looked at short-term forecasts (one quarter ahead to four quarters ahead) of earnings is that analysts provide better forecasts of earnings than models that depend purely upon historical data. The mean relative absolute error, which measures the absolute difference between the actual earnings and the forecast for the next quarter, in percentage terms, is smaller for analyst forecasts than it is for forecasts based upon historical data. Two other studies shed further light on the value of analysts' forecasts. A study[8] in 1978 examine the relative accuracy of forecasts in the Earnings Forecaster, a publication from Standard and Poors that summarizes forecasts of earnings from more than 50 investment firms. This study measured the squared forecast errors by month of the year and computed the ratio of analyst forecast error to the forecast error from time-series models of earnings. It found that the time series models actually outperform analyst forecasts from April until August, but underperform them from September through January. The authors of the study hypothesize that this is because there is more firm-specific information available to analysts during the latter part of the year. The other study by O'Brien (1988) compares consensus analyst forecasts from the Institutions Brokers Estimate System (I/B/E/S) with time series forecasts from one quarter ahead to four quarters ahead. The analyst forecasts outperform the time series model for one-quarter ahead and two-quarter ahead forecasts, do as well as the time series model for three-quarter ahead forecasts and do worse than the time series model for four-quarter ahead forecasts. Thus, the advantage gained by analysts from firm-specific information seems to deteriorate as the time horizon for forecasting is extended. Dreman and Berry examined analyst forecasts from 1974 to 1991 and found that in more than 55% of the forecasts examined, analyst estimates of earnings were off by more than ten percent from actual earnings.[9] One

---

[7] Sell side analysts work for brokerage houses and investment banks and their research is offered to clients of these firms as a service. In contrast, buy side analysts work for institutional investors and their research is generally proprietary.

[8] See Crichfield, Dyckman and Lakonishok (1978).

[9] Dreman, D.N. and M. Berry, 1995, Analyst Forecasting Errors and their Implications for Security Analysis", Financial Analysts Journal, May/June, pg 30-41.

potential explanation given for this poor forecasting is that analysts are routinely over optimistic about future growth. A study in 1988 by Chopra finds that a great deal of this forecast error comes from the failure of analysts to consider large macroeconomic shifts. In other words, as figure 10.7 indicates, analysts tend to overestimate growth at the peak of a recovery and under estimate growth in the midst of a recession.[10] Higgins (1998) compares analyst forecast errors across seven countries and suggests, not surprisingly, that analysts are more accurate and less biased in countries that mandate more financial disclosure.

[10] Chopra, V.K., 1998, Why so much error in analyst forecasts?, Financial Analysts Journal, Nov-Dec, pg 35-42.

Figure 10.7: Earnings Forecast Errors and Economic Growth

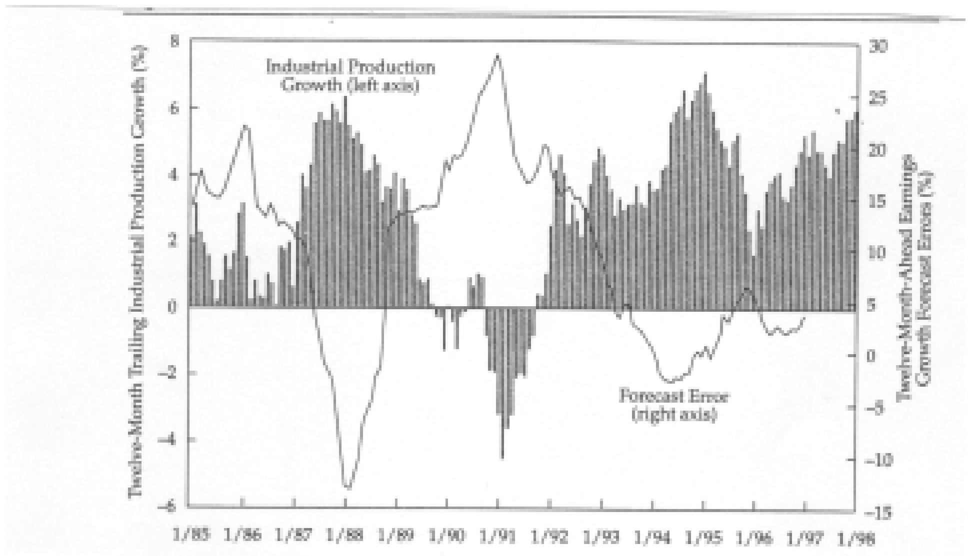

There is little evidence to suggest that analysts provide superior forecasts of earnings when the forecasts are over three or five years. An early study by Cragg and Malkiel compared long-term forecasts by five investment management firms in 1962 and 1963 with actual growth over the following three years to conclude that analysts were poor long term forecasters. This view was contested in 1988 by Vander Weide and Carleton who found that the consensus prediction of five-year growth in the I/B/E/S was superior to historically oriented growth measures in predicting future growth.

There is an intuitive basis for arguing that analyst predictions of growth rates must be better than time-series or other historical-data based models simply because they use more information. The evidence indicates, however, that this superiority in forecasting is surprisingly small for long-term forecasts and that past growth rates play a significant role in determining analyst forecasts.

Market Reaction to Earnings Forecast Revisions

In chapter 7, we considered the price momentum strategies where investors buy stocks that have gone up the most in recent periods, expecting the momentum to carry forward into future periods. You could construct similar strategies based upon earnings momentum. While some of these strategies are based purely upon earnings growth rates, most of them are based upon how earnings measure up to analyst expectations. In fact, one strategy is to buy stocks where analysts are revising earnings forecasts upwards, and hope that stock prices follow these earnings revisions. A number of studies in the United States seem to conclude that it is possible to use forecast revisions made by analysts to earn excess returns. In one of the earliest studies of this phenomenon, Givoly and Lakonishok created portfolios of 49 stocks in three sectors, based upon earnings revisions, and reported earning an excess return on 4.7% over the following four months on the stocks with the most positive revisions. Hawkins, in 1983, reported that a portfolio of stocks with the 20 largest upward revisions in earnings on the I/B/E/S database would have earned an annualized return of 14% as opposed to the index return of only 7%. In another study, Cooper, Day and Lewis report that much of the excess returns is concentrated in the weeks around the revision – 1.27% in the week before the forecast revision, and 1.12% in the week after, and that analysts that they categorize as leaders (based upon timeliness, impact and accuracy) have a much greater impact on both trading volume and prices. In 2001, Capstaff, Paudyal and Rees expanded the research to look at earnings forecasts in other countries and concluded that you could have earned excess returns of 4.7% in the U.K, 2% in France and 3.3% in Germany from buying stocks with the most positive revisions.

Potential Pitfalls

The limitation of an earnings momentum strategy is its dependence on two of the weakest links in financial markets –earnings reports that come from firms and analyst forecasts of these earnings. In recent years, we have become increasing aware not only that of the capacity of firms to manage their earnings but also to manipulate them using questionable accounting ploys. At the same time, we have discovered that analysts' forecasts are biased not only by their closeness to the firm they follow but also because of their investment banking relationships.

Even if the excess returns persist, you also need to consider why they might exist in the first place. To the extent that analysts influence trades made by their clients, they are likely to affect prices when they revise earnings. The more influential they are, the greater the effect they will have on prices, but the question is whether the effect is lasting. One way you may be able to earn higher returns from this strategy is to identify key analysts and build an investment strategy around forecast revisions made by them, rather than looking at consensus estimates made by all analysts.

Finally, you should recognize that it is a short-term strategy that yields fairly small excess returns over investment horizons ranging from a few weeks to a few months. The increasing skepticism of markets towards both earnings reports from firms and forecasts by analysts bodes ill for these strategies. While forecast revisions and earnings surprises by themselves are unlikely to generate lucrative portfolios, they can augment other more long-term screening strategies. One way you may be able to earn higher returns from this strategy is to identify key analysts who are both independent and influential and build an investment strategy around forecast revisions made by them, rather than looking at consensus estimates made by all analysts.

II. Analyst Recommendations

The centerpiece of analyst reports are the recommendations that they make on stocks. You would expect stock prices to react to analyst recommendations, when they are made, if for no other reason than for the fact that some investors follow these recommendations. In this section,, we consider some key empirical facts about analyst recommendations first and then consider how markets react to them. We close with an analysis of whether investors who use these recommendations to make investment decisions can make money off them in the short and the long term.

The Recommendation Game

There are three empirical facts that need to be laid on the table about recommendations before we start examining how markets react to them.

- If we categorize analyst recommendations into buy, sell and hold, the overwhelming number are buy recommendations. In 2001, for instance, buy recommendations outnumbered sell recommendations 7 to 1, but that was actually a drop from the late 1990s, where sell recommendations were often outnumbered by more than 25 to 1.
- Part of the reason for this imbalance between buy and sell recommendations is that analysts often have many more layers beyond buy, sell and hold. Some investment banks, for instance, have numerical rating systems for stocks where stock are classified from 1 to 5 (as is the case with Value Line) whereas others break buy and sell recommendations into sub classes (strong buy, weak buy). What this allows them to do is not only rate stocks more finely, but also to send sell signals without ever saying the word. Thus, an analyst downgrading a stock from a strong buy to a weak buy is sending a sell signal on the stock.
- As with earnings forecasts, there is herd behavior when it comes to recommendations. Thus, when one analyst upgrades a stock from a weak buy to a strong buy, there tends to be a rush of other analyst upgrades in the following days.

The Market Reaction to Recommendations

How do markets react to recommendations made by analysts? A study by Womack examined the stock price response to buy and sell recommendations on the day of the recommendation and in the weeks following. While both buy and sell recommendations affect stock prices, sell recommendations affect prices much more than buy recommendation. This should not be surprising when you remember that buy recommendations vastly outnumber sell recommendations. Interestingly, this study also documents that the price effect of buy recommendations tends to be immediate and there is no evidence of price drifts after the announcement, whereas prices continue to trend down after sell recommendations. Figure 10.8 graphs his findings. Stock prices increase by about 3% on buy recommendations whereas they drop by about 4% on sell recommendations at the time of the recommendations (3 days around reports). In the six months following, prices decline an additional 5% for sell recommendations, while leveling off for buy recommendations.

Figure 10.8: Market Reaction to Recommendations: 1989- 1990

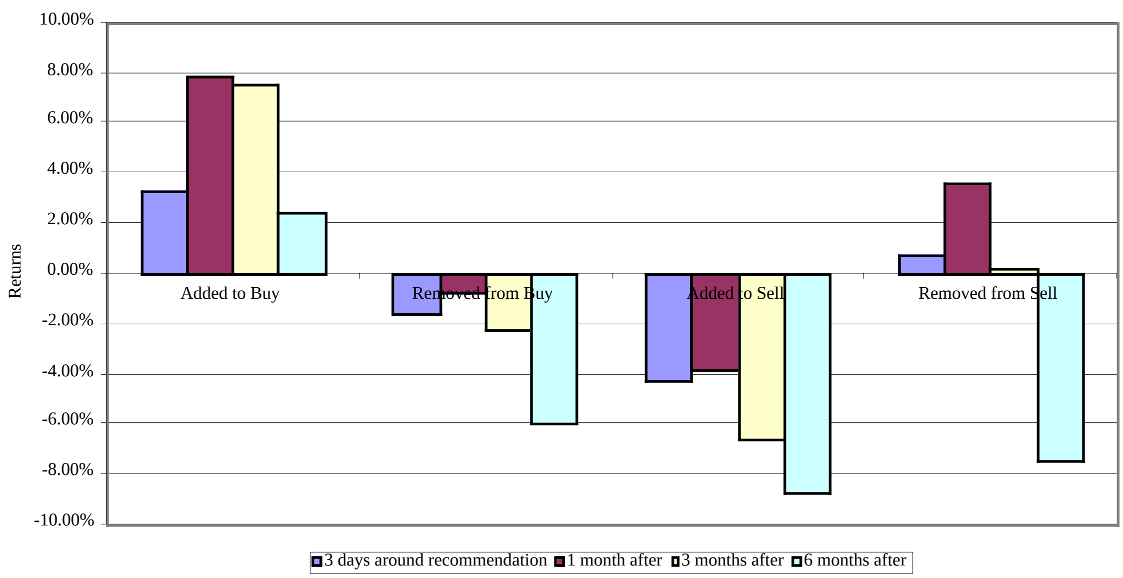

Source: Womack

One of the key issues that equity research analysts were confronted with in the aftermath of the bursting of the dot-com bubble is the extent to which recommendations were perceived to be driven not by views on the stock itself but as cheerleading for investment banking business done by the firms followed by the analysts. Michaely and Womack test this proposition[11] by looking at the stock price performance of buy recommendations after initial public offerings and comparing recommendations made by analysts who work for the underwriters on these offerings and recommendations from analysts who do not. Their findings are summarized in figure 10.9:

---

[11] Michaely, R. and K.L. Womack, Conflicts of Interests and the Credibility of Underwriter Analysts Recommendation, Review of Financial Studies, Winter, 635-686.

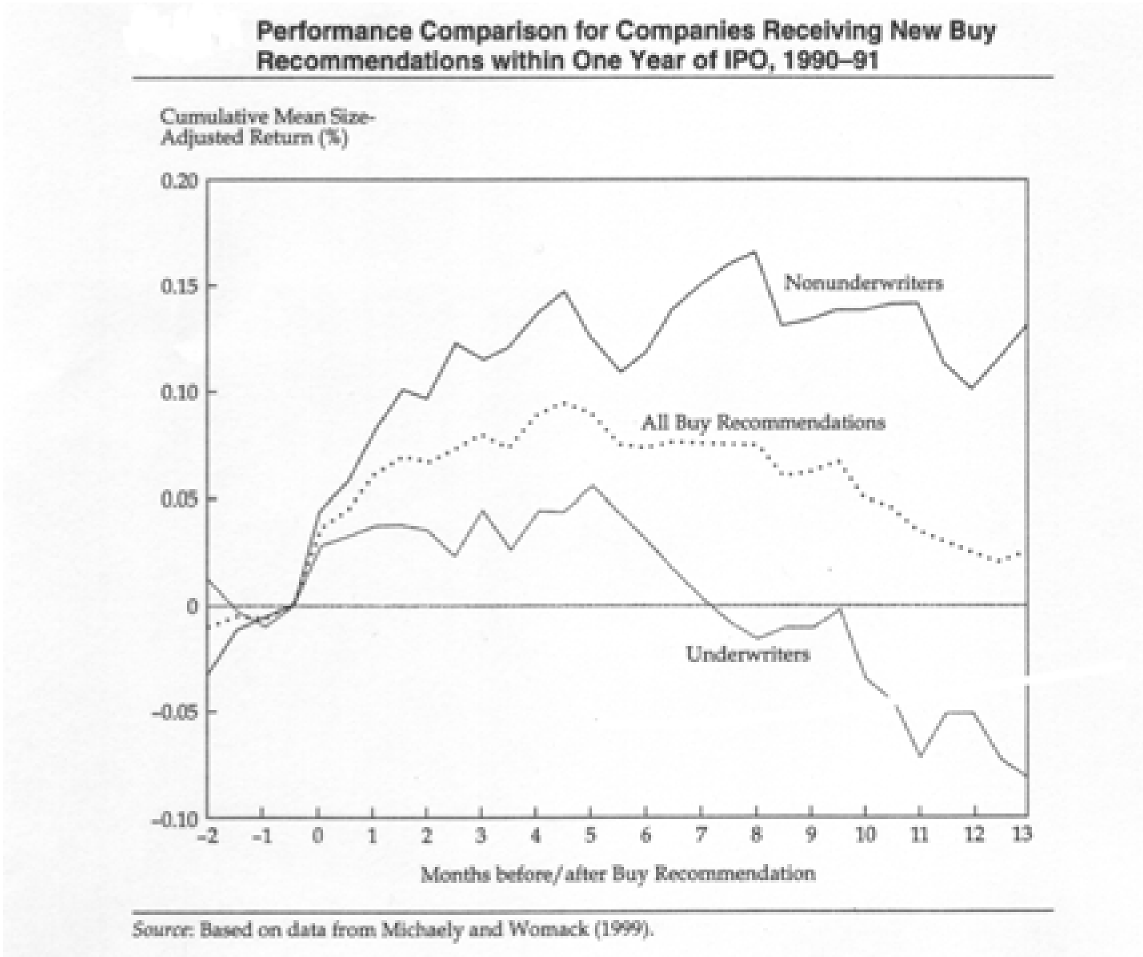

Note that stock prices for recommendations made by non-underwriters do significantly better than the market, but the stocks recommended by underwriters (in those stocks) tend to do poorly. While this may seem obvious, many investors in the late nineties deliberately overlooked the connections between analysts and the firms that they analyzed and paid a significant price for it.[12]

Potential and Perils of Analyst Recommendations

Can you make money off analyst recommendations? The answer seems to be yes, at least in the short term. Even if there were no new information contained in

[12] In June 2002, Merrill Lynch agreed to pay $ 100 million to settle with New York State, after the state uncovered emails sent by Henry Blodgett, Merrill's well know internet analyst, that seemed to disparage stocks internally as he was recommending them to outside clients. The fact that many of these stocks were being taken to the market by Merrill added fuel to the fire. Merrill agreed to make public any potential conflicts of interest it may have on the firms followed by its equity research analysts.

recommendations, there is the self-fulfilling prophecy created by clients who trade on these recommendations, pushing up stock prices after buy recommendations and pushing them down after sell recommendations.[13] If this is the only reason for the stock price reaction, though, the returns are not only likely to be small but could very quickly dissipate, leaving you with large transactions costs and little to show for them.

To incorporate analyst recommendations into an investment strategy, you need to adopt a more nuanced approach. You should begin by identifying the analysts who are not only the most influential but also have the most content (private information) in their recommendations. In addition, you may want to screen out analysts where the potential conflicts of interest may be too large for the recommendations to be unbiased. You should invest based upon their recommendations, preferably at the time the recommendations are made.[14] Assuming that you still attach credence to the views of the recommending analysts, you should watch the analysts for signals that they have changed or are changing their minds. Since these signals are often subtle, you can easily miss them.

Finding the best analysts

How does one go about finding the best analysts following a stock? Do not fall for the hype. The highest profile analysts are not always the best and some analysts are notorious for self promotion. The best sources of information on analysts tend to be outside services without an axe to grind. The Wall Street Journal has a special section on sell-side equity research analysts, where it evaluates analysts on the quality of their recommendations and ranks them on that basis. There are a few online services that track equity research forecasts and recommendations and report on how close actual earnings numbers were to their forecasts.

There are qualitative factors to consider as well. Analysts who have clear, well thought out analyses that show a deep understanding of the businesses that they analyze should be given more weight that analysts who make spectacular recommendations based upon facile analysis. Most importantly, good analysts should just as willing to stand up to the management of companies and disagree with them (and issue sell recommendations).

---

[13] This can be a significant factor. When the Wall Street Journal publishes its Dartboard column, it reports on the stocks being recommended by the analysts its picks. These stocks increase in price by about 4% in the two days after they are picked but reverse themselves in the weeks that follow.

[14] This might not be your choice to make since analysts reveal their recommendations first to their clients. If you are not a client, you will often learn about the recommendation only after the clients have been given a chance to take positions on the stock.

Trading on Public Information

Most of us do not have access to private information about firms, but we all share access to public information about a firm. Some of this public information takes the form of periodic earnings reports and dividend announcements, made four times every year by most firms in the United States and less frequently elsewhere, and some of it is news made by the firm when it announces that it is taking over another firm (or being taken over) or making a major investment or divestiture. In some cases, the information comes from a regulatory authority governing the firm's fortunes, as is the case when the FDA announces that it has approved (or not approved) a drug for treatment. In each of these cases, we would expect the stock price to react to the news contained in the announcement. If the market reaction is appropriate, there is little that we can do to make money off the news, but if the market reaction is not, we may be able to exploit it with specific trading strategies.

Earnings Announcements

When firms make earnings announcements, they convey information to financial markets about their current and future prospects. The magnitude of the information, and the size of the market reaction, should depend upon how much the earnings report exceeds or falls short of investor expectations. In an efficient market, there should be an instantaneous reaction to the earnings report, if it contains surprising information, and prices should increase following positive surprises and down following negative surprises.

Earnings Surprises and Price Reaction

Since actual earnings are compared to investor expectations, one of the key parts of an earnings event study is the measurement of these expectations. Some of the earlier studies used earnings from the same quarter in the prior year as a measure of expected earnings, i.e., firms that report increases in quarter-to-quarter earnings provide positive surprises and those that report decreases in quarter-to-quarter earnings provide negative surprises. In more recent studies, analyst estimates of earnings have been used as a proxy for expected earnings, and compared to the actual earnings. Figure 10.10 provides a graph of price reactions to earnings surprises, classified on the basis of magnitude into different classes from 'most negative' earnings reports (Group 1) to 'most positive' earnings reports (Group 10).

Figure 10.10: Price Reaction to Quarterly Earnings Report

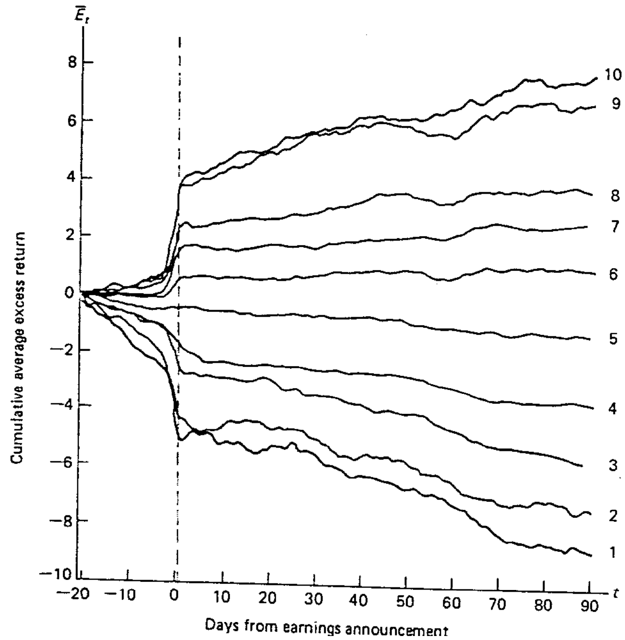

The evidence contained in this graph is consistent with the evidence in most earnings announcement studies -

(a) The earnings announcement clearly conveys valuable information to financial markets; there are positive excess returns (cumulative abnormal returns) after positive announcements and negative excess returns around negative announcements.

(b) There is some evidence of a market reaction in the days immediately prior to the earnings announcement which is consistent with the nature of the announcement, i.e., prices tend to go up on the days before positive announcements and down in the days before negative announcements. This can be viewed either as evidence of insider trading, information leakage or as a consequence of getting the announcement date wrong[15].

(c) There is some evidence, albeit weak, of a price drift in the days following an earnings announcement. Thus, a positive report evokes a positive market reaction on the announcement date, and there are mildly positive excess returns in the days and

---

[15] The Wall Street Journal or COMPUSTAT are often used as information sources to extract announcement dates for earnings. For some firms, news of the announcement may actually cross the news wire the day before the Wall Street Journal announcement, leading to a misidentification of the report date and the drift in returns the day before the announcement.

weeks following the earnings announcement. Similar conclusions emerge for negative earnings reports.

While the study quoted above looked at all earnings announcements, there are studies that indicate that the returns associated with earnings surprises are more pronounced with some types of stocks than with others. For instance,

- A study of value and growth stocks found, instance, that the returns in the three days around earnings announcements were much more positive for value stocks (defined as low PE and PBV stocks) than for growth stocks across all earnings announcements – positive as well as negative. This suggests that you are much more likely to get a positive surprise with a value stock than with a growth stock, indicating perhaps that markets tend to be overly optimistic in their expectations for growth companies.
- Earnings announcements made by smaller firms seem to have a larger impact on stock prices on the announcement date and prices are more likely to drift after the announcement.

Earnings Delays and Price Reaction

The management of a firm has some discretion on the timing of earnings reports and there is some evidence that the timing affects expected returns. A study of earnings reports, classified by the day of the week that the earnings are reported, reveals that earnings and dividend reports on Fridays are much more likely to contain negative information than announcements on any other day of the week. This is shown in figure 10.11.

Figure 10.11: Earnings and Dividend Reports by Day of the Week

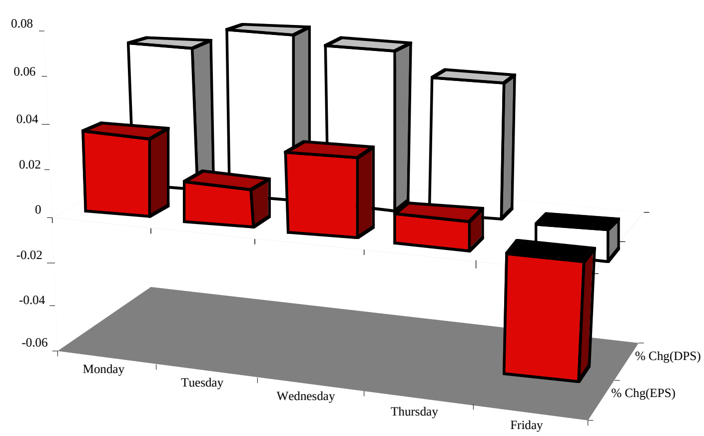

Announcements made on Friday are more likely to contain bad news – earnings drops and dividend cuts - than announcements on any other day of the week, and a significant number of these announcements come out after close of trading on Friday. This may provide an interesting link to the weekend effect described in chapter 7.

There is also some evidence that earnings announcements that are delayed, relative to the expected announcement date,[16] are much more likely to contain bad news than earnings announcements which are early or on time. This is graphed in Figure 10.12.

FIGURE 10.12: Cumulated Abnormal Returns around Earnings Reports: Day 0 is Earnings Announcement Date

Earnings announcements that are more than six days late, relative to the expected announcement date, are much more likely to contain bad news and evoke negative market reactions than earnings announcements that are on time or early. It may be worth the while of investors who build their investment strategy around earnings announcements to keep track of expected earnings announcement dates.

Intraday Price Reaction

Studies have examined the speed with which prices react to earnings announcements in the same day. There, the evidence is mixed. Woodruff and Senchack examined price adjustment by transaction after favorable (surprise>20%) and unfavorable (surprise < -20%)

---

[16] Firms in the United States tend to be consistent about the date each year that they reveal their quarterly earnings. The delay is computed relative to this expected date.

earnings reports, and reported the proportion of the eventual adjustment that has occurred by the hour after the earnings report for each category.

Figure 10.13 : Price Adjustment by Hour after Earnings Report

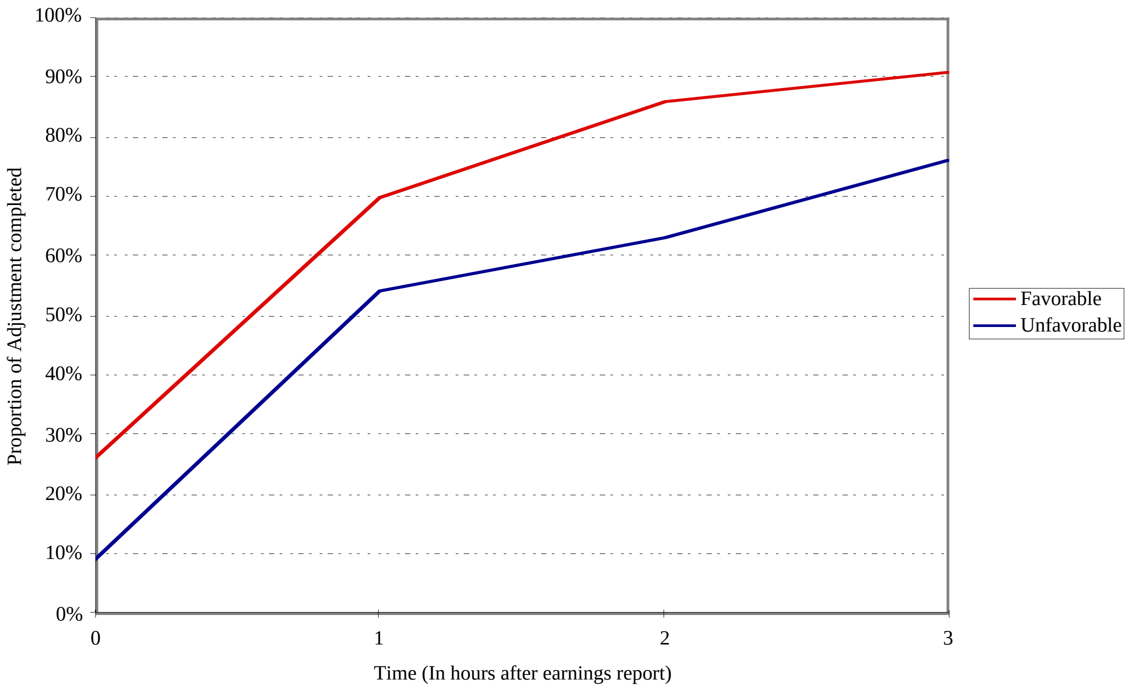

Source: Woodruff and Senchack

As Figure 10.13 illustrates, approximately 91% of the eventual adjustment occurs within three hours of the report for the most positive earnings surprises, while only 76% of the eventual adjustment occurs during the same period for the most negative earnings announcements. This would seem to indicate that markets are much more efficient about assessing good news than bad news. If nothing else, this also illustrates the importance of trading promptly after an earnings announcement. Investors who wait to read about the announcement the next day or even later in the day will find that the bulk of the adjustment has occurred by the time they trade.

Earnings Quality

In recent years, the strategy of investing on earnings surprises has come under some pressure because firms have learnt to play the earnings game. In fact, some firms like Microsoft, Intel and Cisco developed a reputation during the 1990s of being able to consistently beat analyst expectations. While they were rewarded with higher valuations, markets also began building into expectations the capacity of these firms to beat expectations. The earnings reported by these firms were compared not to analyst forecasts but to whispered earnings that were usually set higher than analyst forecasts, based upon past experience. Thus, Intel could report earnings per share of 57 cents, higher than the analyst forecast of 55 cents, but still see its stock price drop because the whispered earnings estimate for the firm was 58 cents.

As firms play the earnings game, the quality of earnings has also diverged across companies. A firm that beats earnings estimates because it has more efficient operating should be viewed more favorably than one that beats estimates because it changed the way it valued inventory. Does the market distinguish appropriately between the two? The evidence seems to indicate that it does not, at least on the date of the announcement, but that it eventually corrects for poor quality earnings. In a study of this phenomenon in 2001, Chan, Chan, Jegadeesh and Lakonishok examined firms that reported high accruals – i.e. the difference between accounting earnings and cash flows and argued that firms report high earnings without a matching increase in cashflow have poorer quality earnings. When they tracked a portfolio composed of these firms, they discovered that the high accrual year was usually the turning point in the fortunes of this firm, with subsequent years bring declining earnings and negative stock returns.

Can you make money off earnings announcements?

Financial markets get much of their firm-specific information from earnings announcements, and there are collectively thousands of earnings announcements each year. There are some portfolio managers whose investment strategies are based primarily or largely upon trading on or after these announcements. One strategy is to buy stocks that report large positive earnings surprises, hoping to benefit from the drift. The evidence indicates that across all stocks, the potential for excess returns from buying after earnings announcements is very small and may very well be non-existent after transactions costs.

How would you refine this strategy to harvest higher returns? You could draw on the empirical evidence and concentrate only on earnings announcements made by smaller, less liquid companies where the drift is more pronounced. In addition, you can try to direct your money towards companies with higher quality earnings surprises by avoiding firms with large accruals (i.e. , firms that report increasing earnings and decreasing cashflows).

Your potential for large returns is greatest if you can forecast which firms are most likely to report large positive earnings surprises and invest in those firms prior to their earnings announcements. Impossible, without insider information, you say. Not quite. You may be able to use a combination of quantitative techniques (time series models that forecast next quarter’s earnings based upon historical earnings) and trading volume (insiders do create blips in the volume) to try to detect these firms. Even if you are right only 55% of the time, you should be able to post high excess returns.

Buy on the rumor, sell on the news

While Wall Street adages should always be taken with a grain of salt, they usually have a kernel of truth to them. This particular one on rumor and news has particular relevance when we look at how prices run up before the news announcement and how little is left on the table after the announcement. An investor who has access to high quality gossip (if that is not an oxymoron) may be able to buy stocks before good news comes out and sell before bad news. But high quality gossip is difficult to come by, especially on Wall Street, where there a dozen false news stories that circulate for every true one.

Acquisitions

The announcements that usually carry the most significance for value are acquisition announcements, simply because of the scale of acquisitions, relative to other investments. Acquisitions are important not only because they affect values substantially but also because they occur often enough in the market to provide the basis for an investment strategy. In this section, we will begin with an analysis of how the announcement of an acquisition affects the market price of the target and acquiring firm on the day of the acquisition, follow up by looking at the post-acquisition performance (operating and stock price) of the acquiring firm and conclude with the question of whether there is anything in this process that can be exploited by an investor for gain.

The Acquisition Date

The big price movements associated with acquisitions occur around the date the acquisition is announced and not when it is actually consummated, though the latter may occur several months later. While much of the attention in acquisitions is focused on the target firms, we will argue that what happens to the acquiring firm is just as interesting, if not more.

Target Firms

The evidence indicates that the stockholders of target firms are the clear winners in takeovers — they earn signifcant excess returns[17] not only around the announcement of the acquisitions, but also in the weeks leading up to it. In 1983, Jensen and Ruback reviewed 13 studies that look at returns around takeover announcements and reported an average excess

---

[17] The excess returns around takeover announcements to target firms are so large that using different risk and return models seems to have no effect on the overall conclusions.

return of 30% to target stockholders in successful tender offers and 20% to target stockholders in successful mergers. In 1988, Jarrell, Brickley, and Netter examined the results of 663 tender offers made between 1962 and 1985 and noted that premiums averaged 19% in the 1960s, 35% in the 1970s and 30% between 1980 and 1985. The price behavior of a typical target firm in an acquisition is illustrated in figure 10.14, from one of the studies,[18] summarizes the target firm stock price in the 10 days before, the day of and the 10 days after an acquisition announcement.

Figure 10.14: Cumulative Excess Return to Target Company Stock

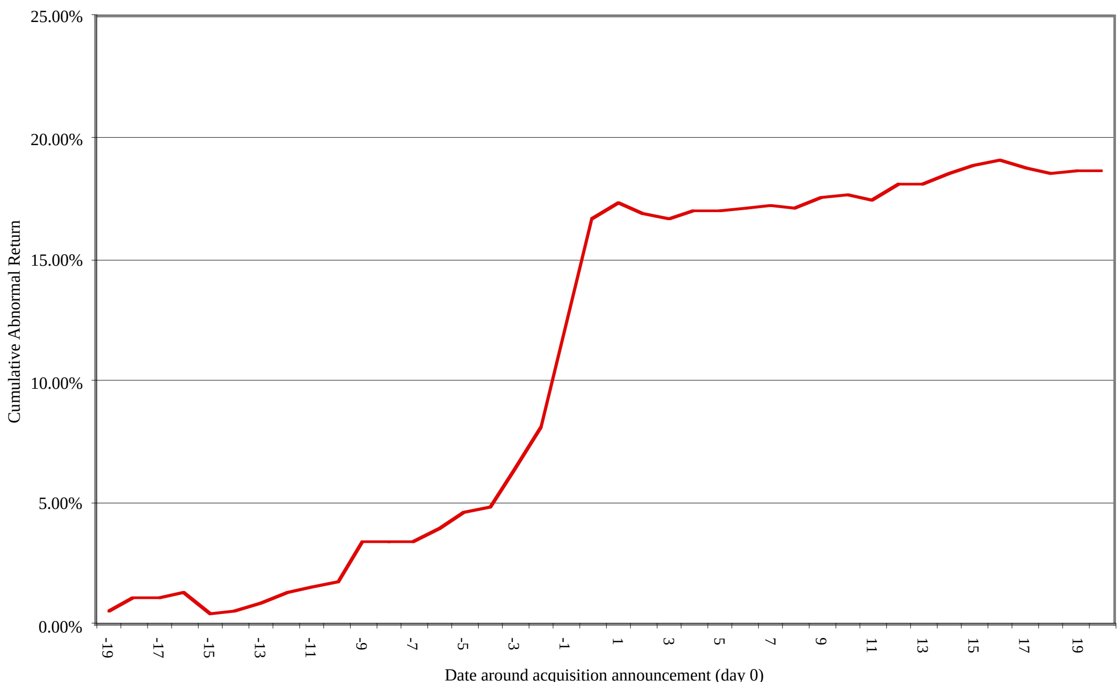

Note that about half the premium associated with the acquisition is already incorporated in the price by the time the acquisition is announced. This suggests that information about acquisitions is leaked to some investors who trade on that information. On the acquisition date, there is a decided jump in the stock price but little evidence of drift thereafter.

If we categorize acquisitions based upon how the acquiring firm pays for them, we find that the stock prices of target firms tend to do much better on the announcement of cash-based acquisitions (where the acquirer uses cash only to pay for the acquired company's stock) than stock based acquisitions. We also find that the premiums in hostile

---

[18] See Dennis and McConnell, 1986, Corporate Mergers and Security Returns, Journal of Financial Economics, v16, 143-188.

acquisitions are larger than the premiums on friendly mergers and that the premium on tender offers is slightly higher than the premium on mergers. Figure 10.15, extracted from one study[19], provides an illustration of the magnitude of the differences:

Target Firm Premiums in Acquisitions

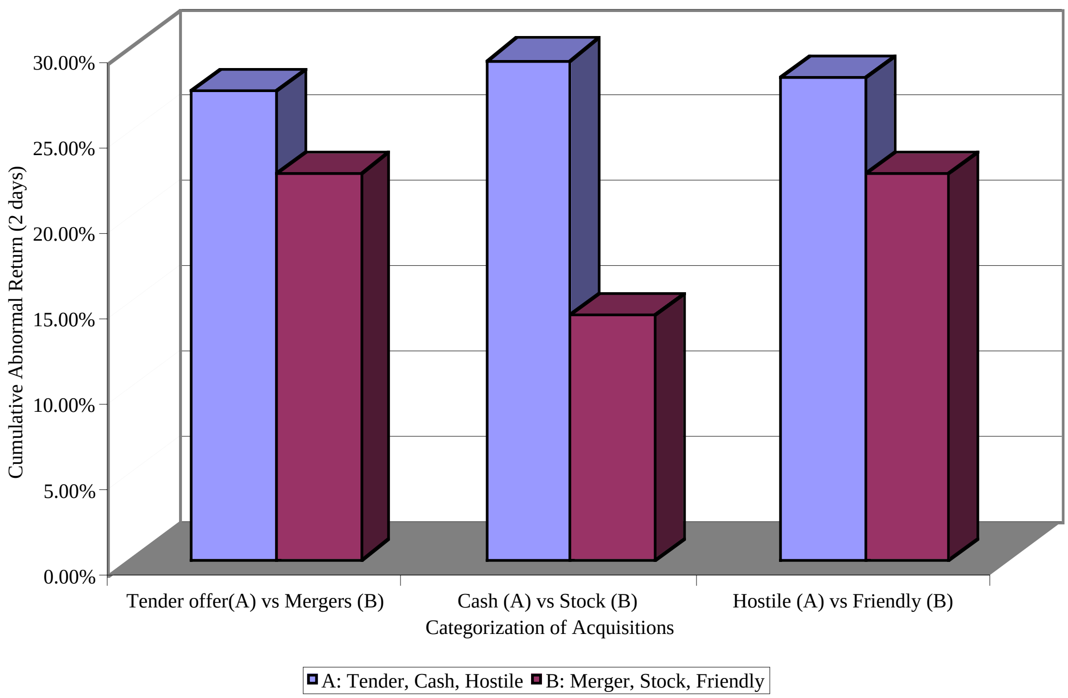

Source: Huang and Walkling

Some attempts at takeovers fail, either because the bidding firm withdraws the offer or because the target firm fights it off. Bradley, Desai, and Kim(1983) analyzed the effects of takeover failures on target firm stockholders and found that, while the initial reaction to the announcement of the failure is negative, albeit statistically insignificant, a substantial number of target firms are taken over within 60 days of the first takeover is failing, earning significant excess returns (50% to 66%).

Bidding Firms

The effect of takeover announcements on bidder firm stock prices is not as clear cut as it is for target firms. Jensen and Ruback report excess returns of 4% for bidding firm stockholders around tender offers and no excess returns around mergers. Jarrell, Brickley and Netter, in their examination of tender offers from 1962 to 1985, note a decline in excess

---

[19] See Huang and Walkling (1987), "Acquisition Announcements and Abnormal Returns". Journal of Financial Economic, v19, 329-350.

returns to bidding firm stockholders from 4.4% in the 1960s to 2% in the 1970s to -1% in the 1980s. Other studies indicate that approximately half of all bidding firms earn negative excess returns around the announcement of takeovers, suggesting that shareholders are skeptical about the perceived value of the takeover in a significant number of cases.

When an attempt at a takeover fails, Bradley, Desai and Kim (1983) report negative excess returns of 5% to bidding firm stockholders around the announcement of the failure. When the existence of a rival bidder in figured in, the studies indicate significant negative excess returns (of approximately 8%) for bidder firm stockholders who lose out to a rival bidder within 180 trading days of the announcement, and no excess returns when no rival bidder exists.

Considering the evidence, it is quite clear that bidding firm stockholders often do not share the enthusiasm that managers in these firms have about mergers and acquisitions. While managers would argue that this is because they are not privy to the information that is available only to insiders, we will see in the next section that many mergers fail and that stockholders are perhaps more prescient than managers.

After the Acquisition

Many studies examine the extent to which mergers and acquisitions succeed or fail after the firms combine. These studies generally conclude that mergers often fail to deliver on their promises of efficiency and synergy, and even those that do deliver seldom create value for the acquirers' stockholders.

McKinsey and Co. examined 58 acquisition programs between 1972 and 1983 for evidence on two questions: (1) Did the return on the amount invested in the acquisitions exceed the cost of capital? (2) Did the acquisitions help the parent companies outperform the competition? They concluded that 28 of the 58 programs failed both tests, and six failed at least one test. In a follow-up study[20] of 115 mergers in the U.K. and the U.S. in the 1990s, McKinsey concluded that 60% of the transactions earned returns on capital less than the cost of capital, and that only 23% earned excess returns. In 1999, KPMG examined 700 of the most expensive deals between 1996 and 1998 and concluded that only 17% created value for the combined firm, 30% were value neutral and 53% destroyed value[21].

---

[20] This study was referenced in an article titled "Merger Mayhem" that appeared in Barrons on April 20, 1998.

[21] KPMG measured the success at creating value by comparing the post-deal stock price performance of the combined firm to the performance of the relevant industry segment for a year after the deal was completed.

A study[22] looked at the eight largest bank mergers in 1995 and concluded that only two (Chase/Chemical, First Chicago/NBD) subsequently outperformed the bank-stock index. The largest, Wells Fargo's acquisition of First Interstate, was a significant failure. In an incisive book on the topic in 1996 titled "The Synergy Trap", Sirower took a detailed look at the promises and failures of synergy and drew the gloomy conclusion that synergy is often promised but seldom delivered.

The most damaging piece of evidence on the outcome of acquisitions is the large number of acquisitions that are reversed within fairly short time periods. Mitchell and Lehn note that 20.2% of the acquisitions made between 1982 and 1986 were divested by 1988. In a study published in1992, Kaplan and Weisbach found that 44% of the mergers they studied were reversed, largely because the acquirer paid too much or because the operations of the two firms did not mesh. Studies that have tracked acquisitions for longer time periods (ten years or more) have found the divestiture rate of acquisitions rises to almost 50%, suggesting that few firms enjoy the promised benefits from acquisitions do not occur. In another study,

Takeover-based Investment Strategies

There are three broad classes of investment strategies that can be constructed around takeovers. The first and most lucrative, if you can pull it off, is to find a way to invest in a target firm before the acquisition is announced. The second is to wait until after the takeover is announced and then try to take advantage of the price drift between the announcement date and the day the deal is consummated. This is often called risk arbitrage and we will take a closer look at it in the next chapter. The third is also a post-announcement strategy, but it is a long-term strategy where you invest in firms that you believe have the pieces in place to deliver the promised synergy or value creation.

Pre-announcement Investing

Looking at the stock price reaction of target firms both immediately prior to and immediately after the acquisition announcement, it is quite clear that the real money to be made in acquisitions comes from investing in firms before they become targets rather than after. Absent inside information, is this doable? There may be a way, and the answer lies in looking at firms that typically become target firms.

---

[22] This study was done by Keefe, Bruyette and Woods, an investment bank. It was referenced in an article titled "Merger Mayhem" in Barrons, April 20, 1998.

Research[23] indicates that the typical target firm in a hostile takeover has the following characteristics:

(1) It has under performed other stocks in its industry and the overall market, in terms of returns to its stockholders in the years preceding the takeover.

(2) It has been less profitable than firms in its industry in the years preceding the takeover.

(3) It has a much lower stock holding by insiders than do firms in its peer groups.

Other studies also provide tantalizing clues about typical target firms. Lang, Walkling and Stulz find, for instance, that stocks that trade at low market values, relative to their replacement costs (a low Tobin's Q) are much more likely to be taken over than firms that trade at high market values. The odds of being taken over also increase if the firm has a smaller market capitalization, doest not have shares with different voting classes and anti-takeover amendments on its books.

There are two ways in which we can use the findings of these studies to identify potential target firms. The first is to develop a set of screens that incorporate the variables mentioned above. You could, for instance, invest in firms with market capitalizations below $ 5 billion, with low insider holdings, depressed valuations (low price to book ratios) and low returns on equity. The second and slightly more sophisticated variant is to estimate the probability of being taken over for every firm in the market using statistical techniques.[24]

Post-announcement Investing

In this strategy, you buy companies after acquisitions or mergers are completed because you believe that they will be able to deliver what they promise at the time of the merger – higher earnings growth and synergy. As we noted in the earlier section on synergy, it shows up in relatively few mergers. Can we identify those mergers that are most likely to succeed and invest only in those? Again, the clues may lie in history.

Some studies find improvements in operating efficiency after mergers, especially hostile ones[25]. In a study in 1992, Healy, Palepu, and Ruback found that the median post-

---

[23] This research was also referenced in chapter 8. A paper by Bhide examines hostile takeovers and their aftermath.

[24] A probit, for instance, resembles a regression but estimates probabilities based upon specified independent variables. In this case, you could run a probit across firms in the market using the variables identified by earlier studies – low ROE, poor stock returns and low market cap – as independent variables. You will get as output the probability of being taken over for each firm in the market. You could follow up by constructing a portfolio of stocks where this probability is highest.

[25]A study by Healy, Palepu and Ruback (1989) looked at the post-merger performance of 50 large mergers

acquisition cash flow returns improve for firms involved in mergers, though 25% of merged firms lag industry averages after transactions. In 1999, Parrino and Harris examined 197 transactions between 1982 and 1987 and categorized the firms based upon whether the management is replaced (123 firms) at the time of the transaction, and the motive for the transaction. They find that

- On average, in the five years after the transaction, merged firms earned 2.1% more than the industry average.
- Almost all this excess return occurred in cases where the CEO of the target firm is replaced within one year of the merger. These firms earned 3.1% more than the industry average, whereas firms, whereas when the CEO of the target firm continued in place the merged firm did not do better than the industry

In addition, a few studies examine whether acquiring related businesses (i.e., synergy-driven acquisitions) provides better returns than acquiring unrelated business (i.e., conglomerate mergers) and come to conflicting conclusions with no consensus.[26] Nail and Megginson examined 260 stock swap transactions and categorized the mergers as either a conglomerate or a 'same-industry" transactions. They found no evidence of wealth benefits for either stockholders or bondholders in conglomerate transactions. However, they did find significant net gains for both stockholders and bondholders in the case of mergers of related firms.

Finally, on the issue of synergy, the KPMG study of the 700 largest deals from 1996 to 1998 concludes the following:

- Firms that evaluate synergy carefully before an acquisition are 28% more likely to succeed than firms that do not.
- Cost-saving synergies associated with reducing the number of employees are more likely to be accomplished than new product development or R&D synergies. For instance, only a quarter to a third of firms succeeded on the latter, whereas 66% of firms were able to reduce headcount after mergers.

Considering all the contradictory evidence contained in different studies[27], we conclude that:

---

from 1979 to 1983 and concluded that merged firms improved their operating performance (defined as EBITDA/Sales) relative to their industries.

[26] Michel and Shaked (1984) and Duofsky and Varadarajan (1987) find that diversification-driven mergers do better than synergy-driven mergers, in terms of risk-adjusted returns. Varadarajan and Ramanujam (1987) find that the latter do better in terms of return on equity.

[27] Some of this evidence is anecdotal and is based upon the study of just a few mergers.

- Mergers of equals (firms of equal size) seem to have a lower probability of succeeding than acquisitions of a smaller firm by a much larger firm[28].
- Cost saving mergers, where the cost savings are concrete and immediate, seem to have a better chance of delivering on synergy than mergers based upon growth synergy.
- Acquisition programs that focus on buying small private businesses for consolidations have had more success than acquisition programs that concentrate on acquiring publicly traded firms.
- Hostile acquisitions seem to do better at delivering improved post-acquisition performance than friendly mergers.

Other Announcements

While earnings and acquisition announcements may offer the best opportunity for trading profits for investors trading on information, the market reacts to other announcements made by firms as well.

Stock Splits

A stock split increases the number of shares outstanding, without changing the current earnings or cash flows of the firm. As a purely cosmetic event, a stock split should not affect the value of the firm or of outstanding equity. Rather, the price per share will go down to reflect the stock split, since there are more shares outstanding. One of the first event studies by Fama, Fisher, Jensen and Roll (Figure 10.16) examined the stock price reaction to 940 stock splits between 1927 and 1959 by cumulating excess returns in the 60 months around the actual split date.

[28] This might well reflect the fact that failures of mergers of equal are much more visible than failures of the small firm/large firm combinations.

Figure 10.16 : Market Reaction to Stock Splits

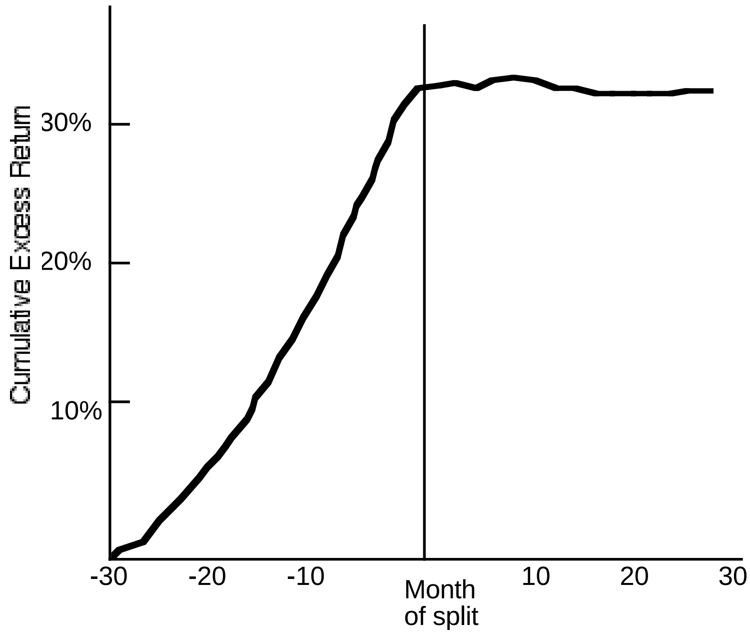

Source: Fama, Fisher, Jensen and Roll

On average, their study found that stock splits tended to follow periods of excess returns; this is not surprising, since splits typically follow price run-ups. They also found no evidence of excess returns around the splits themselves, suggesting that the splits were neutral events. One of the limitations of the study was its use of monthly return rather than daily returns. More recent studies that look at the daily price reaction to stock splits find a mild positive effect – stock prices go up when splits are announced. One study that looked at all two for one stock splits between 1975 and 1990 estimated that stock prices increase, on average, 3.38% on the announcement of a stock split and that the announcement effect is much greater for small stocks (10.04%) than for large stocks (1.01%).[29] Researchers attribute this to a signaling effect – i.e. that only companies that expect their stock prices to go up in the future will announce stock splits.

In recent years, a few studies have pointed out that stock splits may have an unintended negative effect on stockholders by raising transactions costs. For instance, the bid-ask spread[30], which is one component of the transactions costs, is a much larger percentage of the price for a $ 20 stock than it is for a $40 stock. Copeland (1979)

---

[29] Ikenberry, Rankine and Stice (1996) report that stocks that split continue to earn excess returns in the two years after the split – 7.93% in the first year and 12.15% in the second year.

[30] The bid-ask spread refers to the difference between the price at which a security can be bought (the ask price) or the sold (the bid price) at any point in time.

chronicles the increase in transactions costs and the decline in trading volume following splits. This additional cost has to be weighed off against the potential signaling implications of a stock split; investors may view a stock split as a positive signal about future prospects. This may explain the small positive returns some researchers have found around stock split announcement dates.[31]

Dividend Changes

Financial markets examine every action a firm takes for implications for future cash flows and firm value. When firms announce changes in dividend policy, they are conveying information to markets, whether they intend to or not. An increase in dividends is generally viewed as a positive signal, since firms that make these commitments to investors must believe that they have the capacity to generate these cash flows in the future. Decreasing dividends is a negative signal, largely because firms are reluctant to cut dividends. Thus, when a firm takes this action, markets see it as an indication that this firm is in substantial and long-term financial trouble. Consequently, such actions lead to a drop in stock prices. The empirical evidence concerning price reactions to dividend increases and decreases is consistent, at least on average, with this signaling theory. Figure 10.17 summarizes the average excess returns around dividend changes for firms.

[31] See Charest (1978) and Grinblatt, Masulis and Titman (1984).

Figure 10.17: Excess Returns around Announcements of Dividend Changes

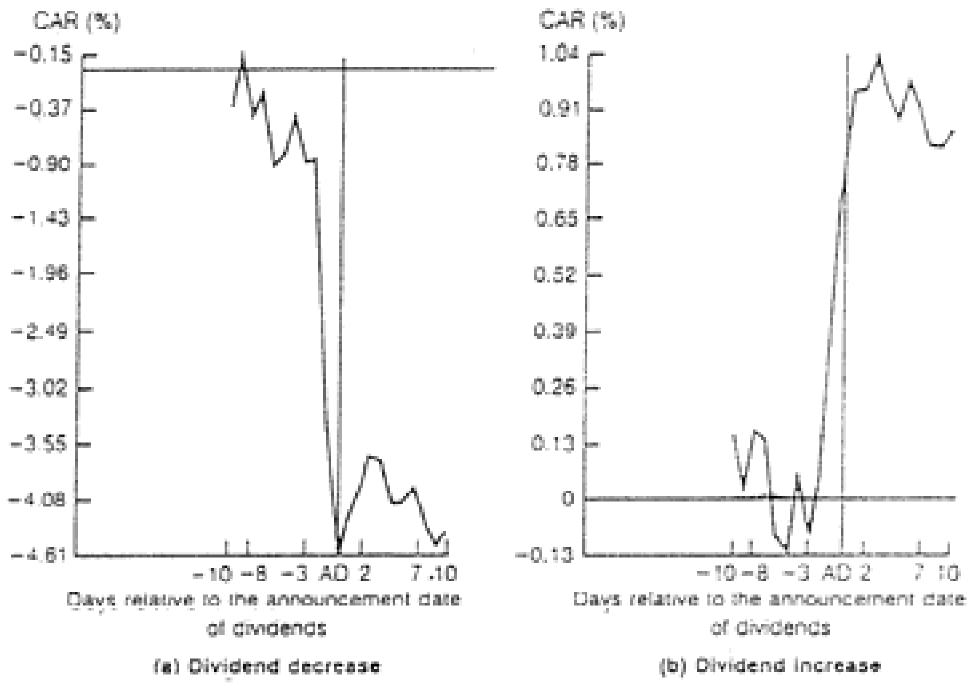

Source: Aharony and Swary

On average, stock prices go up when dividends are increased and go down when dividends are decreased, though the price reaction to the latter seems much more intense – a drop of more than 4.5% on dividend decreases and an increase of only about 1% on dividend increases.

While the price change on the dividend announcement itself might not offer opportunities for investors (unless they have access to inside information), another study by Michaely et al.) looked at the price drift after dividend changes are announced. He found that prices continue to drift up after dividend increases and drift down after dividend decreases for long periods. Investors may be able to take advantage of this drift and augment returns on their portfolios.

Implementing an Information-based Investment Strategy

If you decide to center your investment strategy around information releases – earnings reports, acquisition announcements or other news – you have to recognize that it is much more difficult to deliver the returns on an actual portfolio than it is in a hypothetical portfolio. To succeed at this strategy, you have to

- *Identify the information around which your strategy will be built*: Since you have to trade on the announcement, it is critical that you determine in advance the information that will trigger a trade. To provide an example, you may conclude that your best potential for returns comes from buying small companies that report earnings that are much higher than expected. However, you have to go further and specify what constitutes a small company (market cap less than $ 1 billion? $ 5 billion?) and by how much the actual earnings need to beat expectations (10% higher than expectations? 20% higher than expectations?). This is necessary because you will not have to the time for analysis after the report comes out.

- *Invest in an information system that will deliver the information to you instantaneous*: Many individual investors receive information with a time lag – 15 to 20 minutes after it reaches the trading floor and institutional investors. While this may not seem like a lot of time, the biggest price changes after information announcements occur during these periods.
- *Execute quickly*: Getting an earnings report or an acquisition announcement in real time is of little use if it takes you 20 minutes to trade. Immediate execution of trades is essential to succeeding with this strategy.
- *Keep a tight lid on transactions costs*: Speedy execution of trades usually goes with higher transactions costs, but these transactions costs can very easily wipe out any potential you may see for excess returns, especially because you will be trading a lot (as with any short-term strategy).
- *Know when to sell*: Almost as critical as knowing when to buy is knowing when to sell, since the price effects of news releases may begin to fade or even reverse after a while. Thus, if you buy firms after positive earnings announcements, you have to determine when you will sell their stock at the time you buy it. While this may seem to take away your flexibility, the alternative of holding on to stocks too long, hoping that they will go up can be even more damaging.

If you consider the requirements for success – immediate access to information, and instantaneous and cheap execution – it is not surprising that information-based investing was for long profitable only for institutional investors. In recent years, however, online access to information and trading has made it feasible for individual investors to join the party. This is a mixed blessing, though, since the more investors trade on information, the less returns there are from these strategies.

Conclusion

As investors, we all dream of receiving that important news release ahead of the market and making lucrative profits on it. Information is the key to investment success, and this chapter explores the possibility of acquiring and using information to augment portfolio returns. We began by looking at how market prices move in response to information, and noted that in an efficient market, the price reaction to new information is instantaneous. In such a market, investing in an asset after the information has been released is a neutral strategy. In an inefficient market, you can make money after the information is released, buying after good news, if it is a slow learning market, or selling after good news, if it is an overreacting market.

To examine whether it is possible to use information profitably, we first looked at the two groups of individuals most likely to have access to privileged or private information. Insiders in firms, especially top managers, clearly know more about their firms than investors in markets. Insider trading does seem to provide a signal of future price movements – insider buying seems to precede stock price increases and selling seems to occur ahead of price drops – but the signal is noisy and the returns are small. This may, however, reflect the fact that the really profitable insider trades – the illegal ones – are never filed with the SEC. Equity research analysts also have access to information that most other investors do not have and reflect this information in earnings forecasts and recommendations. Here again, while upwards revisions in earnings and buy recommendations generally precede stock price increases and downward revisions and sell recommendations precede poor stock price performance, the returns are surprisingly small. The difficulty that both insiders and analysts have in converting information to returns should be a cautionary note to any investor considering an information-based investment strategy.

In the second part of the chapter, we looked at earnings reports, acquisition announcements and other firm-specific announcements. These announcements affect prices significantly – positive (negative) earnings reports are associated with price increases (decreases), target company stock prices jump on acquisition announcements and stock prices generally increase when there are stock splits or dividend increases. Unless you can anticipate these news releases, though, this price increase cannot be translated into a large profit. There does seem to be some evidence of a price drift after earnings announcements, and there are investors who try to take advantage of this drift by buying after positive earnings reports and selling after negative reports. To succeed with information based trading, you have to be selective and disciplined in your investment choices and efficient in your execution.

Lessons for Investors

To be a successful trader on information, you need to:

1. Find a reliable source of information: It goes without saying that good information is the key to success with any information-based trading strategy. (To stay on the right side of the law, make sure that your reliable source is not an insider.)
2. Have a clearly defined strategy for trading on information: Since you will have to trade quickly, you will not have the time after information comes out to assess and analyze it. You will need to make a pre-judgment on when you will be trading.
3. Be disciplined: Don't deviate from your trading strategy and stick to the time horizon that you have chosen for yourself. Holding on to a stock for a few days more hoping to recoup your losses can make a bad situation worse.
4. Control your trading costs: Since you will be trading frequently and immediate execution is key, your trading costs can be large. As the funds at your disposal increase, the price impact you have as you trade can be substantial.

CHAPTER 11

A SURE PROFIT: THE ESSENCE OF ARBITRAGE

Arbitrage represents the holy grail of investing because it allows investors to invest no money, take no risk and walk away with sure profits. In other words, it is the ultimate money machine that investors hope to access.

In this chapter, we consider three types of arbitrage. The first is pure arbitrage, where, in fact, you risk nothing and earn more than the riskless rate. For pure arbitrage to be feasible, you need two assets with identical cashflows, different market values at the same point in time and a given point in time in the future at which the values have to converge. This type of arbitrage is most likely to occur in derivatives markets – options and futures- and in some parts of the bond market. The second is near arbitrage, where you have assets that have identical or almost identical cash flows, trading at different prices, but there is no guarantee that the prices will converge and there exist significant constraints on the investors forcing convergence. The third is speculative arbitrage, which may not really be arbitrage in the first place. Here, investors take advantage of what they see as mispriced and similar (though not identical) assets, buying the cheaper one and selling the more expensive one. If they are right, the difference should narrow over time, yielding profits. It is in this category that we consider hedge funds in their numerous forms. As we will see, the peril of this strategy is that the initial assessment of mispricing is usually based upon a view of the world that may or may not be justified.

Pure Arbitrage

The requirement that you have two assets with identical cashflows and different market prices makes pure arbitrage difficult to find in financial markets. First, identical assets are not common in the real world, especially if you are an equity investor. Second, assuming two identical assets exist, you have to wonder why financial markets would allow pricing differences to persist. If in addition, we add the constraint that there is a point in time where the market prices converge, it is not surprising that pure arbitrage is most likely to occur with derivative assets – options and futures and in fixed income markets, especially with default-free government bonds.

Futures Arbitrage

A futures contract is a contract to buy (and sell) a specified asset at a fixed price in a future time period. There are two parties to every futures contract - the seller of the contract, who agrees to deliver the asset at the specified time in the future, and the buyer of the contract, who agrees to pay a fixed price and take delivery of the asset. If the asset that underlies the futures contract is traded and is not perishable, you can construct a pure arbitrage if the futures contract is mispriced. In this section, we will consider the potential for arbitrage first with storable commodities and then with financial assets and then look at whether such arbitrage is possible.

The Arbitrage Relationships

The basic arbitrage relationship can be derived fairly easily for futures contracts on any asset, by estimating the cashflows on two strategies that deliver the same end result – the ownership of the asset at a fixed price in the future. In the first strategy, you buy the futures contract, wait until the end of the contract period and buy the underlying asset at the futures price. In the second strategy, you borrow the money and buy the underlying asset today and store it for the period of the futures contract. In both strategies, you end up with the asset at the end of the period and are exposed to no price risk during the period – in the first, because you have locked in the futures price and in the second because you bought the asset at the start of the period. Consequently, you should expect the cost of setting up the two strategies to exactly the same. Across different types of futures contracts, there are individual details that cause the final pricing relationship to vary – commodities have to be stored and create storage costs whereas stocks may pay a dividend while you are holding them.

a. Storable Commodities

The distinction between storable and perishable goods is that storable goods can be acquired today at the spot price and stored till the expiration of the futures contract, which is the practical equivalent of buying a futures contract and taking delivery at expiration. Since the two approaches provide the same result, in terms of having possession of the commodity at expiration, the futures contract, if priced right, should cost the same as a strategy of buying and storing the commodity. The two additional costs of the latter strategy are as follows.

(a) Since the commodity has to be acquired now, rather than at expiration, there is an added financing cost associated with borrowing the funds needed for the acquisition now.

$$\text{Added Interest Cost} = (\text{Spot price})\left((1+\text{Interest Rate})^{\text{Life of Futures contract}} - 1\right)$$

(b) If there is a storage cost associated with storing the commodity until the expiration of the futures contract, this cost has to be reflected in the strategy as well. In addition, there may be a benefit to having physical ownership of the commodity. This benefit is called the convenience yield and will reduce the futures price. The net storage cost is defined to be the difference between the total storage cost and the convenience yield.

If F is the futures contract price, S is the spot price, r is the annualized interest rate, t is the life of the futures contract and k is the net annual storage costs (as a percentage of the spot price) for the commodity, the two equivalent strategies and their costs can be written as follows.

Strategy 1: Buy the futures contract. Take delivery at expiration. Pay \$F.

Strategy 2: Borrow the spot price (S) of the commodity and buy the commodity. Pay the additional costs.

(a) Interest cost $= S\left((1+r)^t - 1\right)$

(b) Cost of storage, net of convenience yield = S k t

If the two strategies have the same costs,

$$F^* \begin{array}{l} = S\left((1+r)^t - 1\right) + Skt \\ = S\left((1+r)^t + kt\right) \end{array}$$

This is the basic arbitrage relationship between futures and spot prices. Note that the futures price does not depend upon your expectations of what will happen to the spot price over time but on the spot price today. Any deviation from this arbitrage relationship should provide an opportunity for arbitrage, i.e., a strategy with no risk and no initial investment, and for positive profits. These arbitrage opportunities are described in Figure 11.1.

This arbitrage is based upon several assumptions. First, investors are assumed to borrow and lend at the same rate, which is the riskless rate. Second, when the futures contract is over priced, it is assumed that the seller of the futures contract (the arbitrageur) can sell short on the commodity and that he can recover, from the owner of the commodity, the storage costs that are saved as a consequence. To the extent that these assumptions are unrealistic, the bounds on prices within which arbitrage is not feasible expand. Assume, for instance, that the rate of borrowing is r_b and the rate of lending is r_a, and that short seller cannot recover any of the saved storage costs and has to pay a transactions cost of t_s. The futures price will then fall within a bound.

$$(S - t_s)(1 + r_a) < F^* < S((1 + r_b)^t + kt)$$

If the futures price falls outside this bound, there is a possibility of arbitrage and this is illustrated in Figure 11.2.

Figure 11.1: Storable Commodity Futures: Pricing and Arbitrage

$$F^* = S\left((1+r)^t + k\,t\right)$$

	If F > F*		If F < F*	
Time	Action	Cashflows	Action	**Cashflows**
Now:	1. Sell futures contract	0	1. Buy futures contract	0
	2. Borrow spot price at riskfree r	S	2. Sell short on commodity	S
	3. Buy spot commodity	-S	3. Lend money at riskfree rate	-S
At t:	1. Collect commodity; Pay storage cost.	-Skt	1. Collect on loan	$S(1+r)^t$
	2. Deliver on futures contract	F	2. Take delivery of futures contract	-F
	3. Pay back loan	$-S(1+r)^t$	3. Return borrowed commodity; Collect storage costs	+Skt
Net Cash Flow=		$\mathbf{F-S((1+r)^t - kt) > 0}$		$\mathbf{S((1+r)^t + kt) - F > 0}$

Key inputs:

F* = Theoretical futures price

F = Actual futures price

S = Spot price of commodity

r= Riskless rate of interest (annualized)

t = Time to expiration on the futures contract

k = Annualized carrying cost, net of convenience yield (as % of spot price)

Key assumptions

1. The investor can lend and borrow at the riskless rate.
2. There are no transactions costs associated with buying or selling short the commodity.
3. The short seller can collect all storage costs saved because of the short selling.

Figure 11.2: Storable Commodity Futures: Pricing and Arbitrage with Modified Assumptions

Modified Assumptions

1. Investor can borrow at r_b ($r_b > r$) and lend at r_a ($r_a < r$).
2. The transactions cost associated with selling short is t_s (where t_s is the dollar transactions cost).
3. The short seller does not collect any of the storage costs saved by the short selling.

$$F_h^* = S\left((1+r_b)^t + k\,t\right)$$
$$F_l^* = (S-t_s)\,(1+r_a)^t$$

If $F > F_h^*$ | If $F < F_l^*$

Time	Action (If $F > F_h^*$)	Cashflows	Action (If $F < F_l^*$)	Cashflows
Now:	1. Sell futures contract	0	1. Buy futures contract	0
	2. Borrow spot price at r_b	S	2. Sell short on commodity	$S - t_s$
	3. Buy spot commodity	-S	3. Lend money at r_a	$-(S - t_s)$
At t:	1. Collect commodity from storage	-Skt	1. Collect on loan	$(S-t_s)(1+r_a)^t$
	2. Delivery on futures contract	F	2. Take delivery of futures contract	-F
	3. Pay back loan	$-S(1+r_b)^t$	3. Return borrowed commodity; Collect storage costs	0
Net Cash Flow=		$\mathbf{F-S((1+r_b)^t - kt) > 0}$		$\mathbf{(S-t_s)\,(1+r_a)^t - F > 0}$

F_h = Upper limit for arbitrage bound on futures prices

F_l = Lower limit for arbitrage bound on futures prices *b. Stock Index Futures*

Futures on stock indices have become an important and growing part of most financial markets. Today, you can buy or sell futures on the Dow Jones, the S&P 500, the NASDAQ and the Value Line indices, as well as many indices in other countries. An index future entitles the buyer to any appreciation in the index over and above the index futures price and the seller to any depreciation in the index from the same benchmark. To evaluate the arbitrage pricing of an index future, consider the following strategies.

Strategy 1: Sell short on the stocks in the index for the duration of the index futures contract. Invest the proceeds at the riskless rate. This strategy requires that the owners of the stocks that are sold short be compensated for the dividends they would have received on the stocks.

Strategy 2: Sell the index futures contract.

Both strategies require the same initial investment, have the same risk and should provide the same proceeds. Again, if S is the spot price of the index, F is the futures prices, y is the annualized dividend yield on the stock and r is the riskless rate, the arbitrage relationship can be written as follows:

$$F^* = S\,(1 + r - y)^t$$

If the futures price deviates from this arbitrage price, there should be an opportunity from arbitrage. This is illustrated in Figure 11.3.

This arbitrage is also conditioned on several assumptions. First, we assume that investors can lend and borrow at the riskless rate. Second, we ignore transactions costs on both buying stock and selling short on stocks. Third, we assume that the dividends paid on the stocks in the index are known with certainty at the start of the period. If these assumptions are unrealistic, the index futures arbitrage will be feasible only if prices fall outside a band, the size of which will depend upon the seriousness of the violations in the assumptions.

Assume that investors can borrow money at r_b and lend money at r_a and that the transactions costs of buying stock is t_c and selling short is t_s. The band within which the futures price must stay can be written as:

$$(S - t_s)(1 + ra - y)^t < F^* < (S + t_c)(1 + rb - y)^t$$

The arbitrage that is possible if the futures price strays outside this band is illustrated in Figure 11.4.

In practice, one of the issues that you have to factor in is the seasonality of dividends since the dividends paid by stocks tend to be higher in some months than others. Figure

11.5 graphs out dividends paid as a percent of the S&P 500 index on U.S. stocks in 2000 by month of the year.

Figure 11.5: Dividend Yields by Month of Year- 2000

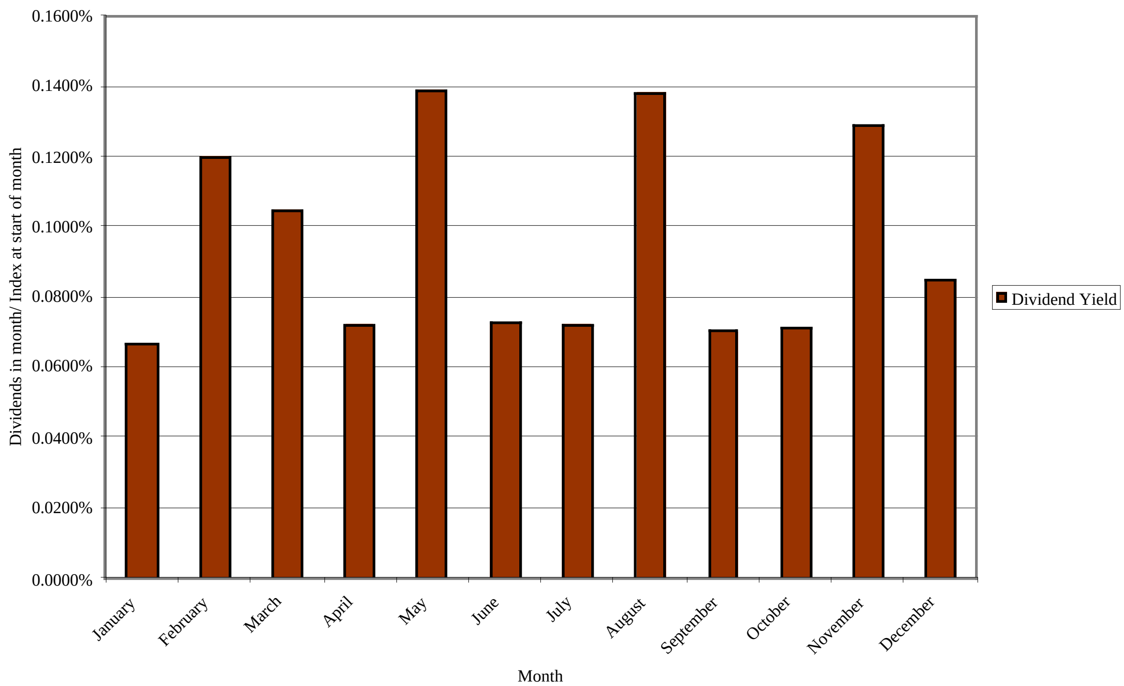

Thus, dividend yields seem to peak in February, May, August and November. An index future coming due in these months is much more likely to be affected by dividend yields especially as maturity draws closer.

Figure 11.3: Stock Index Futures: Pricing and Arbitrage

$F^* = S\ (1+r-y)^t$

·*		If F < F*	
on	Cashflows	Action	**Cashflows**
contract	0	1. Buy futures contract	0
price of index at riskfree r	S	2. Sell short stocks in the index	S
n index	-S	3. Lend money at riskfree rate	-S
lends on stocks	$S((1+y)^t-1)$	1. Collect on loan	$S(1+r)^t$
futures contract	F	2. Take delivery of futures contract	-F
ın	$-S(1+r)^t$	3. Return borrowed stocks; Pay foregone dividends	$- S((1+y)^t-1)$
	$\mathbf{F-S(1+r-y)^t > 0}$		$\mathbf{S\ (1+r-y)^t - F > 0}$

ice r= Riskless rate of interest (annualized)

t = Time to expiration on the futures contract

y = Dividend yield over lifetime of futures contract as % of current index level

ıd borrow at the riskless rate.

s costs associated with buying or selling short stocks.

ith certainty.

Figure 11.4: Stock Index Futures: Pricing and Arbitrage with modified assumptions
Modified Assumptions

b ($r_b > r$) and lend at r_a ($r_a < r$).
sociated with selling short is t_s (where t_s is the dollar transactions cost) and the transactions cost associated wit

$$F_h^* = (S+t_c)(1+r_b-y)^t$$
$$F_l^* = (S-t_s)(1+r_a-y)^t$$

h*

If $F < F_l^*$

n	Cashflows	Action	Cashflows
contract	0	1. Buy futures contract	0
t price at r_b	$S+t_c$	2. Sell short stocks in the index	$S - t_s$
n the index	$-S-t_c$	3. Lend money at r_a	$-(S - t_s)$
dends on stocks	$S((1+y)^t-1)$	1. Collect on loan	$(S-t_s)(1+r_a)^t$
futures contract	F	2. Take delivery of futures contract	$-F$
an	$-(S+t_c)(1+r_b)^t$	3. Return borrowed stocks;	
		Pay foregone dividends	$-S((1+y)^t-1)$
	$\mathbf{F-(S+t_c)(1+r_b-y)^t > 0}$		$\mathbf{(S-t_s)(1+r_a-y)^t - F > 0}$

age bound on futures prices

F_l = Lower limit for arbitrage bound on futures prices *c. Treasury Bond Futures*

The treasury bond futures traded on the Chicago Board of Trade require the delivery of any government bond with a maturity greater than fifteen years, with a no-call feature for at least the first fifteen years. Since bonds of different maturities and coupons will have different prices, the CBOT has a procedure for adjusting the price of the bond for its characteristics. The conversion factor itself is fairly simple to compute and is based upon the value of the bond on the first day of the delivery month, with the assumption that the interest rate for all maturities equals 8% per annum (with semi-annual compounding). For instance, you can compute the conversion factor for a 9% coupon bond with 18 years to maturity. Working in terms of a $100 face value of the bond, the value of the bond can be written as follows, using the interest rate of 8%.

$$\text{PV of Bond} = \sum_{t=0.5}^{t=20} \frac{4.50}{(1.08)^t} + \frac{100}{(1.08)^{20}} = \$111.55$$

The conversion factor for this bond is 111.55. Generally, the conversion factor will increase as the coupon rate increases and with the maturity of the delivered bond.

This feature of treasury bond futures, i.e., that any one of a menu of treasury bonds can be delivered to fulfill the obligation on the bond, provides an advantage to the seller of the futures contract. Naturally, the cheapest bond on the menu, after adjusting for the conversion factor, will be delivered. This *delivery option* has to be priced into the futures contract. There is an additional option embedded in treasury bond futures contracts that arises from the fact that the T.Bond futures market closes at 2 p.m., whereas the bonds themselves continue trading until 4 p.m. The seller does not have to notify the clearing house until 8 p.m. about his intention to deliver. If bond prices decline after 2 p.m., the seller can notify the clearing house of his intention to deliver the cheapest bond that day. If not, the seller can wait for the next day. This option is called the *wild card option*.

The valuation of a treasury bond futures contract follows the same lines as the valuation of a stock index future, with the coupons of the treasury bond replacing the dividend yield of the stock index. The theoretical value of a futures contract should be –

$$F^* = (S - PVC)(1 + r)^t$$

where,

F* = Theoretical futures price for Treasury Bond futures contract

S = Spot price of Treasury bond

PVC = Present Value of coupons during life of futures contract

r = Riskfree interest rate corresponding to futures life

t = Life of the futures contract

If the futures price deviates from this theoretical price, there should be the opportunity for arbitrage. These arbitrage opportunities are illustrated in Figure 11.6.

This valuation ignores the two options described above - the option to deliver the cheapest-to-deliver bond and the option to have a wild card play. These give an advantage to the seller of the futures contract and should be priced into the futures contract. One way to build this into the valuation is to use the cheapest deliverable bond to calculate both the current spot price and the present value of the coupons. Once the futures price is estimated, it can be divided by the conversion factor to arrive at the standardized futures price.

d. Currency Futures

In a currency futures contract, you enter into a contract to buy a foreign currency at a price fixed today. To see how spot and futures currency prices are related, note that holding the foreign currency enables the investor to earn the risk-free interest rate (R_f) prevailing in that country while the domestic currency earn the domestic riskfree rate (R_d). Since investors can buy currency at spot rates and assuming that there are no restrictions on investing at the riskfree rate, we can derive the relationship between the spot and futures prices. *Interest rate parity* relates the differential between futures and spot prices to interest rates in the domestic and foreign market.

$$\frac{\text{Futures Price}_{d,f}}{\text{Spot Price}_{d,f}} = \frac{(1+R_d)}{(1+R_f)}$$

where Futures Price$_{d,f}$ is the number of units of the domestic currency that will be received for a unit of the foreign currency in a forward contract and Spot Price$_{d,f}$ is the number of units of the domestic currency that will be received for a unit of the same foreign currency in a spot contract. For instance, assume that the one-year interest rate in the United States is 5% and the one-year interest rate in Germany is 4%. Furthermore, assume that the spot exchange rate is \$0.65 per Deutsche Mark. The one-year futures price, based upon interest rate parity, should be as follows:

$$\frac{\text{Futures Price}_{d,f}}{\$\,0.65} = \frac{(1.05)}{(1.04)}$$

This results in a futures price of \$0.65625 per Deutsche Mark.

Why does this have to be the futures price? If the futures price were greater than \$0.65625, say \$0.67, an investor could take advantage of the mispricing by selling the futures contract, completely hedging against risk and ending up with a return greater than the riskfree rate. The actions the investor would need to take are summarized in Table 11.1, with the cash flows associated with each action in brackets next to the action.

Table 11.1: Arbitrage when currency futures contracts are mispriced

Forward Rate Mispricing	*Actions to take today*	*Actions at expiration of futures contract*
If futures price > \$0.65625 e.g. \$0.67	1. Sell a futures contract at \$0.67 per Deutsche Mark. (\$0.00) \$ 0.00 2. Borrow the spot price in the U.S. domestic markets @ 5%.	1. Collect on Deutsche Mark investment. (+1.04 DM) 1.04 DM 2. Convert into dollars at futures price. (-1.04 DM/

	(+$0.65) + $ 0.65 3. Convert the dollars into Deutsche Marks at spot price. (-$0.65/+1 DM) - $ 0.65/+ 1 DM 4. Invest Deustche Marks in the German market @ 4%. (-1 DM) - 1 DM	+$0.6968) -1.04 DM to + $ 0.6968 3. Repay dollar borrowing with interest. (-$0.6825) **Profit = $0.6968 - $0.6825 = $ 0.0143**
If futures price < $0.65625 e.g. $0.64	1. Buy a futures price at $0.64 per Deutsche Mark. ($0.00) $ 0.00 2. Borrow the spot rate in the German market @4%. (+1 DM) + 1 DM 3. Convert the Deutsche Marks into Dollars at spot rate. (-1 DM/+$0.65) - 1 DM/ $ 0.65 4. Invest dollars in the U.S. market @ 5%. (-$0.65) - $ 0.65	1. Collect on Dollar investment. (+$0.6825) $ 0.6825 2. Convert into dollars at futures price. (-$0.6825/1.0664 DM) - $ 0.6825 /+1.0664 DM 3. Repay DM borrowing with interest. (1.04 DM) **Profit = 1.0664-1.04 = 0.0264 DM** - 1.04 DM + 0.0264 DM

The first arbitrage of Table 11.1 results in a riskless profit of $0.0143, with no initial investment. The process of arbitrage will push down futures price towards the equilibrium price.

If the futures price were lower than $0.65625, the actions would be reversed, with the same final conclusion. Investors would be able to take no risk, invest no money and still end up with a positive cash flow at expiration. In the second arbitrage of Table 34.3, we lay out the actions that would lead to a riskless profit of .0264 DM.

Special Features of Futures Markets

There are two special features of futures markets that can make arbitrage tricky. The first is the existence of margins. While we assumed, when constructing the arbitrage, that buying and selling futures contracts would create no cashflows at the time of the transaction, you would have to put up a portion of the futures contract price (about 5-10%) as a margin in the real world. To compound the problem, this margin is recomputed every day based upon futures prices that day – this process is called <u>marking to market</u> - and you may be required to come up with more margin if the price moves against you (down, if you are a buyer and up, if you are a seller). If this margin call is not met, your position can be liquidated and you may never to get to see your arbitrage profits.

The second is that the futures exchanges generally impose 'price movement limits' on most futures contracts. For instance, the daily price movement limit on orange juice futures contract on the New York Board of Trade is 5 cents per pound or $750 per contract (which covers 15,000 pounds). If the price of the contract drops or increases by the amount of the price limit, trading is generally suspended for the day, though the exchange reserves the discretion to reopen trading in the contract later in the day. The rationale for introducing price limits is to prevent panic buying and selling on an asset, based upon faulty information or rumors, and to prevent overreaction to real information. By allowing investors more time to react to extreme information, it is argued, the price reaction will be more rational and reasoned. In the process, though, you can create a disconnect between the spot markets, where no price limits exist, and futures markets, where they do.

Price Limits and Contract specifications: Take a look at the price limits and contract specifications on widely traded futures contracts.

Feasibility of Arbitrage and Potential for Success

If futures arbitrage is so simple, you may ask, how in a reasonably efficient market would arbitrage opportunities even exist? In the commodity futures market, for instance, Garbade and Silber (1983) find little evidence of arbitrage opportunities and their findings are echoed in other studies. In the financial futures markets, there is evidence that indicates that arbitrage is indeed feasible but only to a sub-set of investors. Differences in transactions cost seem to explain most of the differences. Large institutional investors, with close to zero transactions costs and instantaneous access to both the underlying asset and futures markets may be able to find arbitrage opportunities, where you and I as individual investors would not. In addition, these investors are also more likely to meet the requirements for arbitrage – being able to borrow at rates close to the riskless rate and sell short on the underlying asset.

Note, though, that the returns are small[1] even to these large investors and that arbitrage will not be a reliable source of profits, unless you can establish a competitive advantage on one of three dimensions. First, you can try to establish a transactions cost

[1] A study of 835 index arbitrage trades on the S&P 500 futures contracts estimated that the average gross profit from such trades was only 0.30%.

advantage over other investors, which will be difficult to do since you are competing with other large institutional investors. Second, you may be able to develop an information advantage over other investors by having access to information earlier than others. Again, though much of the information is pricing information and is public. Third, you may find a quirk in the data or pricing of a particular futures contract before others learn about it. The arbitrage possibilities seem to be greater when futures contracts are first introduced on an asset, since investors take time to understand the details of futures pricing. For instance, it took investors a while to learn to incorporate the effect of uneven dividends into stock index futures and the wild card option into treasury bond futures. Presumably, investors who learnt faster than the market would have been able to take advantage of the mispricing of futures contracts in these early periods and earn excess returns.

Options Arbitrage

As derivative securities, options differ from futures in a very important respect. They represent rights rather than obligations – calls gives you the right to buy and puts gives you the right to sell. Consequently, a key feature of options is that the losses on an option position are limited to what you paid for the option, if you are a buyer. Since there is usually an underlying asset that is traded, you can, as with futures contracts, construct positions that essentially are riskfree by combining options with the underlying asset.

Exercise Arbitrage

The easiest arbitrage opportunities in the option market exist when options violate simple pricing bounds. No option, for instance, should sell for less than its exercise value.

With a call option: Value of call > Value of Underlying Asset – Strike Price

With a put option: Value of put > Strike Price – Value of Underlying Asset

For instance, a call option with a strike price of $ 30 on a stock that is currently trading at $ 40 should never sell for less than $ 10. If it did, you could make an immediate profit by buying the call for less than $ 10 and exercising right away to make $ 10.

In fact, you can tighten these bounds for call options, if you are willing to create a portfolio of the underlying asset and the option and hold it through the option’s expiration. The bounds then become:

With a call option: Value of call > Value of Underlying Asset – Present value of Strike Price

With a put option: Value of put > Present value of Strike Price – Value of Underlying Asset

Too see why, consider the call option in the previous example. Assume that you have one year to expiration and that the riskless interest rate is 10%.

Present value of Strike Price = $ 30/1.10 = $27.27

Lower Bound on call value = $ 40 - $27.27 = $12.73

The call has to trade for more than $12.73. What would happen if it traded for less, say $ 12? You would buy the call for $ 12, sell short a share of stock for $ 40 and invest the net proceeds of $ 28 ($40 – 12) at the riskless rate of 10%. Consider what happens a year from now:

If the stock price > strike price ($ 30): You first collect the proceeds from the riskless investment ($28(1.10) =$30.80), exercise the option (buy the share at $ 30) and cover your short sale. You will then get to keep the difference of $0.80.

If the stock price < strike price ($ 30): You collect the proceeds from the riskless investment ($30.80), buy a share in the open market for the prevailing price then (which is less than $30) and keep the difference.

In other words, you invest nothing today and are guaranteed a positive payoff in the future. You could construct a similar example with puts.

The arbitrage bounds work best for non-dividend paying stocks and for options that can be exercised only at expiration (European options). Most options in the real world can be exercised prior to expiration (American options) and are on stocks that pay dividends. Even with these options, though, you should not see short term options trading violating these bounds by large margins, partly because exercise is so rare even with listed American options and dividends tend to be small. As options become long term and dividends become larger and more uncertain, you may very well find options that violate these pricing bounds, but you may not be able to profit off them.

Replicating Portfolio

One of the key insights that Fischer Black and Myron Scholes had about options in the 1970s that revolutionized option pricing was that a portfolio composed of the underlying asset and the riskless asset could be constructed to have exactly the same cash flows as a call or put option. This portfolio is called the replicating portfolio. In fact, Black and Scholes used the arbitrage argument to derive their option pricing model by noting that since the replicating portfolio and the traded option have the same cash flows, they would have to sell at the same price.

To understand how replication works, let us consider a very simple model for stock prices where prices can jump to one of two points in each time period. This model, which is called a binomial model, allows us to model the replicating portfolio fairly easily. In the figure below, we have the binomial distribution of a stock, currently trading at $ 50 for the next two time periods. Note that in two time periods, this stock can be trading for as much as $ 100 or as little as $ 25. Assume that the objective is to value a call with a strike price of $ 50, which is expected to expire in two time periods:

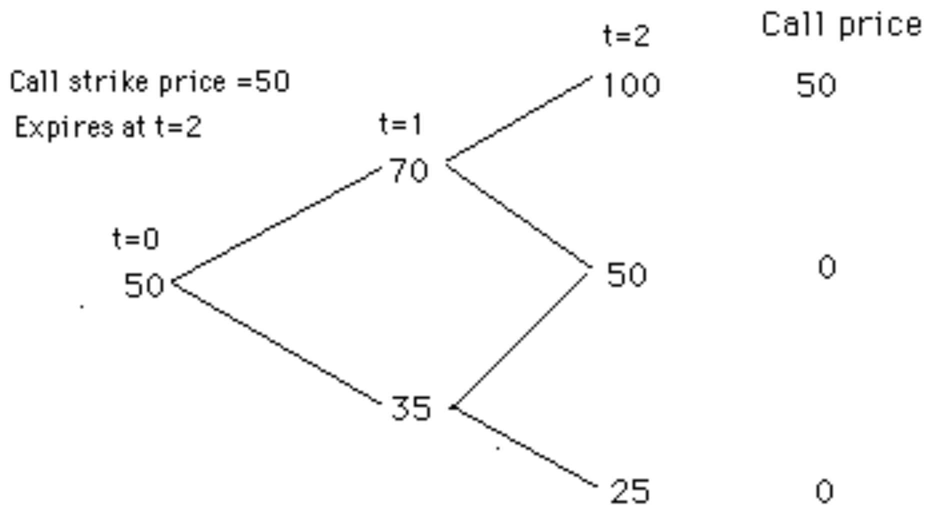

Now assume that the interest rate is 11%. In addition, define

Δ = Number of shares in the replicating portfolio

B = Dollars of borrowing in replicating portfolio

The objective is to combine Δ shares of stock and B dollars of borrowing to replicate the cash flows from the call with a strike price of 50. Since we know the cashflows on the option with certainty at expiration, it is best to start with the last period and work back through the binomial tree.

Step 1: Start with the end nodes and work backwards. Note that the call option expires at t=2, and the gross payoff on the option will be the difference between the stock price and the exercise price, if the stock price > exercise price, and zero, if the stock price < exercise price.

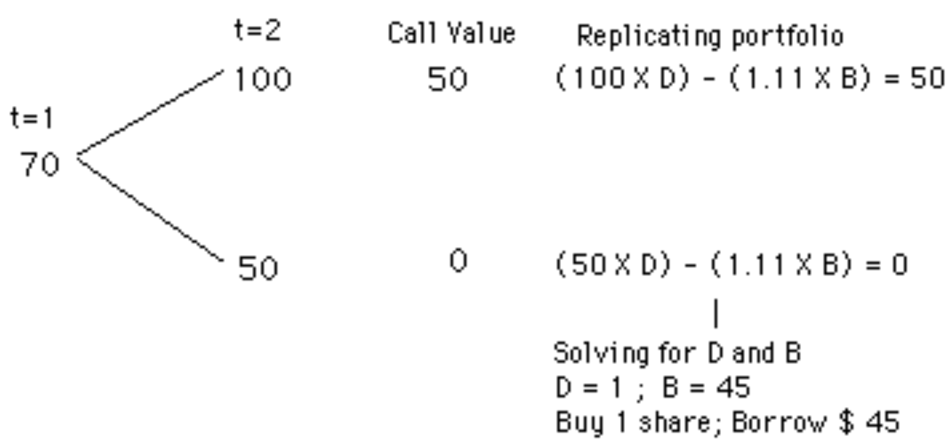

The objective is to construct a portfolio of Δ shares of stock and B in borrowing at t=1, when the stock price is \$ 70, that will have the same cashflows at t=2 as the call option with a strike price of 50. Consider what the portfolio will generate in cash flows under each of the two stock price scenarios, after you pay back the borrowing with interest (11% per period) and set the cash flows equal to the cash flows you would have received on the call.

If stock price = \$ 100: Portfolio Value = $100\,\Delta - 1.11\,B = 50$

If stock price = \$ 50: Portfolio Value = $50\,\Delta - 1.11\,B = 0$

We can solve for both the number of shares of stock you will need to buy (1) and the amount you will need to borrow (\$ 45) at t=1. Thus, if the stock price is \$70 at t=1, borrowing \$45 and buying one share of the stock will give the same cash flows as buying the call. To prevent arbitrage, the value of the call at t=1, if the stock price is \$70, has to be equal to the cost (to you as an investor) of setting up the replicating position:

Value of Call = Cost of Replicating Position = $70\Delta - B = (70)(1) - 45 = 25$

Considering the other leg of the binomial tree at t=1,

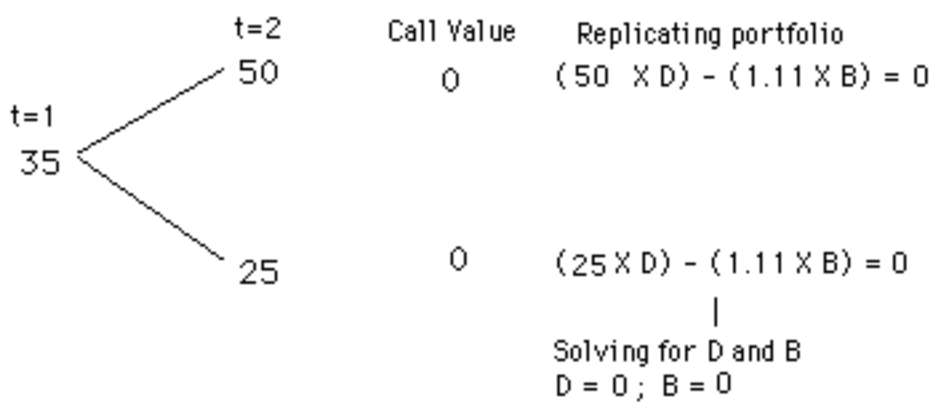

If the stock price is 35 at t=1, then the call is worth nothing.

Step 2: Now that we know how much the call will be worth at t=1 (\$25 if the stock price goes to \$ 70 and \$0 if it goes down to \$ 35), we can move backwards to the earlier time period and create a replicating portfolio that will provide the values that the option will provide.

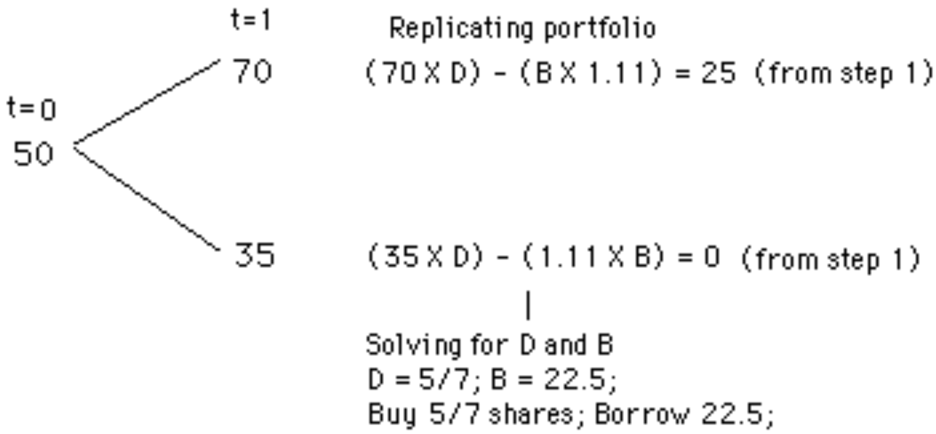

In other words, borrowing $22.5 and buying 5/7 of a share today will provide the same cash flows as a call with a strike price of $50. The value of the call therefore has to be the same as the cost of creating this position.

Value of Call = Cost of replicating position =

$$\left(\frac{5}{7}\right)(\text{Current Stock Price}) - 22.5 = \left(\frac{5}{7}\right)(50) - 22.5 = 13.21$$

Consider for the moment the possibilities for arbitrage if the call traded at less than $13.21, say $ 13.00. You would buy the call for $13.00 and sell the replicating portfolio for $13.21 and claim the difference of $0.21. Since the cashflows on the two positions are identical, you would be exposed to no risk and make a certain profit. If the call trade for more than $13.21, say $13.50, you would buy the replicating portfolio, sell the call and claim the $0.29 difference. Again, you would not have been exposed to any risk.

You could construct a similar example using puts. The replicating portfolio in that case would be created by selling short on the underlying stock and lending the money at the riskless rate. Again, if puts are priced at a value different from the replicating portfolio, you could capture the difference and be exposed to no risk.

What are the assumptions that underlie this arbitrage? The first is that both the traded asset and the option are traded and that you can trade simultaneously in both markets, thus locking in your profits. The second is that there are no (or at least very low transactions costs). If transactions costs are large, prices will have to move outside the band created by these costs for arbitrage to be feasible. The third is that you can borrow at the riskless rate and sell short, if necessary. If you cannot, arbitrage may no longer be feasible.

Arbitrage across options

When you have multiple options listed on the same asset, you may be able to take advantage of relative mispricing – how one option is priced relative to another - and lock in riskless profits. We will look first at the pricing of calls relative to puts and then consider how options with different exercise prices and maturities should be priced, relative to each other.

Put-Call Parity

When you have a put and a call option with the same exercise price and the same maturity, you can create a riskless position by selling the call, buying the put and buying the underlying asset at the same time. To see why, consider selling a call and buying a put with exercise price K and expiration date t, and simultaneously buying the underlying asset at the current price S. The payoff from this position is riskless and always yields a cashflow of K at expiration t. To see this, assume that the stock price at expiration is S*. The payoff on each of the positions in the portfolio can be written as follows:

Position	Payoffs at t if S*>K	Payoffs at t if S*<K
Sell call	-(S*-K)	0
Buy put	0	K-S*
Buy stock	S*	S*
Total	*K*	*K*

Since this position yields K with certainty, the cost of creating this position must be equal to the present value of K at the riskless rate ($K\ e^{-rt}$).

$$S+P-C = K\ e^{-rt}$$

$$C - P = S - K\ e^{-rt}$$

This relationship between put and call prices is called put call parity. If it is violated, you have arbitrage.

If $C-P > S - Ke^{-rt}$, you would sell the call, buy the put and buy the stock. You would earn more than the riskless rate on a riskless investment.

If $C-P < S - Ke^{-rt}$, you would buy the call, sell the put and sell short the stock. You would then invest the proceeds at the riskless rate and end up with a riskless profit at maturity.

Note that violations of put call parity create arbitrage opportunities only for options that can be exercised only at maturity (European options) and may not hold if options can be exercised early (American options).

Does put-call parity hold up in practice or are there arbitrage opportunities? One study[2] examined option pricing data from the Chicago Board of Options from 1977 to 1978 and found potential arbitrage opportunities in a few cases. However, the arbitrage opportunities were small and persisted only for short periods. Furthermore, the options examined were American options, where arbitrage may not be feasible even if put-call parity is violated. A more recent study by Kamara and Miller of options on the S&P 500 (which are European options) between 1986 and 1989 finds fewer violations of put-call parity and the deviations tend to be small, even when there are violations.

Mispricing across Strike Prices and Maturities

A spread is a combination of two or more options of the same type (call or put) on the same underlying asset. You can combine two options with the same maturity but different exercise prices (bull and bear spreads), two options with the same strike price but different maturities (calendar spreads), two options with different exercise prices and maturities (diagonal spreads) and more than two options (butterfly spreads). You may be able to use spreads to take advantage of relative mispricing of options on the same underlying stock.

> *Strike Prices*: A call with a lower strike price should never sell for less than a call with a higher strike price, assuming that they both have the same maturity. If it did, you could buy the lower strike price call and sell the higher strike price call, and lock in a riskless profit. Similarly, a put with a lower strike price should never sell for more than a put with a higher strike price and the same maturity. If it did, you could buy the higher strike price put, sell the lower strike price put and make an arbitrage profit.
>
> *Maturity*: A call (put) with a shorter time to expiration should never sell for more than a call (put) with the same strike price with a long time to expiration. If it did, you would buy the call (put) with the shorter maturity and sell the call (put) with the longer maturity (i.e, create a calendar spread) and lock in a profit today. When the first call expires, you will either exercise the second call (and have no cashflows) or sell it (and make a further profit).

Even a casual perusal of the option prices listed in the newspaper each day should make it clear that it is very unlikely that pricing violations that are this egregious will exist in a liquid options market.

---

[2] See Klemkosky, R.C. and B.G. Resnick, 1979, Put-Call Parity and Market Efficiency, Journal of Finance, v 34, pg 1141-1155.

Fixed Income Arbitrage

Fixed income securities lend themselves to arbitrage more easily than equity because they have finite lives and fixed cash flows. This is especially so, when you have default free bonds, where the fixed cash flows are also guaranteed. Consider one very simple example. You could replicate a 10-year treasury bond's cash flows by buying zero-coupon treasuries with expirations matching those of the coupon payment dates on the treasury bond. For instance, if you invest $ 100 million in a ten-year treasury bond with an 8% coupon rate, you can expect to get cashflows of $ 4 million every 6 months for the next 10 years and $ 100 million at the end of the tenth year. You could have obtained exactly the same cashflows by buying zero-coupon treasuries with face values of $ 4 million, expiring every 6 months for the next ten years, and an additional 10-year zero coupon bond with a face value of $ 100 million. Since the cashflows are identical, you would expect the two positions to trade for the same price. If they do not trade at the same price, you would buy the cheaper position and sell the more expensive one, locking in the profit today and having no cashflow or risk exposure in the future.

With corporate bonds, you have the extra component of default risk. Since no two firms are exactly identical when it comes to default risk, you may be exposed to some risk if you are using corporate bonds issued by different entities. In fact, two bonds issued by the same entity may not be equivalent because of differences in how they are secured and structured. There are some arbitrageurs who argue that bond ratings are a good proxy for default risk, and that buying one AA rated bond and selling another should be riskless, but bond ratings are not perfect proxies for default risk. In fact, you see arbitrage attempted on a wide variety of securities with promised cashflows, such as mortgage backed bonds. While you can hedge away much of the cashflow risk, the nature of the cashflow claims will still leave you exposed to some risk. With mortgage backed bonds, for instance, the unpredictability of prepayments by homeowners has exposed many "riskless" positions to risk.

Is there any evidence that investors are able to find treasuries mispriced enough to generate arbitrage profits? Grinblatt and Longstaff, in an assessment of the treasury strips program – a program allowing investors to break up a treasury bond and sell its individual cash flows – note that there are potential arbitrage opportunities in these markets but find little evidence of trading driven by these opportunities. A study by Balbas and Lopez of the Spanish bond market may shed some light on this question. Examining default free and option free bonds in the Spanish market between 1994 and 1998, they conclude that there were arbitrage opportunities especially surrounding innovations in financial markets. We would extend their findings to argue that opportunities for arbitrage with fixed income securities are probably greatest when new types of bonds are introduced – mortgage backed securities in the early 1980s, inflation- indexed treasuries in the late 1990s and the treasury strips program in the late 1980s. As investors become more informed about these bonds and how they should be priced, arbitrage opportunities seem to subside.

Determinants of Success

The nature of pure arbitrage – two identical assets that are priced differently – makes it likely that it will be short lived. In other words, in a market where investors are on the look out for riskless profits, it is very likely that small pricing differences will be exploited quickly, and in the process, disappear. Consequently, the first two requirements for success at pure arbitrage are access to real-time prices and instantaneous execution. It is also very likely that the pricing differences in pure arbitrage will be very small – often a few hundredths of a percent. To make pure arbitrage feasible, therefore, you can add two more conditions. The first is access to substantial debt at favorable interest rates, since it can magnify the small pricing differences. Note that many of the arbitrage positions require you to be able to borrow at the riskless rate. The second is economies of scale, with transactions amounting to millions of dollars rather than thousands. Institutions that are successful at pure arbitrage often are able to borrow several times their equity at the riskless rate to fund arbitrage transactions, using the guaranteed profits on the transaction as collateral.

With these requirements, it is not surprising that individual investors have generally not been able to succeed at pure arbitrage. Even among institutions, pure arbitrage is feasible only to a few, and even to those, it is a transient source of profits in two senses. First, you cannot count on the existence of pure arbitrage opportunities in the future, since it requires that markets repeat their errors over time. Second, the very fact that some institutions make profits from arbitrage attracts other institutions into the market, reducing the likelihood of future arbitrage profits. To succeed in the long term with arbitrage, you will need to be constantly on the lookout for new arbitrage opportunities.

Near Arbitrage

In near arbitrage, you either have two assets that are very similar but not identical, which are priced differently, or identical assets that are mispriced, but with no guaranteed price convergence. No matter how sophisticated your trading strategies may be in these scenarios, your positions will no longer be riskless.

Same Security, Multiple Markets

In today's global markets, there are a number of stocks that are listed on more than one market. If you can buy the same stock at one price in one market and simultaneously sell it at a higher price in another market, you can lock in a riskless profit. As we will see in this section, things are seldom this simple.

Dual and Multiple Listings

Many large companies such as Royal Dutch, General Electric and Microsoft trade on multiple markets on different continents. Since there are time periods during the day when there is trading occurring on more than one market on the same stock, it is conceivable (though not likely) that you could buy the stock for one price in one market and sell the same stock at the same time for a different (and higher price) in another market. The stock will trade in different currencies, and for this to be a riskless transaction, the trades have to at precisely the same time and you have to eliminate any exchange rate risk by converting the foreign currency proceeds into the domestic currency instantaneously. Your trade profits will also have to cover the different bid-ask spreads in the two markets and transactions costs in each.

There are some exceptional cases, where the same stock trades in different markets in one country. Swaicki and Hric examine 84 Czech stocks that trade on the two Czech exchanges – the Prague Stock Exchange (PSE) and the Registration Places System (RMS)- and find that prices adjust slowly across the two markets, and that arbitrage opportunities exist (at least on paper) –the prices in the two markets differ by about 2%. These arbitrage opportunities seem to increase for less liquid stocks. While the authors consider transactions cost, they do not consider the price impact that trading itself would have on these stocks and whether the arbitrage profits would survive the trading.

Depository Receipts

Many Latin American and European companies have American Depository Receipts (ADRs) listed on the U.S. market. These depository receipts create a claim equivalent to the one you would have had if you had bought shares in the local market and should therefore trade at a price consistent with the local shares. What makes them different and potentially riskier than the stocks with dual listings is that ADRs are not always directly comparable to the common shares traded locally – one ADR on Telmex, the Mexican telecommunications company, is convertible into 20 Telmex shares. In addition, converting an ADR into local shares can be both costly and time consuming. In some cases, there can be differences in voting rights as well. In spite of these constraints, you would expect the price of an ADR to closely track the price of the

Most widely traded ADRs: Take a look at the 50 most widely traded ADRs on the U.S. market.

shares in the local market, albeit with a currency overlay, since ADRs are denominated in dollars.

In a study conducted in 2000 that looks at the link between ADRs and local shares, Kin, Szakmary and Mathur conclude that about 60 to 70% of the variation in ADR prices can be attributed to movements in the underlying share prices and that ADRs overreact to the U.S, market and under react to exchange rates and the underlying stock. However, they also conclude that investors cannot take advantage of the pricing errors in ADRs because convergence does not occur quickly or in predictable ways. With a longer time horizon and/or the capacity to convert ADRs into local shares, though, you should be able to take advantage of significant pricing differences.

Closed End Funds

In a conventional mutual fund, the number of shares increases and decreases as money comes in and leaves the fund, and each share is priced at net asset value – the market value of the securities of the fund divided by the number of shares. Closed end mutual funds differ from other mutual funds in one very important respect. They have a fixed number of shares that trade in the market like other publicly traded companies, and the market price can be different from the net asset value.

In both the United States and the United Kingdom, closed end mutual funds have shared a very strange characteristic. When they are created, the price is usually set at a premium on the net asset value per share. As closed end funds trade, though, the market price tends to drop below the net asset value and stay there. Figure 11.7 provides the distribution of price to net asset value for all closed end funds in the United States in early 2002.

Figure 11.7: Discounts/Premiums on Closed End Funds- June 2002

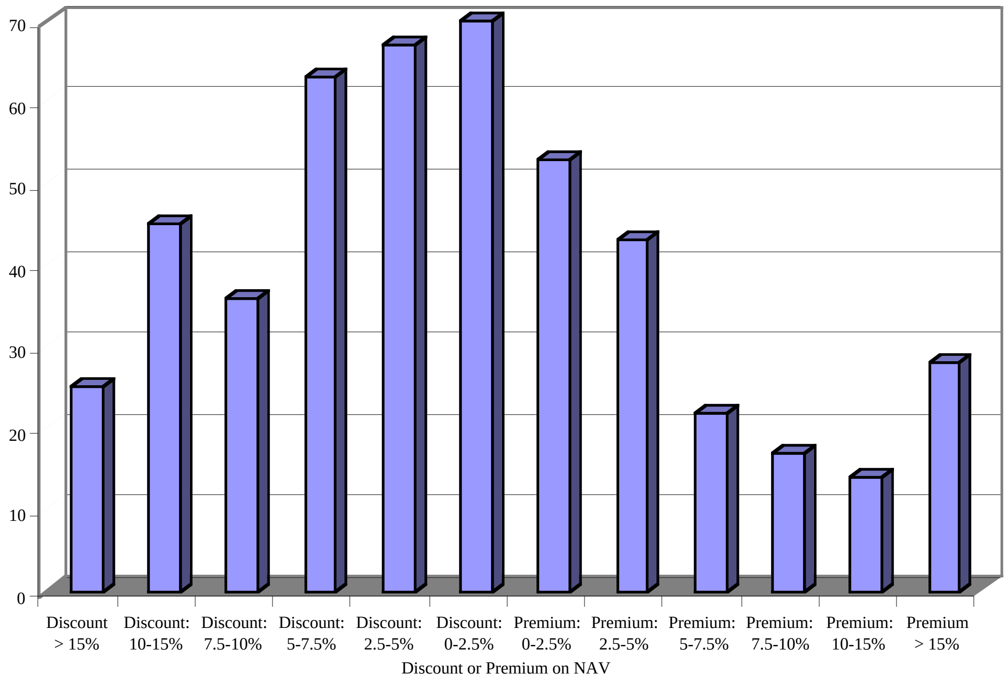

Note that almost 70% of the closed end funds trade at a discount to net asset value and that the median discount is about 5%.

So what, you might ask? Lots of firms trade at less than the estimated market value of their assets. That might be true, but closed end funds are unique for two reasons. First, the assets are all traded stocks and the market value is therefore known at any point in time and not an estimate. Second, liquidating a closed end fund's assets should not be difficult to do, since you are selling securities to the market. Thus, liquidation should neither be costly nor time consuming. Given these two conditions, you may wonder why you should not buy closed end funds that trade at a discount and either liquidate them yourself or hope that some one else will liquidate them. Alternatively, you may be able to push a closed-end fund to open-end and see prices converge on net asset value. Figure 11.8 reports on the performance of closed-end funds when they open end, based upon a study of 94 UK closed-end funds that open ended:

Figure 11.8: Relative Discount on Closed End Funds that Open End

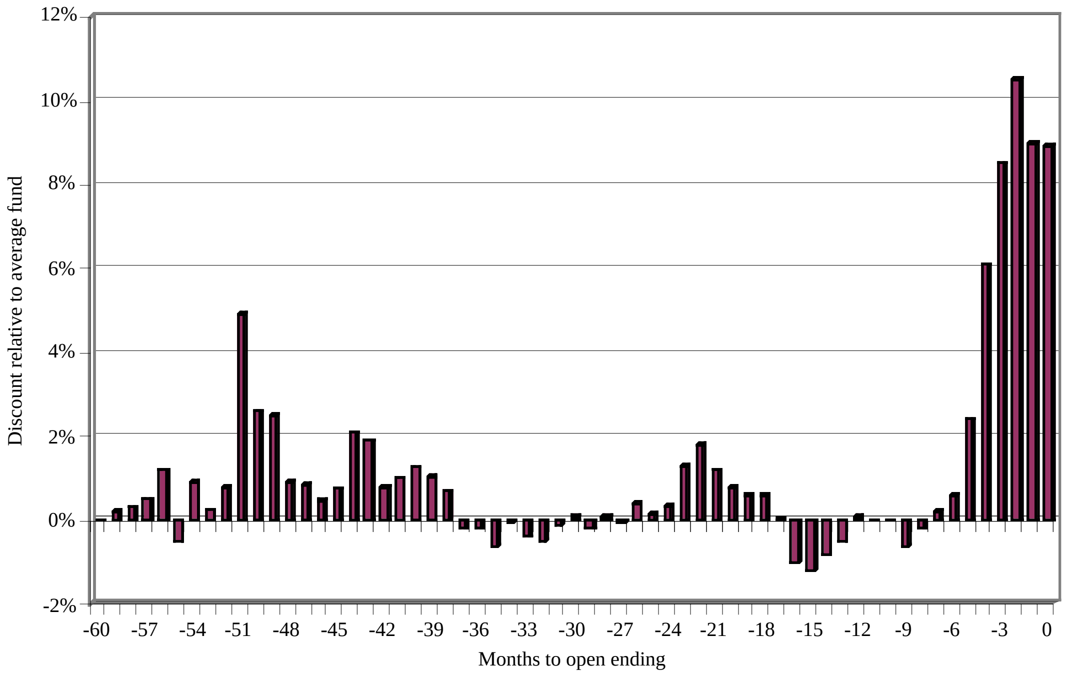

Note that as you get closer to the open-ending date (day 0), the discount becomes smaller relative to the average closed-end fund. For instance, the discount goes from being on par with the discount on other funds to being about 10% lower than the typical closed-end fund.

So what is the catch? In practice, taking over a closed-end fund while paying less than net asset value for its shares seems to be very difficult to do for several reasons- some related to corporate governance and some related to market liquidity. The potential profit is also narrowed by the mispricing of illiquid assets in closed end fund portfolios (leading to an overstatement of the NAV) and tax liabilities from liquidating securities. There have been a few cases of closed end funds being liquidated, but they remain the exception. What about the strategy of buying discounted funds and hoping that the discount disappears? This strategy is clearly not riskless but it does offer some promise. In one of the first studies of this strategy in 1978, Thompson studied closed end funds from 1940 to 1975 and reported that you could earn an annualized excess return of 4% from buying discounted funds. A study in 1986 by Anderson reports excess returns from a strategy of buying closed end funds whose discounts had widened and selling funds whose discounts had narrowed – a contrarian strategy applied to closed end

Most discounted closed end funds: Take a look at the 50 closed end funds with the largest discounts.

funds. In 1995, Pontiff reported that closed end funds with a discount of 20% or higher earn about 6% more than other closed end funds. This, as well as studies in the UK, seem to indicate a strong mean reversion component to discounts at closed funds. Figure 11.9, which is from a study of the discounts on closed end funds in the UK, tracks relative discounts on the most discounted and least discounted funds over time:

Figure 11.9: Discounts on most discounted and least discounted funds over time

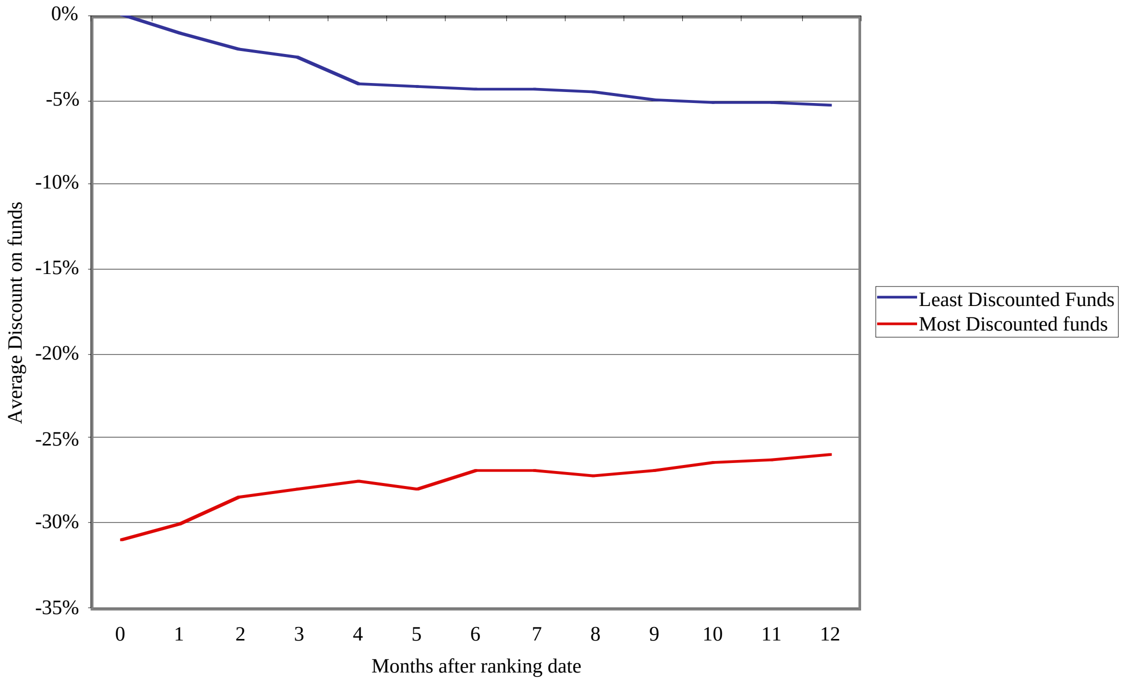

Source: Minio-Paluello (1998)

Note that the discounts on the most discounted funds decrease whereas the discounts on the least discounted funds increase, and the difference narrows over time.

Convertible Arbitrage

A convertible bond has two securities embedded in it – a conventional bond and a conversion option on the company's stock. When companies have convertible bonds or convertible preferred stock outstanding in conjunction with common stock, warrants, preferred stock and conventional bonds, it is entirely possible that you could find one of these securities mispriced relative to the other, and be able to construct a near-riskless strategy by combining two or more of the securities in a portfolio.

In the simplest form of this strategy, note that since the conversion option is a call option on the stock, you could construct a conversion option by combining the underlying stock and the treasury bond (a replicating portfolio). Adding a conventional bond to this should create the equivalent of the convertible bond. Once you can do this, you can take advantage of differences between the pricing of the convertible bond and synthetic convertible bond and potentially make arbitrage profits. In the more complex forms, when you have warrants, convertible preferred and other options trading simultaneously on a firm, you could look for options that are mispriced relative to each other, and then buy the cheaper option and sell the more expensive one.

In practice, there are several possible impediments. First, many firms that issue convertible bonds do not have straight bonds outstanding, and you have to substitute in a straight bond issued by a company with similar default risk. Second, companies can force conversion of convertible bonds, which can wreak havoc on arbitrage positions. Third, convertible bonds have long maturities. Thus, there may be no convergence for long periods, and you have to be able to maintain the arbitrage position over these periods. Fourth, transactions costs and execution problems (associated with trading the different securities) may prevent arbitrage.

Determinants of Success

Studies that have looked at closed end funds, dual listed stocks and convertibles all seem to conclude that there are pockets of inefficiency that can exploited to make money. However, there is residual risk in all off these strategies, arising sometimes from the fact that the assets are not perfectly identical (convertibles versus synthetic convertibles) or because there are no mechanisms for forcing the prices to converge (closed end funds).

So, what would you need to succeed with near arbitrage strategies? The first thing to note is that these strategies will not work for small investors or for very large investors. Small investors will be stymied both by transactions costs and execution problems. Very large investors will quickly drive discounts to parity and eliminate excess returns. If you decide to adopt these strategies, you need to refine and focus your strategies on those opportunities where convergence is most likely. For instance, if you decide to try to exploit the discounts of closed-end funds, you should focus on the closed end funds that are most discounted and concentrate especially on funds where there is the potential to bring pressure on management to open end the funds. You should also avoid funds with substantial illiquid or non-traded stocks in their portfolios, since the net asset values of these funds may be significantly overstated.

The Limits of Arbitrage

In a perfect world (at least for financial economists), any relative mispricing of assets attracts thousands of investors who borrow risklessly and take advantage of the arbitrage. In the process, they drive it out of existence. In the real world, it is much more likely that any assets that are mispriced are not perfectly identical (thus introducing some risk into the mix) and that only a few large investors have the capacity to access low-cost debt and take advantage of arbitrage opportunities. It is entirely possible then that near arbitrage opportunities will be left unexploited because these large investors are unwilling to risk their capital in these investments. Vishny and Shleifer provide a fascinating twist on this argument. They note that the more mispriced assets become on a relative basis, the greater the risk to arbitrageurs that the mispricing will move against them. Hence, they argue that arbitrageurs will pull back from investing in the most mispriced assets, especially if there are thousands of other traders in the market who are pushing prices in the opposite direction.

Speculative Arbitrage

The word arbitrage is used much too loosely in investments and there are a large number of strategies that are characterized as arbitrage, but actually expose investors to significant risk. In fact, the strategies covered in this section would probably be better characterized as pseudo arbitrage strategies.

Paired Arbitrage

In classic arbitrage, you buy an asset at one price and sell an exactly identical asset at a different (and higher) price. In paired arbitrage, you buy one stock (say GM) and sell another stock that you view as very similar (say Ford), and argue that you are not that exposed to risk. Clearly, this strategy is not riskless since no two equities are exactly identical, and even if they were very similar, there may be no convergence in prices.

Let us consider first how you pair up stocks. The conventional practice among those who have used this strategy on Wall Street has been to look for two stocks whose prices have historically moved together – i.e., have high correlation over time. This often leads to two stocks in the same sector, such as GM and Ford. Once you have paired the stocks, you compute the spread between them and compare this spread to historic norms. If the spread is too wide, you buy the cheaper stock and short the more expensive stock. In many cases, the strategy is self-financing. For example, if Ford is trading at $ 20 and GM is trading at $ 40 and you believe that GM is overpriced relative to Ford, you would buy two shares of Ford and sell short one share of GM. If you are right, and the spread narrows between the shares, you will profit on your paired position.

Can such a simplistic strategy, based entirely upon past prices, make excess returns? In 1999, Gatev, Goetzmann and Rouwenhorst tested a variety of trading rules based upon pairs trading from 1982-1997, using the following process:

- Screening first for only stocks that traded every day, the authors found a matching partner for each stock by looking for the stock with the minimum squared deviation in normalized price series[3]. Intuitively, note that if two stocks move together all the time, the squared distance in returns should be zero. Once they had paired all the stocks, they studied the pairs with the smallest squared deviation separating them.
- With each pair, they tracked the normalized prices of each stock and took a position on the pair, if the difference exceeded the historical range by two standard deviations, buying the cheaper stock and selling the more expensive one.

Over the 15 year period, the pairs trading strategy did significantly better than a buy-and-hold strategy. Strategies of investing in the top 20 pairs earned an excess return of about 6% over a 6-month period, and while the returns drop off for the pairs below the top 20, you continue to earn excess returns. When the pairs are constructed by industry group (rather than just based upon historical prices), the excess returns persist but they are smaller. Controlling for the bid-ask spread in the strategy reduces the excess returns by about a fifth, but the returns are still significant.

While the overall trading strategy looks promising, there are two points worth emphasizing that should also act as cautionary notes about this strategy. The first is that the study quoted above found that the pairs trading strategy created negative returns in about one out of every six periods, and that the difference between pairs often widened before it narrowed. In other words, it is a risky investment strategy that also requires the capacity to trade instantaneously and at low cost. The second is a quote from a well known quantitative analyst, David Shaw, who bemoaned the fact that by the late 1990s, the pickings for quantitative strategies (like pairs trading) had become slim because so many investment banks were adopting the strategies. As the novelty has worn off, it seems unlikely that the pairs trading will generate the kinds of profits it generated during the 1980s.

Merger Arbitrage

As we noted in the last chapter, the stock price of a target company jumps on the announcement of a takeover. However, it trades at a discount usually to the price offered by the acquiring company. The difference between the post-announcement price and the offer

---

[3] If you use absolute prices, a stock with a higher price will always look more volatile. You can normalize the prices around 1 and use these series.

price is called the arbitrage spread, and there are investors who try to profit off this spread in a strategy called merger or risk arbitrage. If the merger succeeds, the arbitrageur captures the arbitrage spreads, but if it fails, he or she could make a substantial loss. In a more sophisticated variant in stock mergers (where shares of the acquiring company are exchanged for shares in the target company), the arbitrageur will sell the acquiring firm's stock in addition to buying the target firm's stock.

To begin with, we should note that the term risk arbitrage is extremely misleading, It is clearly not arbitrage in the classic sense since there are no guaranteed profits and it is not quite clear why the prefix "risk" is attached to it. Notwithstanding this quarrel with terminology, we can examine whether risk arbitrage delivers the kinds of returns we often hear about anecdotally, and if it does, is it compensation for risk (that the merger may not go through) or is it an excess return? Mitchell and Pulvino (2000) use a sample of 4750 mergers and acquisitions to examine this question. They conclude that there are excess returns associated with buying target companies after acquisition announcements of about 9.25% annually, but that you lost about two thirds of these excess returns if you factor in transactions costs and the price impact that you have when you trade (especially on the less liquid companies).

While the overall strategy returns look attractive, Mitchell and Pulvino also point to one unappealing aspect of this strategy. The strategy earns moderate positive returns much of the time, but earns large negative returns when it fails. Thus, they argue that this strategy has payoffs that resemble those you would observe if you sell puts – when the market goes up, you keep the put premium but when it goes down, you lost much more. Does this make it a bad strategy? Not at all, but it points to the dangers of risk arbitrage when it is restricted to a few big-name takeover stocks (as it often is)- an investor who adopts this strategy is generally is one big failure away from going under. If you use leverage to do risk arbitrage, the dangers are multiplied.

Determinants of Success

The fact that we categorize the strategies in this section as speculative arbitrage is not meant to be a negative comment on the strategies. We believe that these are promising investment strategies that have a history of delivering excess returns but they are not riskfree. More ominously, it is easy for those who have done pure arbitrage to drift into near arbitrage and then into speculative arbitrage as they have funds to invest. In some cases, their success at pure or near arbitrage may bring in funds which require this shift. In doing so, however, there are two caveats that have to be kept in mind:

- The use of financial leverage has to be scaled to reflect the riskiness of the strategy. With pure arbitrage, you can borrow 100% of what you need to put the strategy into play. In futures arbitrage, for instance, you borrow 100% of the spot price and borrow the commodity. Since there is no risk, the leverage does not create any damage. As you move to near and speculative arbitrage, this leverage has to be reduced. How much it has to be reduced will depend upon both the degree of risk in the strategy and the speed with which you think prices will converge. The more risky a strategy and the less certain you are about convergence, the less debt you should take on.
- These strategies work best if you can operate without a market impact. As you get more funds to invest and your strategy becomes more visible to others, you run the risk of driving out the very mispricing that attracted you to the market in the first place.

In many ways, the rise and fall of Long Term Capital Management (see below) should stand as testimony to how even the brightest minds in investing can sometimes either miss or willfully ignore these realities. Long Term Capital Management's eventual undoing can be traced to many causes but the immediate cause was the number of speculative arbitrage positions they put in place – pairs trading, interest rate bets – with tremendous leverage.

The Fall of Long Term Capital

Investors considering arbitrage as their preferred investment philosophy should pay heed to the experiences of Long Term Capital Management (LTCM). The firm, which was founded in the early 1990s by ex-Salomon trader, John Merriweather, promised to bring together the best minds in finance to find and take advantage of arbitrage opportunities around the world. Delivering on the first part of the promise, Merriweather lured the best bond traders from Salomon and brought on board two Nobel prize winners – Myron Scholes and Bob Merton. In the first few years of its existence, the firm also lived up to the second part of the promise, earning extraordinary returns for the elite of Wall Street. In those years, LTCM was the envy of the rest of the street as it used low cost debt to lever up its capital and invest in pure and near arbitrage opportunities.

As the funds at their disposal got larger, the firm had to widen its search to include pseudo arbitrage investments. By itself, this would not have been fatal but the firm continued to use the same leverage on these riskier investments as it did on its safe investments. It bet on paired trades in Europe and decreasing spreads in country bond markets, arguing that the sheer number of investments in had in its portfolio would create diversification – if it lost on one investment, it would gain on another. In 1997, the strategy unraveled as collapses in one market (Russia) spread into other markets as well. As the portfolio dropped in value, LTCM found itself facing the downside of its size and high leverage. Unable to unwind its large positions without affecting market prices and facing the pressures of lenders, LTCM faced certain bankruptcy. Fearing that it would bring down other investors in the market, the Federal Reserve engineered a bailout of the firm.

What are the lessons that we can learn from the fiasco? Besides the cynical one that it is good to have friends in high places, you could argue that the fall of LTCM teaches us that

(a) Size can be a double-edged sword. While it gives you economies of scale in transactions costs and lowers the cost of funding, it also makes it more difficult for you to unwind positions that you have taken.

(b) Leverage can make low-risk positions into high-risk investments, since small moves in the price can translate into large changes in equity

(c) The most brilliant minds in the world and the best analytical tools cannot insulate you from the vagaries of the market.

For an excellent analysts of LTCM, see "When Genius Failed" by Roger Lowenstein

Long Short Strategies – Hedge Funds

In the last few years, hedge funds have become one of the fastest growing parts of the money management business. Largely unregulated, headed by outsized personalities like George Soros and Julian Robertson and seemingly delivering huge returns to their investors, hedge funds have become serious players in the money management game. At the outset of this section, we have to note that it is probably not quite accurate to categorize hedge funds as having any specific strategy, since you can have hedge funds specializing in almost every strategy we have listed in this book. You can have value and growth investing hedge funds, hedge funds that specialize in market timing, hedge funds that invest on information and hedge funds that do convertible arbitrage. The reason that we consider it in this chapter is because it lends itself particularly well to arbitrage strategies, which require that you buy some assets and sell short on others at the same time. In this section, we will take a closer look at hedge fund strategies and how well they really have performed.

Background and History

What makes a fund into a hedge fund? The common characteristic shared by all hedge funds is that they not only buy assets that they feel are undervalued, but that they simultaneously sell short on assets that they believe to be overvalued. Defined this way, hedge funds have probably been around as long as stock markets have been in existence, though they have traditionally been accessible only to the very wealthy. In the last decade, however, hedge funds have taken a larger and larger market share of total investment funds. While the magnitude of the funds under hedge fund management is disputed, it is estimated that 500 to 600 billion dollars in assets were managed by about 6000 hedge funds in early 2002.

Performance

Are the storied returns to investing in hedge funds true? Are small investors who are often shut out from investing in hedge funds losing out because of this? To answer these questions, we need to look not at anecdotal evidence or the performance of the best hedge funds, but at all hedge funds. In a study of all offshore hedge funds from 1989 to 1995, Brown, Goetzmann and Ibbotson (1999) chronicled the returns in table 11.2.

Table 11.2: Offshore Hedge Fund Returns – 1989 to 1995

Year	*No of funds in sample*	*Arithmetic Average Return*	*Median Return*	*Return on S&P 500*	*Average Annual Fee (as % of money under management)*	*Average Incentive Fee (as % of excess returns)*
1988-89	78	18.08%	20.30%		1.74%	19.76%
1989-90	108	4.36%	3.80%		1.65%	19.52%
1990-91	142	17.13%	15.90%		1.79%	19.55%
1991-92	176	11.98%	10.70%		1.81%	19.34%
1992-93	265	24.59%	22.15%		1.62%	19.10%
1993-94	313	-1.60%	-2.00%		1.64%	18.75%
1994-95	399	18.32%	14.70%		1.55%	18.50%
Entire Period		13.26%		16.47%%		

[a]Returns are net of fees

There are several interesting numbers in this table. First, the average hedge fund earned a lower return (13.26%) over the period than the S&P 500 (16.47%), but it also had a lower standard deviation in returns (9.07%) than the S & P 500 (16.32%). Thus, it seems to offer a better payoff to risk, if you divide the average return by the standard deviation – this is the commonly used Sharpe ratio for evaluating money managers. Second, these funds are much more expensive than traditional mutual funds, with much higher annual fess and annual incentive fees that take away one out of every five dollars of excess returns.

As noted earlier, hedge funds come in all flavors. A study in 1999 by Ackermann, McEnally and Ravenscraft looked at the returns on a variety of hedge funds, and their findings are summarized in figure 11.10 below:

Figure 11.10: Hedge Funds: Average Returns and Standard Deviations - 1989-1995

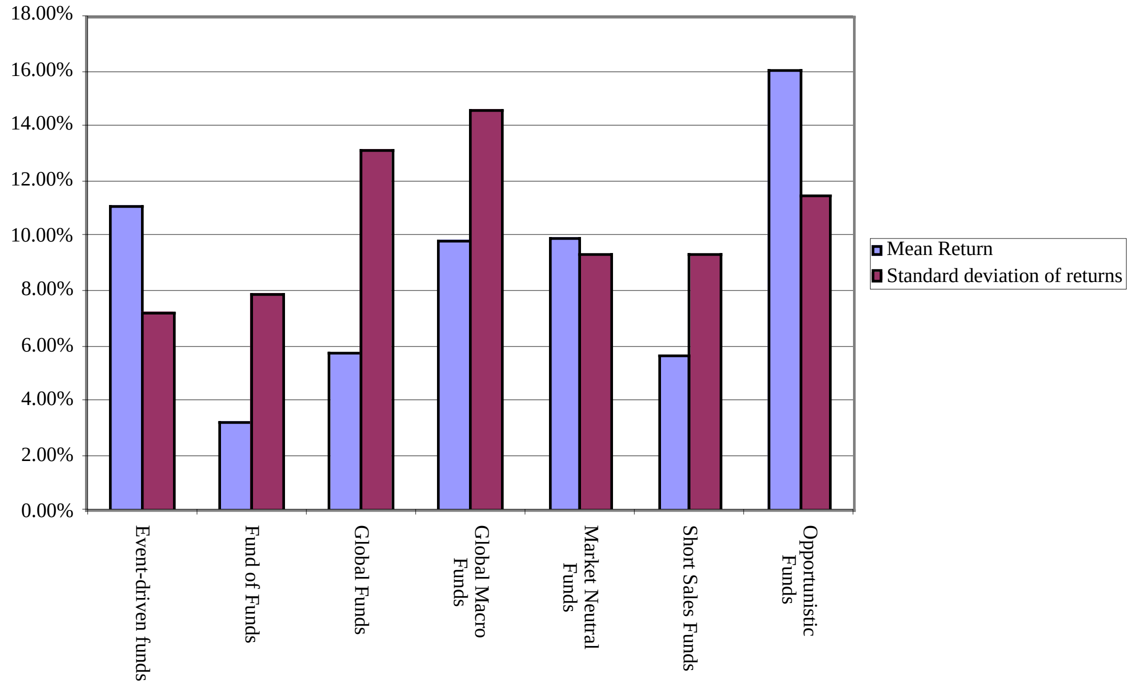

Source: Ackermann, McEnally and Ravenscraft

Note that event-driven and opportunistic funds (which look for arbitrage opportunities) provide the best Sharpe ratios and fund of funds (which invest across what they claim are the best mutual funds) offer the worst.

Liang examined 2016 hedge funds from 1990 to 1999. While his overall conclusions matched those of Brown et al., i.e. that these hedge funds earned a lower return than the S&P 500 (14.2% versus 18.8%), they were less risky and had higher Sharpe ratios (0.41 for the hedge funds versus 0.27 for the S&P 500), he also noted that there a large number of hedge funds die each year. Of the 2016 funds over the period for instance, only 1407 remained live at the end of the period.

In summary, what are we to make of these findings? First, the biggest advantage of hedge funds does not seem to lie in high returns but in lower risk. Every study that we have quoted finds that hedge funds under perform the market, especially over bull markets, but that they are more efficient in their risk taking.[4] Second, the high failure rate of hedge funds has to be factored into any investment strategy. An investor considering hedge funds should either hold several hedge funds or view a hedge fund as a supplemental investment. Third, the fee structure is tilted towards management and the expenses are larger for hedge funds. Since many hedge fund strategies are successful only on the small scale, it will be interesting to see what happens to returns as the business grows. In our view, success at attracting more money may ultimately prove fatal to this business.

Conclusion

For many practitioners, the promise of being able to invest no money, take no risk and still make a profit remains alluring. That is essentially what arbitrage allows you to do. In pure arbitrage, you buy a security at one price and sell an exactly identical security at a higher price, thus ensuring yourself a riskless profit. The markets where pure arbitrage is feasible tend to be the derivatives markets, since you can construct equivalent securities using the underlying asset and lending or borrowing. In futures markets, for instance, you can attain equivalent results by borrowing money, buying the underlying asset (storable commodity, stock index or bond) and storing it until the maturity of a futures contract, or buying the futures contract directly. In the options markets, you can replicate a call option by borrowing money and buying the underlying asset and a put option, by selling short on the underlying asset and lending at the riskless rate. Pure arbitrage may also be feasible with default free bonds. If opportunities exist for pure arbitrage, they are likely to be few and far between and available only to a subset of large institutional investors with very low transactions costs and the capacity to take on very high leverage at close to riskless rates.

Near arbitrage refers to scenarios where the two assets being bought and sold are either not exactly identical or where there is no point in time where prices have to converge. We considered three examples – equities on the same company that trade in different markets at different prices, closed end funds that trade at a discount on the net asset value of the securities in the fund and convertible bonds that trade at prices that are inconsistent with the prices of other securities –warrants and convertible preferred - issued by the company. In each, the mispricing may be obvious but the profits are not guaranteed because you cannot force convergence (liquidating the closed end fund or converting the ADR into local

---

[4] A contrary viewpoint is offered by three hedge fund managers at AQR Capital Management in Chicago. They argue that the returns on many hedge funds are based upon self-assessments of value for illiquid securities. The resulting smoothing out of returns creates the illusion of low or no correlation with the market and low standard deviations.

shares) or because you cannot create exactly identical securities (convertible arbitrage). You may still be able to construct low risk strategies that earn high returns, but riskless profits are not feasible.

In the final section, we examine what we term speculative or pseudo arbitrage. In pairs trading, two stocks that tend to move together are paired up and traded when the price difference moves out of a historical range – you sell the more expensive stock and you buy the cheaper one. In merger arbitrage, you buy stocks of target firms after acquisitions are announced and hope to make the difference between the price post-announcement and the offer price. While both these strategies offer promising returns, neither is close to being riskless.

In closing, we look at the part of the money management business that is closest in philosophy to arbitrage– the hedge fund business. While they share the common characteristic of a long-short strategy, hedge funds come in all varieties. On average, they have returns that are lower than a strategy of buying and holding stocks, but have risk that is significantly lower. As hedge funds use this finding to bring in more money, they run the risk of being the victims of their own success, since it is not clear whether many hedge fund strategies will scale up – i.e. a convertible arbitrage strategy that works with $ 100 million may not work with a billion.

Lessons for investors

To be a successful arbitrageur, you need to:

- Understand that near arbitrage is more likely than pure arbitrage: Pure arbitrage opportunities, if they exist, are most likely to be found in derivatives and government bond markets and will be very quickly exploited. Near arbitrage, where you have two almost-identical assets that are priced differently or where you have no forced convergence, should be more common.
- Have excellent execution capabilities and low execution costs: Arbitrage requires you to trade large quantities instantaneously in two or more markets.
- Have access to low cost debt: The pricing differences between two similar assets, even if they exist, are likely to be very small and can be made into substantial returns only by using leverage.

If you decide to move to pseudo arbitrage, you need to

- Keep leverage under control: As you move from pure to near arbitrage and from near to pseudo arbitrage, the risk in your strategy will increase and you should reduce the financial leverage in your strategy accordingly.
- Recognize that size is a double edged sword: As you get more funds to invest, you may be able to reduce your execution costs, but you will also have a much more difficult time getting in and out of your positions quickly and without a price impact.

CHAPTER 12

THE IMPOSSIBLE DREAM? TIMING THE MARKET

It is every investor's dream to time the market and this occurs for obvious reasons. A successful market timer does not have to any skill at picking stocks since market timing alone will deliver extraordinary returns. In fact, we will begin this chapter by looking at the immense payoff that can come from timing the market well. This payoff to timing the market makes all of us easy victims for the next market-timing strategy. In this chapter, we consider a range of market timing strategies ranging from technical indicators to fundamental indicators to societal indicators. We look at the assumptions underlying each indicator and why they sometimes help us predict market movements.

In the final section of this chapter, we examine why market timing is so difficult to succeed at, relative to other investment philosophies. We also consider what allows market timing to succeed sometimes and whether we can replicate their success.

Market Timing: Payoff and Costs

The question of whether market timing has a big payoff and what its costs are arouses strong views from both practitioners and academics. While academics are fairly unified in their belief that market timing is not worth the time and resources that are expended on it, practitioners feel deeply on both sides of the issue. We will begin by looking at the payoffs to market timing and then consider the costs.

The Payoff to Market Timing

In a 1986 article, a group of researchers[1] raised the shackles of many an active portfolio manager by estimating that as much as 93.6% of the variation in quarterly performance at professionally managed portfolios could be explained by the mix of stocks, bonds and cash at these portfolios.[2] In a different study in 1992, Shilling examined the effect on your annual returns of being able to stay out of the market during bad months. He concluded that an investor who would have missed the 50 weakest months of the market between 1946 and 1991 would have seen his annual returns almost double from 11.2% to 19%. Ibbotson examined the relative importance of asset allocation and security selection of 94 balanced mutual funds and 58 pension funds, all of which had to make both asset allocation and security selection decisions. Using ten years of data through 1998, Ibbotson

---

[1] See "Determinants of Portfolio Performance" by Brinson, Hood and Beebower.

[2] This is a much quoted and misquoted study. A survey by Nutall and Nutall found that of 50 writers who quoted this study, 37 misread it to indicate that 93% of the total return came from asset allocation.

finds that about 40% of the differences in returns across funds can be explained by their asset allocation decisions and 60% by security selection. When it comes to the level of returns, almost all of the returns can be explained by the asset allocation decision. Collectively, these studies suggest that the asset allocation decision has important consequences for your returns, and its importance increases with your time horizon.

While how much of actual portfolio returns are due to asset allocation is open to debate, there can be no denying its importance to overall portfolio returns. While the researchers looked at the allocation across financial assets – stocks, bonds and bills - alone, we would define the asset allocation decision much more broadly to include real assets, including real estate, and in the most general case, human capital. The asset allocation decision follows logically from the assessment of the risk preferences, cash needs and tax status of the investor. The portfolio manager has to decide on the mix of assets that maximizes the after-tax returns subject to the risk and cash flow constraints of the investor. This is what we would term the passive approach to asset allocation, where the investor's characteristics determine the right mix for the portfolio. In coming up with the mix, we draw on the lessons of diversification; asset classes tend to be influenced differently by macro economic events such a recessions or changes in inflation, and do not move in tandem. This, in turn, implies that diversifying across asset classes will yield better trade offs between risk and return than investing in any one asset class. The same can be said about expanding portfolios to include both domestic and foreign assets.

There is, however, an active component to asset allocation, which leads portfolio managers to deviate from the passive mix defined above, and one component is market timing. To the extent that portfolio managers believe that they can time markets, i.e., determine which markets are likely to go up more than expected and which less than expected, they will alter the asset mixes accordingly. Thus, a portfolio manager who believes that the stock market is over valued and is ripe for a correction, while real estate is under valued, may reduce the proportion of the portfolio that is allocated to equities and increase the proportion allocated to real estate. It should be noted that there are some who differentiate between these actions, that they call tactical asset allocation, and more drastic switches from stock to cash, which they call market timing. We see only a difference in degree and will draw no such distinction.

The Cost of Market Timing

If market timing were costless, you could argue that everyone should try to time markets, given the huge returns to getting it right. There are, however, significant costs associated with trying to time markets (and getting it wrong):

- In the process of switching from stocks to cash and back, you may miss the best years of the market. In an article, titled "The Folly of Stock Market Timing", Jeffrey examined the effects of annually switching from stock to cash and back from 1926 to 1982 and concluded that the potential downside vastly exceeds the potential upside. In his article on market timing in 1975, Bill Sharpe suggested that unless you can tell a good year from a bad year 7 times out of 10, you should not try market timing. This result is confirmed by Chua, Woodward and To, who use Monte Carlo simulations on the Canadian market and confirm you have to be right 70-80% of the time to break even from market timing.
- These studies do not consider the additional transactions costs that inevitably flow from market timing strategies, since you will trade far more extensively with these strategies. At the limit, a stock/cash switching strategy will mean that you will have to liquidate your entire equity portfolio if you decide to switch into cash and start from scratch again the next time you want to be in stocks.
- A market timing strategy will also increase your potential tax liabilities. To see why, assume that you have a strategy of selling your stocks after two good years in the market, based upon the empirical findings that a bad year is more likely to follow. You will have to pay capital gains taxes when you sell your stocks, and over your lifetime as an investor, you will pay far more in taxes.

In Summary

The perceived payoff from market timing is large and apparent, whereas the costs are often less visible. This must explain why so many portfolio managers and investors, their protestations to the contrary, engage in some market timing. In addition, the high profile of market strategists at all of the major investment firms suggests that the asset allocation decision is perceived to be an important one.

Its appeal to investors, notwithstanding, market timing remains an elusive dream for most. Looking back at market history, there have been far fewer successful market timers than successful stock selectors, and it is not clear whether even the few successes that can be attributed to market timing are more attributable to luck. Why is it so difficult to succeed at market timing? One very important reason is that there are fewer potential differential advantages that investors can build on when it comes to timing markets. For instance, it is unlikely that one can acquire an informational advantage over other investors at timing markets, but it is still possible, with sufficient research and private information, to get an informational advantage at picking stocks. Market timers contend that they can take existing information and use it more creatively or in better models to arrive at predictions for markets, but such approaches can be easily imitated, and imitation is the kiss of death for successful investment strategies.

Market Timing Approaches

There as probably as many market timing approaches as there are investors. Some of these approaches are based upon non-financial indicators, some on macroeconomic variables such as interest rates and business cycles and some draw on the valuation tools that we used to analyze individual stocks – discounted cashflow and relative valuation models.

Market Timing based upon Non-financial Indicators

Through the decades, there are some investors who have claimed to foretell the market's future by looking at non-financial indicators. Some of these indicators, such as whether the NFC or AFC team wins the Super Bowl are clearly of dubious origin and would fall into a category that we title spurious indicators. Other indicators such as the hemline index, which relates stock prices to the length of hemlines on skirts, fall into the grouping of "feel good indicators" that measure the overall mood of people in the economy, who after all are both the consumers who act as the engine for the economy and as investors determining prices. Finally, there are the "hype indicators" that measure whether market prices are becoming disconnected from reality.

Spurious Indicators

Millions of investors track what happens to their stocks and to the market every day and it is not surprising that they find other occurrences that seem to predict what the market will do in the next period. Consider one very widely talked-about indicator – who wins the Super Bowl.[3] In the 35 years that the Super Bowl has been played from 1966 to 2001, the winner has come from the National Football Conference (or is an old pre-merger NFL team like the Steelers or Colts) 25 years, and the market has risen in 22 out of the 25 years. In the 10 years that an American Football Conference team has won, the market has fallen 7 times. In fact, there are academic researchers who claim that the success rate of 83% (29 out of 35 years) is far too high to due to chance. [4]

---

[3] For those unfamiliar with the Super Bowl, it is played between the winner of the American Football Conference (AFC) and the winners of the National Football Conference (NFC). It is played on the last Sunday in January.

[4] See Krueger and Kennedy, who claim to have been first to spot the correlation.

So why not invest in the market after observing who wins the Super Bowl? There are several potential problems. First, we disagree that chance cannot explain this phenomenon. When you have hundreds of potential indicators that you can use to time markets, there will be some that show an unusually high correlation purely by chance. Second, a forecast of market direction (up or down) does not really qualify as market timing, since how much the market goes up clearly does make a difference. Third, you should always be cautious when you can find no economic link between a market timing indicator and the market. There is no conceivable reason who wins the Super Bowl should affect or be correlated with overall economic performance. Indicators such as these may make for amusing anecdotes at parties but can be lethal to your portfolio as market timing devices.

Feel Good Indicators

When people feel optimistic about the future, it is not just stock prices that are affected by this optimism. Often, there are social consequences as well, with styles and social mores affected by the fact that investors and consumers feel good about the economy. In the 1920s, for instance, you had the Great Gatsby and the go-go years, as people partied and the markets zoomed up. In the 1980s, in another big bull market, you had the storied excesses of Wall Street, documented in books like Liars Poker and movies like Wall Street. It is not surprising, therefore, that people have discovered linkages between social indicators and Wall Street. Consider, for instance, an index that has been around for decades called the hemline index that finds a correlation between the hemlines on women's skirts and the stock market. This politically incorrect index is based on the notion that shorter dresses and skirts are associated with rising stock prices whereas longer dresses are predictors of stock market decline. Assuming the index works, we would argue that you are seeing a manifestation of the same phenomenon. As people get more upbeat, fashions do seem to get more daring (with higher hemlines following) and markets also seem to go up. You could undoubtedly construct other indices that have similar correlations. For instance, you should expect to see a high correlation between demand at highly priced restaurants at New York City (or wherever young investment bankers and traders go) and the market.

The problem with feel good indicators, in general, is that they tend to be contemporaneous or lagging rather than leading indicators. In other words, the hemlines don't drop before the markets drop but in conjunction with or after a market drop. As an investor, these indicators are off little use, since your objective is to get out before the market drops and to get in before the market goes up.

Hype Indicators

It is said that Joseph Kennedy, a well known speculator on stocks in his own time, knew it was time to get out of the market when he heard his shoe-shine boy talking about stocks. In our own time, there are some who believe that the market peaked when financial channel CNBC's ratings exceeded those of long-running soap operas. In fact, one recent indicator called the "cocktail party chatter" indicator tracks three measures – the time elapsed at a party before talk turns to stocks, the average age of the people discussing stocks and the fad component of the chatter. According to the indicator, the less time it takes for the talk to turn to stocks, the lower the average age of the market discussants and the greater the fad component, the more negative you should be about future stock price movements.

Harking back to our discussion of bubbles, remember that propagation is critical to bubbles getting bigger. In our media world, this will involve print, television and the internet and an overflow into day-to-day conversations. Thus, the discussion at the water cooler in a typical business is more likely to be about stocks than about football or other such daily (and more normal) obsessions, when markets are buoyant.

While hype indicators, of all non-financial indicators, offer the most promise as predictors of the market, they do suffer from several limitations. For instance, defining what constitutes abnormal can be tricky in a world where standards and tastes are shifting – a high rating for CNBC may be indicate too much hype or may be just reflecting of the fact that viewers find financial markets to be both more entertaining and less predictable than a typical soap opera. Even if we decide that there is an abnormally high interest in the market today and you conclude (based upon the hype indicators) that stocks are over valued, there is no guarantee that stocks will not get more overvalued before the correction occurs. In other words, hype indicators may tell you that a market is overvalued, but they don't tell you when the correction will occur.

Market Timing based upon Technical Indicators

In chapter 7, we examined a number of chart patterns and technical indicators used by analysts to differentiate between under and over valued stocks. Many of these indicators are also used by analysts to determine whether and by how much the entire market is under or over valued. In this section, we consider some of these indicators.

Past Prices

In chapter 7, we looked at evidence of negative long term correlation in stock prices – stocks that have gone up the most in recent periods are more likely to go down in future periods. Studies do not seem to find similar evidence when it comes to the overall market. If markets have gone up significantly in the most recent years, there is no evidence that market returns in future years will be negative. If we consolidate stock returns from 1871 to 2001, into five-year periods, we find a positive correlation of .2085 between five-year period returns – in other words, positive returns over the last five years are more likely to be followed by positive returns than negative returns in the next 5 years. In table 12.1, we report on the probabilities of an up-year and a down-year following a series of scenarios, ranging from 2 down years in a row to 2 up years in a row, based upon actual stock price data from 1871 to 2001.

Table 12.1:Market Performance

Priors	*Number of occurrences*	*% of positive returns*	*Average return*
After two down years	19	57.90%	2.95%
After one down year	30	60.00%	7.76%
After one up year	30	83.33%	10.92%
After two up years	51	50.98%	2.79%

It is true that markets are more likely to go down after two years of positive performance than under any other scenario, but there is also evidence of price momentum, with the odds of an up year increasing if the previous year was an up year. Does this mean that we should sell all our stocks after two good years? We don't think so, for two reasons. First, the probabilities of up and down years do change but note that the likelihood of another good year remains more than 50% even after 2 consecutive good years in the market. Thus, the cost of being out of the market is substantial with this market timing strategy. Second, the fact that the market is overpriced does not mean that all stocks are over priced. As a stock picker, you may be able to find under valued stocks even in an over priced market.

Another price-based indicator that receives attention at least from the media at the beginning of each calendar year is the January indicator. The indicator posits that as January goes, so goes the year – if stocks are up, the market will be up for the year, but a bad beginning usually precedes a poor year.[5] According to the venerable Stock Trader's Almanac that is compiled every year by Yale Hirsch, this indicator has worked 88% of the time. Note, though that if you exclude January from the year's returns and compute the returns over the remaining 11 months of the year, the signal becomes much weaker and returns are negative only 50% of the time after a bad start in January. Thus, selling your stocks after stocks have gone down in January may not protect you from poor returns.

[5] Note that there are narrower versions of the January indicator, using just the first 5 or 10 days of January.

Trading Volume

There are some analysts who believe that trading volume can be a much better indicator of future market returns than past prices. Volume indicators are widely used to forecast future market movements. In fact, price increases that occur without much trading volume are viewed as less likely to carry over into the next trading period than those that are accompanied by heavy volume. At the same time, very heavy volume can also indicate turning points in markets. For instance, a drop in the index with very heavy trading volume is called a selling climax and may be viewed as a sign that the market has hit bottom. This supposedly removes most of the bearish investors from the mix, opening the market up presumably to more optimistic investors. On the other hand, an increase in the index accompanied by heavy trading volume may be viewed as a sign that market has topped out. Another widely used indicator looks at the trading volume on puts as a ratio of the trading volume on calls. This ratio, which is called the put-call ratio is often used as a contrarian indicator. When investors become more bearish, they sell more puts and this (as the contrarian argument goes) is a good sign for the future of the market.

Technical analysts also use money flow, which is the difference between uptick volume and downtick volume, as predictor of market movements. An increase in the money flow is viewed as a positive signal for future market movements whereas a decrease is viewed as a bearish signal. Using daily money flows from July 1997 to June 1998, Bennett and Sias find that money flow is highly correlated with returns in the same period, which is not surprising. While they find no predictive ability with short period returns – five day returns are not correlated with money flow in the previous five days – they do find some predictive ability for longer periods. With 40-day returns and money flow over the prior 40 days, for instance, there is a link between high money flow and positive stock returns.

Chan, Hameed and Tong extend this analysis to global equity markets. They find that equity markets show momentum – markets that have done well in the recent past are more likely to continue doing well,, whereas markets that have done badly remain poor performers. However, they find that the momentum effect is stronger for equity markets that have high trading volume and weaker in markets with low trading volume.

Volatility

In recent years, a number of studies have uncovered a relationship between changes in market volatility and future returns. One study by Haugen, Talmor and Torous in 1991 found that increases in market volatility cause an immediate drop in stock prices but that stock returns increase in subsequent periods. They looked at daily price volatility from 1897 through 1988 and look for time periods where the volatility has increased or decreased significantly, relative to prior periods. [6] They then look at returns both at the time of the volatility change and in the weeks following for both volatility increases and decreases, and their results are summarized in Figure 12.1:

Figure 12.1: Returns around volatility changes

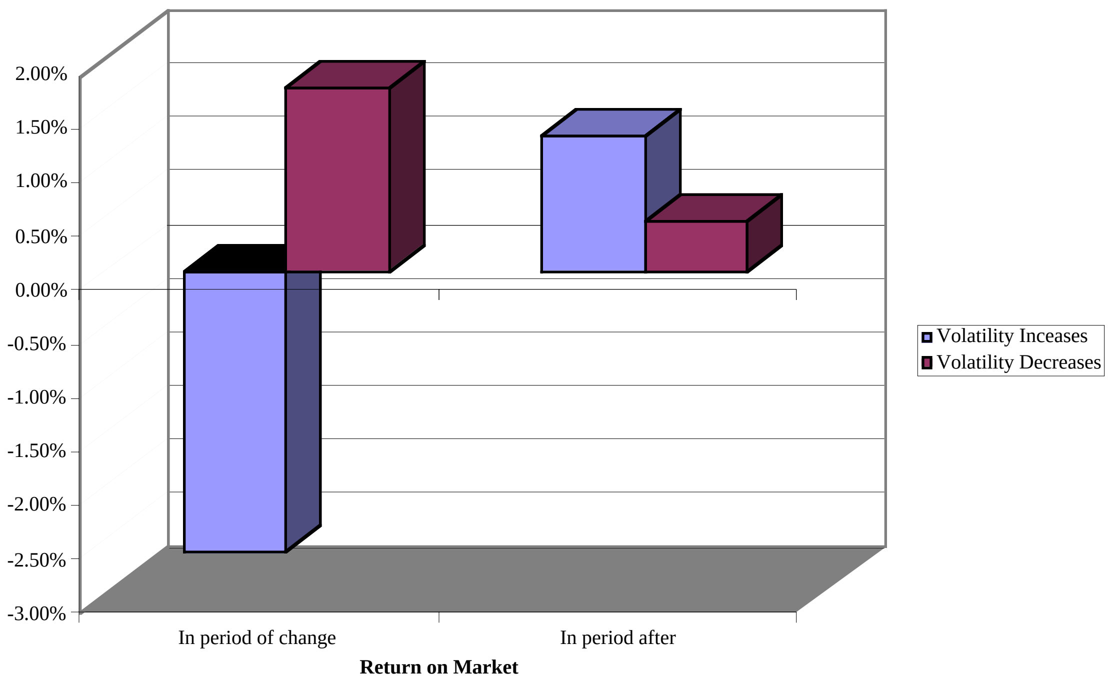

Source: Haugen, Talmor and Torous

Note that volatility increases cause stock prices to drop but that stock prices increase in the following four weeks. With volatility decreases, stock prices increase at the time of the volatility change, and they continue to increase in the weeks after, albeit at a slower pace.

Does this mean that you should buy stocks after an increase in volatility? Not necessarily. The increase in returns in the weeks following a volatility increase may just reflect the reality that stocks are riskier. However, if you believe that a surge in volatility is temporary and that stock volatility will revert back to normal levels, a strategy of buying stocks after an increase in equity market volatility may bear fruit.

Other Technical Indicators

There are a number of non-price indicators that are used by analysts to forecast future market movements. As with stock-specific technical indicators, market-wide

---

[6] Daily price volatility is estimated over four week windows. If the volatility in any four week window exceeds (falls below) the volatility in the previous four-week window (at a statistical significance level of 99%), it is categorized as an increase (decrease) in volatility.

indicators are often used in contradictory ways by momentum and contrarian analysts, with an increase in a specific indicator being viewed as bullish by one group and bearish by the other. Since we did cover technical indicators in depth in chapter 7, we will make only a short mention of some of these indicators in this section, categorized into price and sentiment indicators:

- Price indicators include many of the pricing patterns that we discussed in chapter 8. Just as support and resistance lines and trend lines are used to determine when to move in and out of individual stocks, they are also used to decide when to move in and out of the stock market.
- Sentiment indicators try to measure the mood of the market. One widely used measure is the confidence index which is defined to be the ratio of the yield on BBB rated bonds to the yield on AAA rated bonds. If this ratio increases, investors are becoming more risk averse or at least demanding a higher price for taking on risk, which is negative for stocks. Another indicator that is viewed as bullish for stocks is aggregate insider buying of stocks. When this measure increases, according to its proponents, stocks are more likely to go up.[7] Other sentiment indicators include mutual fund cash positions and the degree of bullishness among investment advisors/newsletters. These are often used as contrarian indicators – an increase in cash in the hands of mutual funds and more bearish market views among mutual funds is viewed as bullish signs for stock prices.[8]

While many of these indicators are used widely, they are mostly backed with anecdotal rather than empirical evidence.

Market Timing based upon Normal Ranges (Mean Reversion)

There are many investors who believe that prices tend to revert back to what can be called normal levels after extended periods where they might deviate from these norms. With the equity market, the normal range is defined usually in terms of price earnings (PE) ratios whereas with the bond market, a normal range of interest rates is used to justify betting on market direction.

---

[7] Chowdhury, Howe and Lin (1993) find a positive correlation between aggregate insider buying and market returns but report that a strategy based upon the indicator would not earn enough to cover transactions costs.

[8] See “Investor sentiment and Stock Returns’ by Fisher and Statman, Financial Analysts Journal, March/april 2000. They examined three sentiment indicators – the views of Wall Street strategists, investment newsletters and individual investors - and concluded that there is indeed evidence supporting a contrarian investment strategy

Is there a normal range for PE ratios?

Buy if the PE drops below 12 and sell if it rises above 18. You will see variations of this advice in many market timing newsletters. A more academic version of this argument was made by Campbell and Shiller who looked at PE ratios from 1871 to recent years and concluded that stocks revert back to a PE ratio of about 16 times normalized earnings. They defined normalized earnings as the average earnings over the previous 10 years. The implicit belief here is that there is a normal range for PE ratio and that if the PE rises above the top end of the range, stocks are likely to be overvalued, whereas if they fall below the bottom of the range, they are likely to be overvalued. While the approach is straightforward, where does the normal range of PE ratios come from? In most cases, it seems to come from looking at history and attaching a subjective judgment on the upper and lower limits. A slightly more sophisticated approach to estimating a range would require us to estimate the standard deviation in PE ratios over time and use it to compute a range – two standard deviations on either side of the average would give you a range outside which you should fall only 5% of the time by chance.

Consider, for instance, figure 12.2 which presents PE ratios for the S&P 500 going back to 1960.

Figure 12.2: PE Ratio for S&P 500: 1960-2001

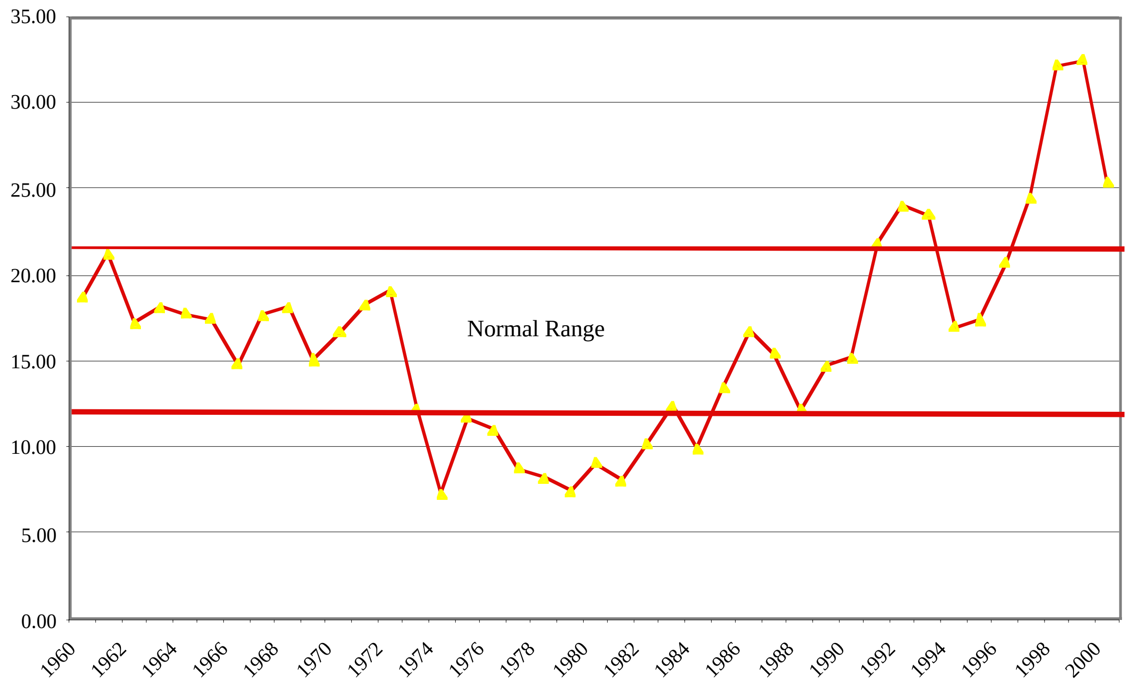

We have attempted to draw a normal range for interest rates in the United States, based upon history, though it indicates the subjective judgments that we had to make along the way. Based upon our band, stocks would be considered as overvalued if they traded at a PE ratio greater than 22 or less than 12.

The limitations of this approach should be obvious. In addition to trusting history to repeat itself, we are making two other assumptions. The first is that we can identify a normal trading range by looking at historical data. As you can see from the graph, you will not get any consensus – someone else looking at this graph might end up with a different band for PE. The second assumption is that the fundamentals have not shifted significantly over time. If interest rates are much lower today than they have been historically, you would expect stocks to trade at much higher PE ratios than they have historically. How much higher? We will look at this question in more detail in the later parts of this chapter.

> *PE Ratios for U.S stocks*: Take a look at the PE ratio for the S&P 500 each year going back to 1960.

Normal Range of Interest Rates

Some analysts hypothesize that market interest rates move within a normal range. Under this hypothesis, when interest rates approach the high end of the range, they are more likely to decrease, and when they approach the low end of the range, they are more likely to increase. This hypothesis is corroborated by two pieces of evidence:

1. Slope of the Yield Curve: The yield curve, which reflects future expectations about interest rates, is more likely to be downward sloping when interest rates are high than when there are low. Thus, investors are more likely to expect interest rates to come down if they are high now and go up, if they are low now. Table 12.2 below summarizes the frequency of downward sloping yield curves as a function of the level of interest rates.[9]

Table 12.2: Yield Curves and the Level of Interest Rates

	1-year Corporate Bond Rate	*Slope of Yield Curve*		
		Positive	*Flat*	*Negative*
	Above 4.40%	0	0	20
1900-70	3.25% - 4.40%	10	10	5
	Below 3.25%	26	0	0
1971-2000	Above 8.00%	4	1	3
	Below 8.00%	15	6	1

[9] Some of this table is extracted from Wood (1984).

This evidence is consistent with the hypothesis that maintains interest rates move within a normal range; when they approach the upper end (lower end) of the normal range, the yield curve is more likely to be downward sloping (upward sloping).

2. Interest rate level and expected change: More significantly, investors' expectations about future interest rate movements seem to be borne out by actual changes in interest rates. When changes in interest rates are regressed against the current level of interest rates, there is a negative and significant relationship between the level of the rates and the change in rates in subsequent periods, i.e., there is a much greater likelihood of a drop in interest rates next period if interest rates are high in this one, and a much greater chance of rates increasing in future periods if interest rates are low in this one. For instance, using treasury bond rates from 1970 to 1995 and regressing the change in interest rates (Δ Interest Rate_t) in each year against the level of rates at the end of the prior year (Interest Rate_{t-1}), we arrive at the following results:

$$\Delta \text{ Interest Rate}_t = 0.0139 - 0.1456 \text{ Interest Rate}_{t-1} \qquad R^2=.0728$$
$$(1.29) \qquad (1.81)$$

This regression suggests two things. One is that the change in interest rates in this period is negatively correlated with the level of rates at the end of the prior year; if rates were high (low), they were more likely to decrease (increase). Second, for every 1% increase in the level of current rates, the expected drop in interest rates in the next period increases by 0.1456%.

This evidence has to be considered with some caveats. The first is that the proportion of interest rate changes in future periods explained by the current level of rates is relatively small (about 7.28%); there are clearly a large number of other factors, most of which are unpredictable, that affect interest rate changes. The second is that the normal range of interest rates, which is based upon past experience, might shift if the underlying expectations of inflation change dramatically as they did in the 1970s in the United States. Consequently, many firms that delayed borrowing in the early part of that decade, because they thought that interest rates were at the high end of the range, found themselves facing higher and higher rates in each of the following years.

Hindsight is 20/20

Market timing always seems simple when you look back in time. After the fact, you can always find obvious signals of market reversals – bull markets turning to bear markets or vice versa. Thus, in 2001, there were investors who looked back at 1999 and bemoaned the fact that they missed getting out of stocks when the market topped at the end of that year. At that time, though, the signs were not so obvious. There were analysts who argued that the market was overvalued and indicators that supported that point of view, but there were just as many analysts, if not more, who saw the market continuing to rise and had supporting models.

.In practice, there is almost never a consensus among investors on whether markets have hit bottom or peaked at the time that it occurs. It is an interesting fact that optimism about the future is greatest just as markets top out and the market mood is darkest just as markets turn around. To succeed at market timing, you cannot wait until a bottom has been established before buying or for a market top before selling. If you do, you will miss much of the subsequent payoff.

Market Timing based upon Fundamentals

Just as the prices of individual stocks must reflect their cashflows, growth potential and risk, entire markets (equity, bond and real asset) have to reflect the fundamentals of these assets. If they do not, you can argue that they are misvalued. In this section, we consider two ways in which we can bring fundamentals into market timing models. In the first, we try to develop market timing strategies based upon the level of fundamental variables – interest rates and economic growth, for instance. In the second, we try to extend the valuation techniques developed for individual stocks to markets.

Fundamental Indicators

You can try to time markets by developing simple signals based upon macro economic variables. In this section, we will consider some of these signals – some old and some new – that have been used by portfolio managers as market timing tools.

Short term Interest Rates

Buy stocks when short-term rates (treasury bills) are low and sell them when short term rates are high, or so goes the conventional wisdom. But is there a basis to this advice? In table 12.3, we examine stock returns under four treasury bill scenarios –after a decline in rates of more than 1% over the prior year, a drop of between 0 and 1%, an increase in rates of less than 1% and an increase of more than 1% between 1928 and 2001.

Table 12.3: Stock Returns and Treasury Bill Rates

		In following year	
Change in T.Bill rate	*Number of years*	*% of up years*	*Average Annual returns*
Drop by more than 1%	10	70%	10.58%
Drop between 0 and 1%	24	75%	13.17%

Increase between 0 and 1%	26	69.23%	11.94%
Incrase more than 1%	13	61.54%	8.90%

In this case, there is surprisingly strong empirical evidence backing up the proposition that a drop in the treasury bill rate seems to predict high stock market returns. Generally speaking, markets are more likely to go up in years after the treasury bill rate has decreased and earn higher returns for investors.[10]

This result has been confirmed by a number of academic studies. In 2001, Ang and Baekart documented that treasury bill rates dominate other variables as a predictor of short term stock market movements. A 1989 study by Breen, Glosten and Jagannathan evaluated a strategy of switching from stock to cash and vice versa, depending upon the level of the treasury bill rate and conclude that such a strategy would have added about 2% in excess returns to an actively managed portfolio.

In a 2002 study that does raise cautionary notes about this strategy, Abhyankar and Davies examine the correlation between treasury bill rates and stock market returns in sub-periods from 1929 to 2000. They find that almost all of the predictability of stock market returns comes from the 1950-1975 time period, and that short term rates have had almost no predictive power since 1975. They also conclude that short rates have more predictive power with the durable goods sector and with smaller companies than they do with the entire market.

In conclusion, then, you should be aware of how high or low short term rates are when you invest in the market, but the value of short term rates as a predictor of stock market movements has decreased over the last few decades. Its remaining predictive power seems to be restricted to the short term and to sub-sectors of the market.

Treasury Bond Rate

Intuitively, it is the treasury bond rate – the long-term riskless rate – that should have a much stronger impact on stock prices, since it offers a direct alternative to investing in stocks for the long term. If you can make 8% investing risklessly in treasuries for the next 30 years, why would you settle for less when investing in stocks? Thus, we should expect to see stock prices go up if the treasury bond rate comes down and go down, if the rate goes up. Figure 12.3 presents a scatter plot of returns on stock returns each year and the T.Bond rate at the end of the prior year:

---

[10] You could do a similar study using the level of treasury bill rates, but treasury bill rates were much lower prior to the second world war.

Figure 12.3: T.Bond Rates and Stock Returns

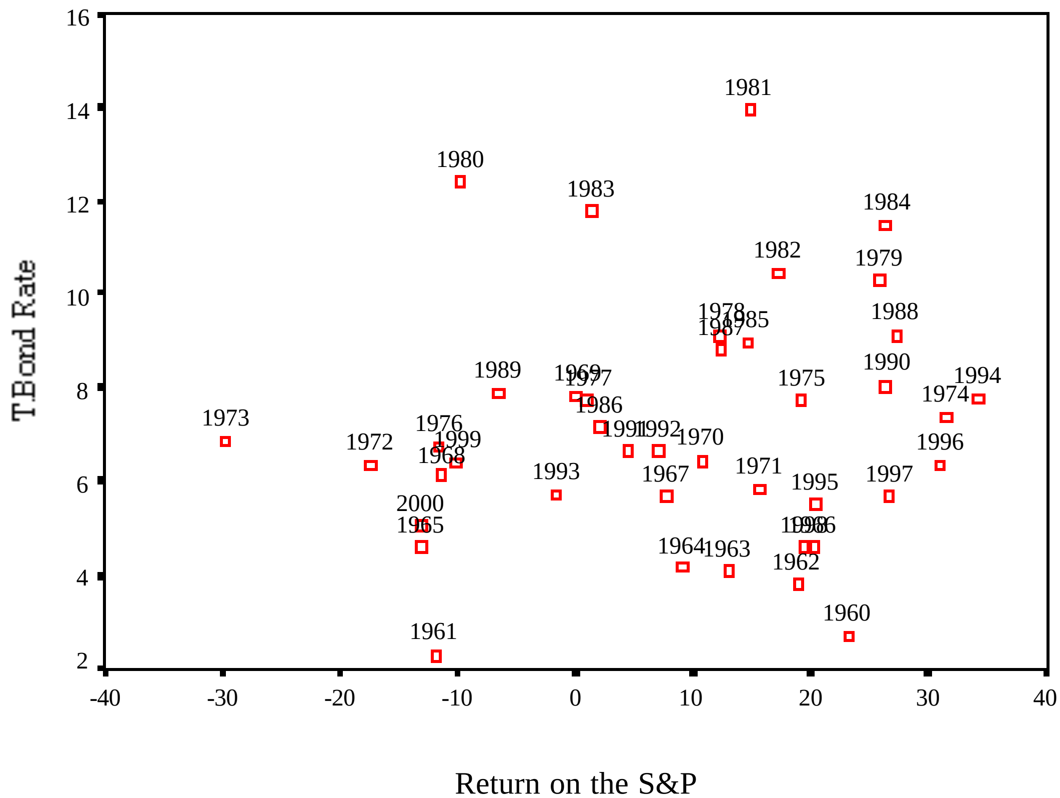

In 1981, for instance, the treasury bond rate at the start of the year was 14% and the return on the stock index during the year was 15%. In 1961, the treasury bond rate was 2% and the return on stocks during the year was –11%. If there is a relationship between treasury bond rates at the start of a period and stock returns during the period, it is not strong enough to be obvious and there seems to be little support for the proposition that stock returns are high when interest rates are low and low when interest rates are high. In fact, stocks did very well in 1982, even though interest rates were very high at the beginning of the year and very badly in 1961, notwithstanding the fact that the treasury bond rate was only 2% at the end of the prior year.

This link between treasury bond rates and stock returns should become even stronger if we consider how much we can earn as a return on stocks. You could define this return narrowly as the dividend yield (dividends/current stock prices) or use a much broader measure, such as earnings yield, which looks at the overall earnings on the market as a percent of the current level of the index. The earnings yield is the inverse of the price earnings ratio and is used widely by market strategists. Rather than focus on the level of the treasury bond rate, market strategists often look at the difference between earnings yields and the treasury bond rate. In simpler terms, they believe that it is best to invest in stocks when earnings yields are high, relative to the treasury bond rate. In fact, there are some strategists who believe that stocks are over valued when the earnings yield is lower than the treasury bond rate. To examine this proposition, we looked at the difference between the earnings yield and the T.Bond rate at the end of every year from 1960 to 2000 and the returns on the S&P 500 in the following year (see table 12.4)

Table 12.4: Earnings Yield, T.Bond Rates and Stock Returns: 1960 –2001

Earnings yield - T.Bond Rate	*Number of years*	*Average*	*Standard Deviation*	*Maximum*	*Minimum*
> 2%	8	11.33%	16.89%	31.55%	-11.81%
1 -2%	5	-0.38%	20.38%	18.89%	-29.72%
0-1%	2	19.71%	0.79%	20.26%	19.15%
-1-0%	6	11.21%	12.93%	27.25%	-11.36%
-2-1%	15	9.81%	17.33%	34.11%	-17.37%
< -2%	5	3.04%	8.40%	12.40%	-10.14%

The relationship is tenuous at best. When the earnings yield exceeds the treasury bond rate by more than 2%, which has occurred in 8 out of the 41 years, the return on the S& P 500 in the following year has averaged 11.33%. However, the returns are almost as good when the earnings yield has lagged the treasury bond rate by zero to 1%. It is true that the annual returns are only 3.04% in the five years following periods when the earnings yield was lower than the treasury bond rate by more than 2%, but the annual returns were also negative in the 5 years when the earnings yield exceeded the treasury bond rate by 1-2%. Thus, there seems to be little historical support for using earnings yield and treasury bond rates to predict future stock market movements.

Business Cycles

As with treasury bonds, there is an intuitive link between the level of stock prices and economic growth. You would expect stocks to do much better in economic booms than during recessions. What makes this relationship tricky, however, is that market movements are based upon predictions of changes in economic activity in the future, rather than levels of activity. In other words, you may see stock prices rising in the depths of a recession, if investors expect the economy to begin recovering in the next few months. Alternatively, you may see stock prices drop even in the midst of robust economic growth, if the growth does not measure up to expectations. In figure 12.4, we have graphed the S&P 500 index and GDP growth going back to 1960:

Figure 12.4: Real GDP Growth and Stock Retur

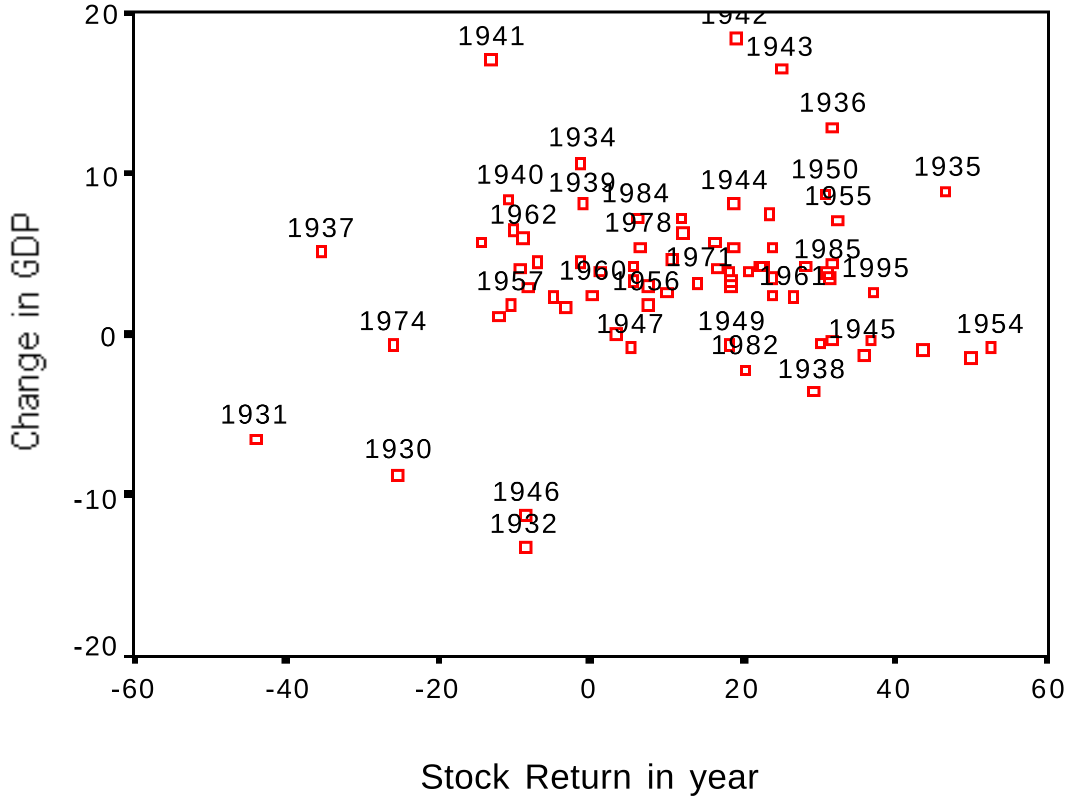

There is a positive relationship between GDP growth during a year and stock returns during the year, but there is also a lot of noise in the relationship. Even if the relationship were strong enough to pass muster, you cannot use it for market timing unless you can forecast real economic growth. The real question then becomes whether you can make forecasts of future stock market movements after observing economic growth in the last year. To examine whether there is any potential payoff to investing after observing economic growth in the prior year, we looked at the relationship between economic growth in a year and stock returns in the following year, using data from 1929 to 2001 in table 12.5:

Table 12.5:Real Economic Growth as a predictor of Stock Returns: 1960 – 2001

		Returns in Next Year			
GDP Annual Growth	*Number of years*	*Average Return*	*Standard deviation in returns*	*Best Year*	*Worst Year*
>5%	23	10.84%	21.37%	46.74%	-35.34%
3.5%-5%	22	14.60%	16.63%	52.56%	-11.85%
2-3.5%	6	12.37%	13.95%	26.64%	-8.81%

0-2%	5	19.43%	23.29%	43.72%	-10.46%
<0%	16	9.94%	22.68%	49.98%	-43.84%
All years	72	12.42%	19.50%	52.56%	-43.84%

There seems to be no clearly discernible relationship between returns next year and GDP growth this year. It is true that the years with negative GDP growth are followed by the lowest stock returns, but the average stock returns in this scenario are barely higher than the average returns you would have earned if you had bought after the best economic growth years (growth exceeds 5%).

If you can forecast future growth in the economy, it can be useful at two levels. One is in overall market timing, since you will steer more of your funds into stocks prior to better-than-expected economic growth and away from stocks when you foresee the economy slowing. You can also use the information to over invest in those sectors that are most sensitive to the economic cycle – automobile and housing stocks, for instance – if you believe that robust economic growth is around the corner.

Intrinsic Value Models

One way in which we can take the individual fundamentals that we considered in the last section and consolidate them into one market view is to do an intrinsic valuation of the entire market. What, you might ask, is an intrinsic valuation? Back in chapter 4, we consider how an individual stock can be valued using a discounted cash flow model as the present value of expected cashflows in the future. A market is composed of individual assets, and if individual assets can be valued using discounted cashflow models, we see no reason why the entire market cannot be valued as the present value of expected cashflows. In this section, we consider how best to extend discounted cashflow models to valuing the market, and the value that may be added from doing so.

Extending DCF Models to the Market

Consider, for instance, the dividend discount model that we introduced in chapter 4. We argued that the value of a stock can be written as the present value of the expected dividends from owning the stock, discounted back at the cost of equity. Extending this argument to an index, the value of an index can also be written as the present value of the expected dividends on the index. Thus, if the dividends on the entire stock index are expected to be $ 40 next year, the expected growth rate in perpetuity is expected to be 4% and the cost of equity for the average risk stock is expected to be 9%, you could value the index as follows:

Value of index = Expected dividends next year / (Cost of equity – Expected growth rate)
= 40 / (.09 - .04) = 800

As with an individual stock, this model can be extended to allow for high growth. Thus, if you expected dividends to grow 10% a year for the next 5 years and then expect the growth rate to drop to 4% in perpetuity, the value of the index can be computed in Table 12.6.

Table 12.6: Valuing an Index with High Growth

	Dividends	Terminal value	Present Value
1	$40.00		$36.70
2	$44.00		$37.03
3	$48.40		$37.37
4	$53.24		$37.72
5	$58.56	$1,218.13	$829.76
	Value of Index =		$978.59

Note that the dividends grow at 10% until year 5 and that the terminal value of the index is based upon a 4% growth rate forever.

Terminal value – 58.56 (1.04)/(.09-.04) = $1,218.13

We noted one limitation of dividend discount models is that companies may not pay out what they can afford to in dividends or may choose alternative ways of returning cash to stockholders (stock buybacks, for instance). You can modify this model by replacing dividends with potential dividends (free cashflows to equity for the index) or by augmenting dividends with stock buybacks on the index.

Some Caveats

While the building blocks for discounted cashflow valuation may remain the same for individual stocks and the markets, there are some cautionary notes that need to be added when valuing entire markets.

- While we allowed for the possibility of high growth in the last section, you should be much more cautious about assuming high growth, both in terms of the growth rate and how long high growth will continue - for a market than you would be for an individual stock, especially when the market is broadly based. Consider, for instance, the S&P 500. Since it includes the 500 companies with the largest market capitalization, arguing that earnings for these companies will grow at a rate much higher than the growth rate of the economy implies that the profit margins of these companies will increase over time. While this is feasible in the short term, especially if the economy is coming out of a recession or if firms are restructuring, we do not see how this can be sustained in the long term.

- The cost of equity that we are considering here is the cost of equity for the entire index. If we are considering a broadly based equity index, this cost of equity should reflect the riskless rate and the risk premium that investors demand for investing in equities as a class.

On the plus side, you should have less trouble forecasting earnings and dividends for an index than you should with individual stocks. After all, you have the luxury of diversification. In other words, you may over estimate earnings on some stocks and under estimate earnings on other stocks, but your overall measure of earnings can still be fairly precise.

Illustration 12.1: Valuing the S&P 500 using a dividend discount model: January 1, 2001

On January 1, 2001, the S&P 500 index was trading at 1320. The dividend yield on the index based upon dividends paid in 2000 was only 1.43%, but including stock buybacks (from 2000) increases the composite dividend yield (dividends + stock buybacks) to 2.50%. Analysts were estimating that the earnings of the stocks in the index would grow 7.5% a year for the next 5 years. Beyond year 5, the expected growth rate is expected to be 5%, the nominal growth rate in the economy. The treasury bond rate was 5.1% and we will use a market risk premium of 4%, leading to a cost of equity of 9.1%:

*Intrinsic value of S*P 500 index*: Take a look at the most recent valuation of the S&P 500.

Cost of equity = 5.1% + 4% = 9.1%

The expected dividends (and stock buybacks) on the index for the next 5 years can be estimated from the current dividends and expected growth of 7.50%.

Current dividends = 2.50% of 1320 = 33.00

	1	*2*	*3*	*4*	*5*
Expected Dividends =	\$35.48	\$38.14	\$41.00	\$44.07	\$47.38
Present Value =	\$32.52	\$32.04	\$31.57	\$31.11	\$30.65

The present value is computed by discounting back the dividends at 9.1%. To estimate the terminal value, we estimate dividends in year 6 on the index:

Expected dividends in year 6 = \$47.38 (1.05) = \$49.74

$$\text{Terminal value of the index} = \frac{\text{Expected Dividends}_6}{r - g} = \frac{\$49.74}{0.091 - 0.05} = \$1213$$

Present value of Terminal value = $\frac{\$1213}{1.091^5} = \$785$

The value of the index can now be computed:

Value of index = Present value of dividends during high growth + Present value of terminal value = \$32.52+32.04+31.57+\$31.11+ \$30.65+ \$785 = 943

Based upon this, we would have concluded that the index was over valued at 1320.

How well do intrinsic valuation models work?

How well would a strategy of buying the index when it is intrinsically undervalued and selling when it is intrinsically overvalued do? It is difficult to answer this question because it depends upon the inputs you estimate for the intrinsic valuation model and your time horizon. Generally speaking, the odds of succeeding increase as the quality of your inputs improves and your time horizon lengthens. Eventually, markets seem to revert back to intrinsic value but eventually can be a long time coming.

There is, however, a significant cost associated with using intrinsic valuation models when they find equity markets to be overvalued. If you take the logical next step of not investing in stocks when they are overvalued, you will have to invest your funds in either other securities that you believe are fairly valued (such as short term government securities) or in other asset classes. In the process, you may end up out of the stock market for extended periods while the market is, in fact, going up. For instance, most intrinsic value models would have suggested that the equity market in the United States was overvalued starting in 1994. If you had followed through and not invested in equities until 2002 (when the models suggested that valuations were fair again), you would have lost far more (by not investing in the bull market between 1994 and 2000) than you would have gained (by not investing in the down markets of 2001 and 2002).

The problem with intrinsic value models is their failure to capture permanent shifts in attitudes towards risk or investor characteristics. This is because so many of the inputs for these models come from looking at the past. Thus, the risk premium used to come up with the cost of equity may have been estimated looking at historical data on stock and bond returns and dividends may reflect what companies did last year. If one or both have changed as a consequence of shifts in the market, you will get a misleading signal from intrinsic valuation models. In fact, many investors who used intrinsic value models bought stocks during the early 1970s as stock prices dropped and failed to take into account the seismic shifts created by the high inflation of that period.

Relative Value Models

In relative value models, you examine how markets are priced relative to other markets and to fundamentals. How is this different from intrinsic value models? While the two approaches shares some characteristics, this approach is less rigid, insofar as it does not require that you work within the structure of a discounted cashflow model. Instead, you either make comparisons of markets over time (the S&P in 2002 versus the S&P in 1990) or different markets at the same point in time (U.S. stocks in 2002 versus European stocks in 2002).

Comparisons Across Time

In its simplest form, you can compare the way stocks are priced today to the way they used to be priced in the past and draw conclusions on that basis. Thus, as we noted in the section on historic norms, many analysts argue that stocks today, priced at 25 times earnings, are too expensive because stocks historically have been priced at 15-16 times earnings.

While reversion to historic norms remains a very strong force in financial markets, we should be cautious about drawing too strong a conclusion from such comparisons. As the fundamentals (interest rates, risk premiums, expected growth and payout) change over time, the PE ratio will also change. Other things remaining equal, for instance, we would expect the following.

- An increase in interest rates should result in a higher cost of equity for the market and a lower PE ratio.
- A greater willingness to take risk on the part of investors will result in a lower risk premium for equity and a higher PE ratio across all stocks.
- An increase in expected growth in earnings across firms will result in a higher PE ratio for the market.

In other words, it is difficult to draw conclusions about PE ratios without looking at these fundamentals. A more appropriate comparison is therefore not between PE ratios across time, but between the actual PE ratio and the predicted PE ratio based upon fundamentals existing at that time.

Illustration 12.2: PE Ratios across time for the S&P 500

Figure 12.5 summarizes the Earnings/Price ratios for S&P 500, treasury bond rates and the difference between bond and bill rates at the end of each year from 1960 to 2000.

Figure 12.5: S&P 500- Earnings Yield, T.Bond rate and Yield spread

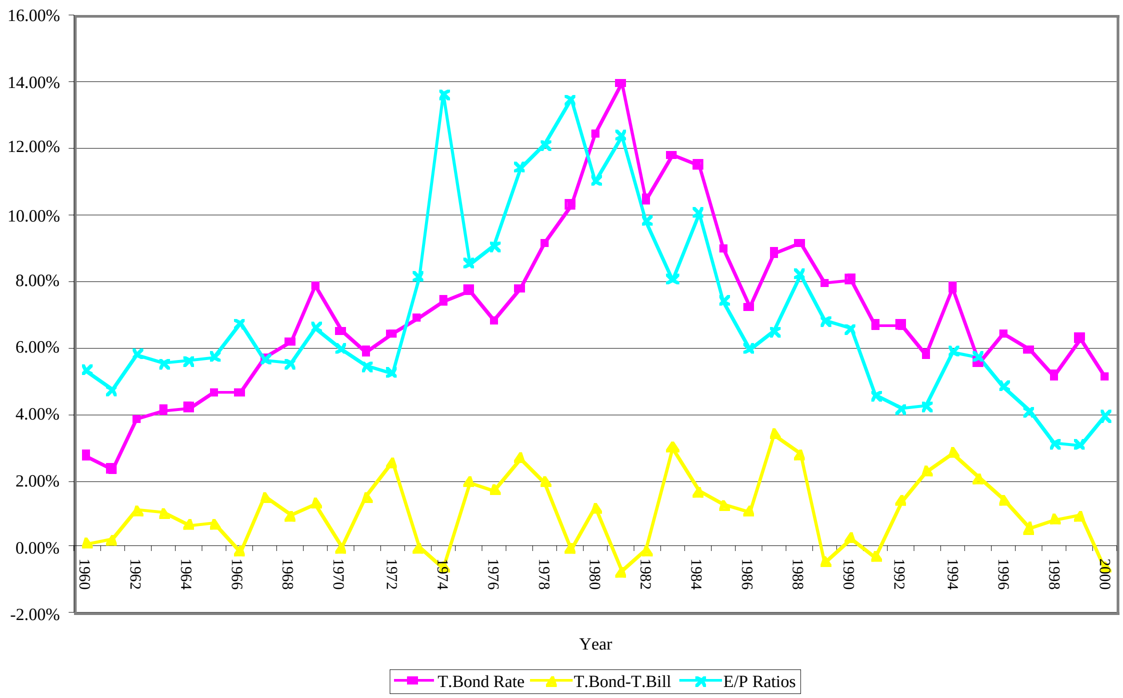

You do not need to be a statistician to note that earnings to price ratios are high (and PE ratios are low) when the treasury bond rates are high, and the earnings to price ratios decline when treasury bond rates drop. This strong positive relationship between E/P ratios and T.Bond rates is evidenced by the correlation of 0.6854 between the two variables. In addition, there is evidence that the term structure also affects the E/P ratio. In the following regression, we regress E/P ratios against the level of T.Bond rates and the yield spread (T.Bond - T.Bill rate), using data from 1960 to 2000.

E/P = 0.0188 + 0.7762 T.Bond Rate - 0.4066 (T.Bond Rate-T.Bill Rate) $R^2 = 0.495$
(1.93) (6.08) (-1.37)

Other things remaining equal, this regression suggests that

- Every 1% increase in the T.Bond rate increases the E/P ratio by 0.7762%. This is not surprising but it quantifies the impact that higher interest rates have on the PE ratio.
- Every 1% increase in the difference between T.Bond and T.Bill rates reduces the E/P ratio by 0.4066%. Flatter or negative sloping term yield curves seem to correspond to lower PE ratios and upwards sloping yield curves to higher PE ratios. While, at first sight, this may seem surprising, the slope of the yield curve, at least in the United

States, has been a leading indicator of economic growth with more upward sloped curves going with higher growth.

Based upon this regression, we predict E/P ratio at the beginning of 2001, with the T.Bill rate at 4.9% and the T.Bond rate at 5.1%.

$$E/P_{2000} = 0.0188 + 0.7762\ (0.051) - 0.4066\ (0.051-0.049) = 0.0599 \text{ or } 5.99\%$$

$$PE_{2000} = \frac{1}{E/P_{2000}} = \frac{1}{0.0599} = 16.69$$

Since the S&P 500 was trading at a multiple of 25 times earnings in early 2001, this would have indicated an over valued market. This regression can be enriched by adding other variables, which should be correlated to the price-earnings ratio, such as expected growth in GNP and payout ratios, as independent variables. In fact, a fairly strong argument can be made that the influx of technology stocks into the S&P 500 over the last decade, the increase in return on equity at U.S. companies over the same period and a decline in risk premiums could all explain the increase in PE ratios over the period.

> *Relativc value of S*P 500 index*: Take a look at the most recent relative valuation of the S&P 500.

Comparisons across Markets

Comparisons are often made between price-earnings ratios in different countries with the intention of finding undervalued and overvalued markets. Markets with lower PE ratios are viewed as under valued and those with higher PE ratios are considered over valued. Given the wide differences that exist between countries on fundamentals, it is clearly misleading to draw these conclusions. For instance, you would expect to see the following, other things remaining equal:

- Countries with higher real interest rates should have lower PE ratios than countries with lower real interest rates.
- Countries with higher expected real growth should have higher PE ratios than countries with lower real growth.
- Countries that are viewed as riskier (and thus command higher risk premiums) should have lower PE ratios than safer countries

Countries where companies are more efficient in their investments (and earn a higher return on these investments) should trade at higher PE ratios.

Illustration 12.3: Comparing PE ratios across markets

Table 12.7 summarizes PE ratios across different countries in July 2000, together with dividend yields and interest rates (short term and long term) at the time.

Table 12.7: PE Ratios for Developed Markets – July 2000

Country	*PE*	*Dividend Yield*	*2-yr rate*	*10-yr rate*	*10yr - 2yr*
UK	22.02	2.59%	5.93%	5.85%	-0.08%
Germany	26.33	1.88%	5.06%	5.32%	0.26%
France	29.04	1.34%	5.11%	5.48%	0.37%
Switzerland	19.6	1.42%	3.62%	3.83%	0.21%
Belgium	14.74	2.66%	5.15%	5.70%	0.55%
Italy	28.23	1.76%	5.27%	5.70%	0.43%
Sweden	32.39	1.11%	4.67%	5.26%	0.59%
Netherlands	21.1	2.07%	5.10%	5.47%	0.37%
Australia	21.69	3.12%	6.29%	6.25%	-0.04%
Japan	52.25	0.71%	0.58%	1.85%	1.27%
United States	25.14	1.10%	6.05%	5.85%	-0.20%
Canada	26.14	0.99%	5.70%	5.77%	0.07%

A naive comparison of PE ratios suggests that Japanese stocks, with a PE ratio of 52.25, are overvalued, while Belgian stocks, with a PE ratio of 14.74, are undervalued. There is, however, a strong negative correlation between PE ratios and 10-year interest rates (-0.73) and a positive correlation between the PE ratio and the yield spread (0.70). A cross-sectional regression of PE ratio on interest rates and expected growth yields the following.

PE Ratio = 42.62 – 360.9 (10-year rate) + 846.6 (10-year rate– 2-year rate) R^2=59%

(2.78) (-1.42) (1.08)

The coefficients are of marginal significance, partly because of the small size of the sample. Based upon this regression, the predicted PE ratios for the countries are shown in Table 12.8.

Table 12.8:Predicted PE Ratios for Developed Markets – July 2000

Country	*Actual PE*	*Predicted PE*	*Under or Over Valued*
UK	22.02	20.83	5.71%
Germany	26.33	25.62	2.76%

France	29.04	25.98	11.80%
Switzerland	19.6	30.58	-35.90%
Belgium	14.74	26.71	-44.81%
Italy	28.23	25.69	9.89%
Sweden	32.39	28.63	13.12%
Netherlands	21.1	26.01	-18.88%
Australia	21.69	19.73	9.96%
Japan	52.25	46.70	11.89%
United States	25.14	19.81	26.88%
Canada	26.14	22.39	16.75%

From this comparison, Belgian and Swiss stocks would be the most undervalued, while U.S. stocks would have been most over valued.

Illustration 12.7: An Example with Emerging Markets

This example is extended to examine PE ratio differences across emerging markets at the end of 2000. In table 12.9, the country risk factor is estimated, for the emerging markets.[11] It is scaled from zero (safest) to one hundred (riskiest).

Table 12.9: PE Ratios and Key statistics: Emerging Markets

Country	*PE Ratio*	*Interest Rates*	*GDP Real Growth*	*Country Risk*
Argentina	14	18.00%	2.50%	45
Brazil	21	14.00%	4.80%	35
Chile	25	9.50%	5.50%	15
Hong Kong	20	8.00%	6.00%	15
India	17	11.48%	4.20%	25
Indonesia	15	21.00%	4.00%	50
Malaysia	14	5.67%	3.00%	40
Mexico	19	11.50%	5.50%	30
Pakistan	14	19.00%	3.00%	45
Peru	15	18.00%	4.90%	50
Phillipines	15	17.00%	3.80%	45
Singapore	24	6.50%	5.20%	5

[11] These estimates come the Economist.

South Korea	21	10.00%	4.80%	25
Thailand	21	12.75%	5.50%	25
Turkey	12	25.00%	2.00%	35
Venezuela	20	15.00%	3.50%	45

Interest Rates: Short term interest rates in these countries

The regression of PE ratios on these variables provides the following –

PE = 16.16 – 7.94 Interest Rates + 154.40 Real Growth - 0.112 Country Risk R^2=74%

(3.61) (-0.52) (2.38) (-1.78)

Countries with higher real growth and lower country risk have higher PE ratios, but the level of interest rates seems to have only a marginal impact. The regression can be used to estimate the price earnings ratio for Turkey.

Predicted PE for Turkey = 16.16 – 7.94 (0.25) + 154.40 (0.02) - 0.112 (35) = 13.35

At a PE ratio of 12, the market can be viewed as slightly under valued.

Determinants of Success

Can you time markets by comparing stock prices now to prices in the past or to how stocks are priced in other markets? Though you can make judgments about market under or overvaluation with these comparisons, there are two problems with this analysis.

- Since you are basing your analysis by looking at the past, you are assuming that there has not been a significant shift in the underlying relationship. As Wall Street would put it, paradigm shifts wreak havoc on these models.
- Even if you assume that the past is prologue and that there will be reversion back to historic norms, you do not control this part of the process. In other words, you may find stocks to be over valued on a relative basis, but they become more over valued over time. In other words, convergence is neither timed nor even guaranteed.

Relative valuation of emerging markets: Take a look at the most recent relative valuation across emerging markets.

How can you improve your odds of success? First, you can try to incorporate into your analysis those variables that reflect the shifts that you believe have occurred in markets. For instance, if you believe that the influx of pension fund money into the equity markets over the last two decades has changed the fundamental pricing relationship, you can include the percent of stock held by pension funds into your regression. Second, you can have a longer time horizon, since you improve your odds on convergence.

The Information Lag with Fundamentals

If you are considering timing the market using macroeconomic variables such as inflation or economic growth, you should also take into account the time lag before you will get this information. Consider, for instance, a study that shows that there is high positive correlation between GDP growth in a quarter and the stock market's performance in the next. An obvious strategy would be to buy stocks after a quarter of high GDP growth and sell after a quarter of negative or low GDP growth. The problem with the strategy is that the information on GDP growth will not be available to you until you are two months into the next quarter.

If you use a market variable such as the level of interest rates or the slope of the yield curve to make your market forecasts, you are in better shape since this information should be available to you contemporaneously with the stock market. In building these models, you should be careful and ensure that you are not building a model where you will have to forecast interest rates in order to forecast the stock market. To test for a link between the level of interest rates and stock market movements, you would look at the correlation between interest rates at the beginning of each year and stock returns over the year. Since you can observe the former before you make your investment decision, you would have the basis for a viable strategy if you find a correlation between the two. If you had run the test between the level of interest rates at the end of each year and stock returns during the year, implementing an investment strategy even if you find a correlation would be problematic since you would have to forecast the level of interest rates first.

The Evidence on Market Timing

While we have looked at a variety of ways in which investors try to time markets from technical indicators to fundamentals, we have not asked a more fundamental question: Do those who claim to time markets actually succeed? In this section, we consider a broad range of investors who try to time markets and examine whether they succeed.

Mutual Fund Managers

Most equity mutual funds do not lay claims to market timing, but, in our view, they do try to time markets at the margin. We will begin by looking at whether they succeed on average. There are some mutual funds that claim market timing as their primary skill and these funds are called tactical asset allocation funds. We will look at the track records of these funds and pass judgment on whether their claims hold up.

Overall Evidence

How do we know that mutual funds try to time markets? While all equity mutual funds need to hold some cash – investments in treasuries and commercial paper – to meet redemption needs and for day-to-day operations, they collectively hold much more cash than is necessary. In fact, the only explanation for the cash balances that we observe at equity mutual funds is that mutual funds use them to signal their views of future market movements – they hold more cash when they are bearish and less cash when they are bullish. In figure 12.6 below, we present the average cash balance at mutual funds, each from 1980 to 2001 and the returns on the S&P 500 each year.

Figure 12.6: Mutual Fund Cash Holdings and Stock Returns

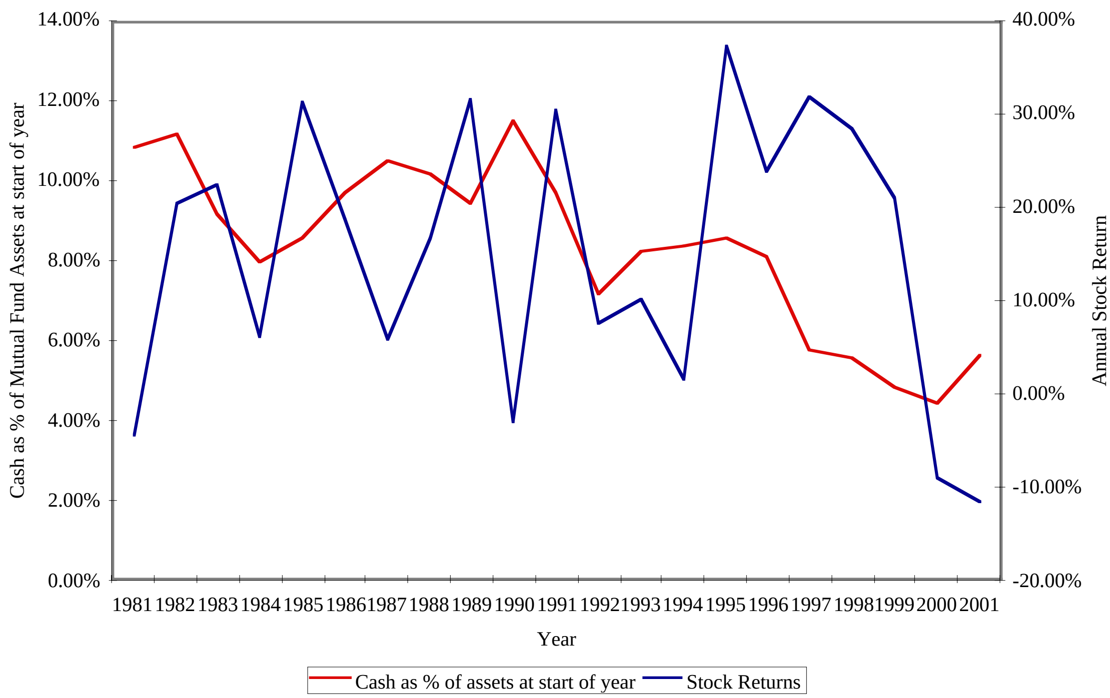

Source: Investment Company Institute

Note that the cash balances seem to increase after bad years for the market and decrease after good years, but there is little predictive power in the level of cash holdings. The question of whether mutual funds are successful at market timing has been examined widely in the literature going back four decades. A study in 1966 by Treynor and Mazuy suggested that we look at whether the betas of funds increase when the market return is large in absolute terms by running a regression of the returns on a fund against both market and squared market returns:

$$R_{\text{Fund, Period } t} = a + b\ \text{Return}_{\text{Market},t} + c\ \text{Return}_{\text{Market},t}^{2}$$

If a fund manager has significant market timing abilities, they argued, the coefficient "c" on squared returns should be positive. This approach, when tested out on actual mutual fund returns, yields negative values for the coefficient on squared returns, indicating negative market timing abilities rather than positive ones. In 1981, Merton and Henrikkson modified this equation to consider whether funds earned higher returns in periods when the market was positive, and found little evidence of market timing as well.

Tactical Asset Allocation and other Market timing Funds

In the aftermath of the crash of 1987, a number of mutual funds sprung up claiming that they could have saved investors the losses from the crash by steering them out of equity markets prior to the crash. These funds were called tactical asset allocation funds and made no attempt to pick stocks. Instead, they argued that they could move funds between stocks, treasury bonds and treasury bills in advance of major market movements and allow investors to earn high returns. Since 1987, though, the returns delivered by these funds has fallen well short of their promises. Figure 12.7 compares the returns on a dozen large tactical asset allocation funds over 5-year and 10-year periods (1987-97) to both the overall market and to fixed mixes – 50% in both stocks and bonds, and 75% stocks/25% bonds. We call the last two couch potato mixes, reflecting the fact that we are making no attempt to time the market.

Figure 12.7

Performance of Unsophisticated Strategies versus Asset Allocation Funds

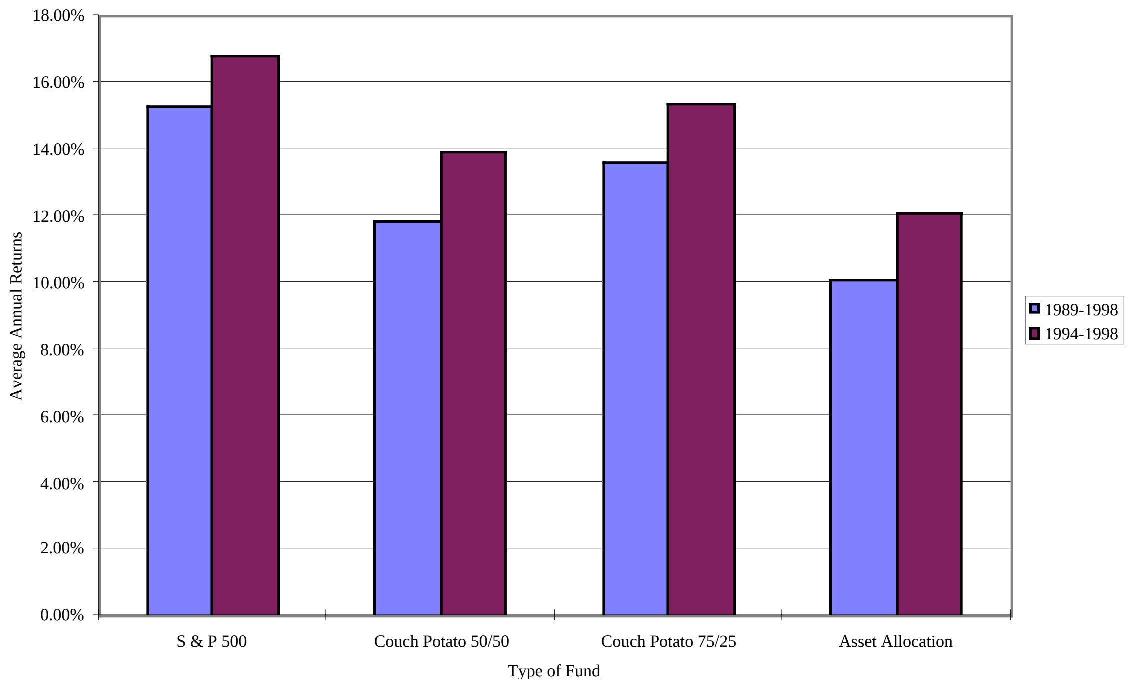

Source: Money Magazine

One critique of this study may be its focus on a few tactical asset allocation funds. In 1998, Becker, Ferson, Myers and Schill examined a much larger sample more than 100 asset allocation funds between 1990 and 1995 and also find little evidence of success at market timing at these funds.

Investment Newsletters

There are hundreds of investment newsletters that investors subscribe to for sage advice on investing. Some of these investment newsletters are centered on suggesting individual stocks for investors but some are directed towards timing the market. For a few hundred dollars, we are told, we too can be privy to private signals of market movements.

Campbell and Harvey (1996) examined the market timing abilities of investment newsletters by examining the stock/cash mixes recommended in 237 newsletters from 1980 to 1992. If investment newsletters are good market timers, you should expect to see the proportion allocated to stocks increase prior to the stock market going up. When the returns earned on the mixes recommended in these newsletters is compared to a buy and hold strategy, 183 or the 237 newsletters (77%) delivered lower returns than the buy and hold strategy. One measure of the ineffectuality of the market timing recommendations of these investment newsletters lies in the fact that while equity weights increased 58% of the time before market upturns, they also increased by 53% before market downturns. There is some evidence of continuity in performance, but the evidence is much stronger for negative performance than for positive. In other words, investment newsletters that give bad advice on market timing are more likely to continue to give bad advice than are newsletters that gave good advice to continue giving good advice.[12]

The only hopeful evidence on market timing comes from a study of professional market timers who are investment advisors. These timers provide explicit timing recommendations only to their clients, who then adjust their portfolios accordingly - shifting money into stocks if they are bullish and out of stocks if they are bearish. A study by Chance and Hemler (2001) looked at 30 professional market timers who were monitored by MoniResearch Corporation, a service monitors the performance of such advisors, and found evidence of market timing ability. It should be noted that the timing calls were both short term and frequent. One market timer had a total of 303 timing signals between 1989 and 1994, and there were, on average, about 15 signals per year across all 30 market timers. Notwithstanding the high transactions costs associated with following these timing signals, following their recommendations would have generated excess returns for investors.[13]

Market Strategists

The market strategists at major investment banks represent perhaps the most visible symbols of market timing. Their prognostications about the market are widely disseminated not only by their investment banks but also by the media. Abby Cohen (Goldman Sachs), Jeff Applegate (Lehman Brothers) and Byron Wien (Morgan Stanley) are all widely known. While much of what market strategists say about markets cannot be easily categorized as bullish or bearish– good market strategists are difficult to pin down when it comes to explicit forecasts – they also make specific recommendations on preferred asset allocation mixes that are presented in the Wall Street Journal. Table 12.10 provides the asset allocation mixes recommended by major investment banks in June 2002.

Table 12.10: Asset Allocation Mixes – Investment Bank Strategists

Firm	*Strategist*	*Stocks*	*Bonds*	*Cash*
A.G. Edwards	Mark Keller	65%	20%	15%

[12] A good market timing newsletter is likely to repeat its success about 50% of the time. A poor market timing newsletter has a 70% chance of repeating its poor performance.

[13] The study looked at excess returns after transactions costs but before taxes. By its very nature, this strategy is likely to generate large tax bills, since almost all of your gains will be taxed at the ordinary tax rate.

Banc of America	Tom McManus	55%	40%	5%
Bear Stearns & Co.	Liz MacKay	65%	30%	5%
CIBC World Markets	Subodh Kumar	75%	20%	2%
Credit Suisse	Tom Galvin	70%	20%	10%
Goldman Sach & Co.	Abby Joseph Cohen	75%	22%	0%
J.P. Morgan	Douglas Cliggott	50%	25%	25%
Legg Mason	Richard Cripps	60%	40%	0%
Lehman Brothers	Jeffrey Applegate	80%	10%	10%
Merrill Lynch & Co.	Richard Bernstein	50%	30%	20%
Morgan Stanley	Steve Galbraith	70%	25%	5%
Prudential	Edward Yardeni	70%	30%	0%
Raymond James	Jeffrey Saut	65%	15%	10%
Salomon Smith	John Manley	75%	20%	5%
UBS Warburg	Edward Kerschner	80%	20%	0%
Wachovia	Rod Smyth	75%	15%	0%

How do these allocation mixes yield market predictions? One way is to look at the percent allocated to stocks. More bullish market strategists will recommend a larger proportion of the portfolio be invested in stocks, whereas bearish strategists will overweight cash and bonds. The other is to look at changes in holdings recommended by the same strategist from period to period – an increase in the proportion allocated to stocks would indicate more bullishness. On both dimensions, the market timing skills of strategists are questionable. The Wall Street Journal, in addition to reporting the asset allocation mixes of strategists also compares the returns that would have been generated by following each bank’s allocation advice to the returns you would have made by being fully invested in stocks over 1-year, 5-year and 10-year periods. To counter the argument that it is unfair to compare a 100% equity portfolio to a asset allocation mix, the Journal also reports on the returns on a robot mix – a fixed allocation across stocks, bonds and bills. Figure 12.8 summarizes the returns on all three, as well as the returns you would have earned by following the strategist who had the best mixes over the period and the one with the worst mixes:

Figure 12.8

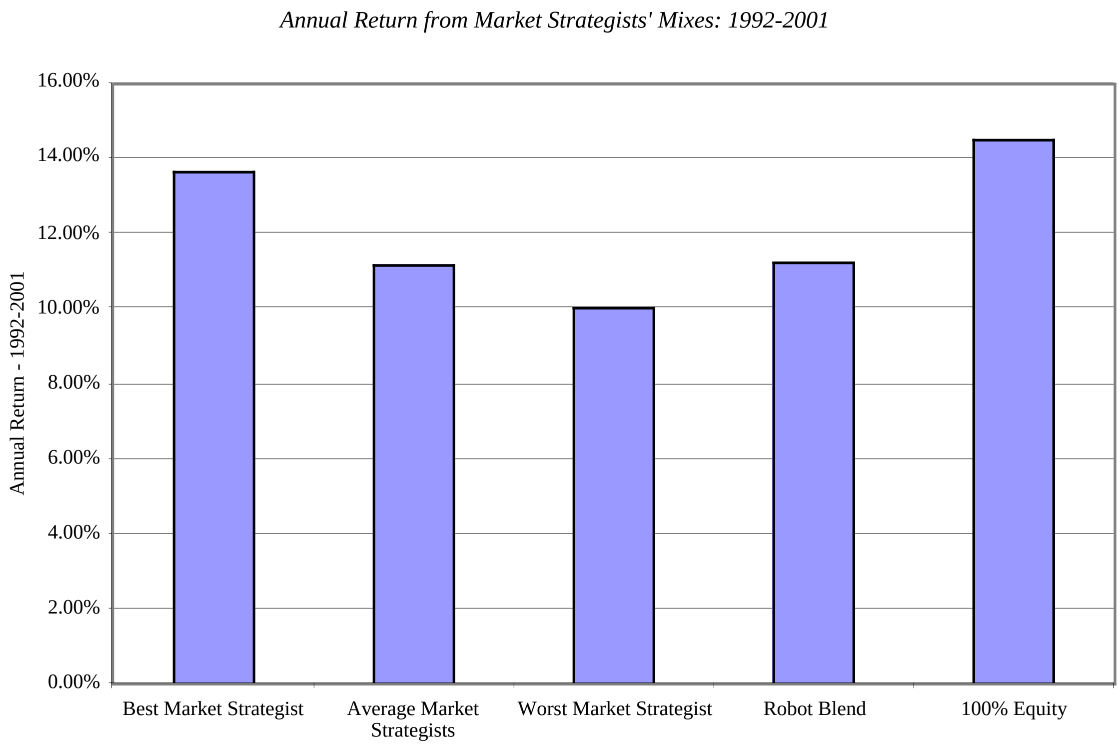

Source: Wall Street Journal

Note that the returns on the robot mix are higher than the average returns generated by following the average market strategists. Of the 16 banks that the Wall Street Journal tracks, only 5 would have generated returns higher than the robot mix over the period and even those would have well within a statistical margin for error. Finally, even the best strategist's asset mix would have underperformed a strategy of being fully invested in stocks. Overall, the evidence indicates that the market timing skills of leading market strategies are vastly overstated.

Market Timers: From Livermore to Acampora

Market timers are the meteors of the investment universe. While they attract a great deal of attention when they shine, they fade quickly. Looking at the high profile market timers (Market Gurus) over time, from Jesse Livermore in the early part of this century to Ralph Acampora, Prudential's flamboyant market strategist, in the 1990s, we find a diverse group.[14] Some were chartists, some used fundamentals and some were mysterious about their methods, but there are three common characteristics that they seem to share:

1. A capacity to see the world in black and white: Market gurus do not prevaricate. Instead, they make bold statements that seem outrageous when they make them about where the market will be 6 months or a year from now. Acampora, for instance, made his reputation with his call that the Dow would hit 7000 when it was at 3500.
2. A correct call on a big market move: All market timers make their reputation by calling at least one big market move. For Livermore, it was the market crash of 1929 and for Acampora, it was the bull market of the 1990s.
3. Outside personalities: Market gurus are born showmen (or show women), who use the media of their time as megaphones to publicize not only their market forecasts but the news of their successes. In fact, part of their success can be attributed to their capacity to make other investors act on their predictions, making these predictions, at least in the near term, self-fulfilling prophecies.

So why do great market gurus stumble? The very same factors that contribute to their success seem to underlie their failures. Their absolute conviction in their market timing abilities and their past successes seems to feed into more outrageous calls that ultimately destroy their reputations. Joe Granville, one of the market gurus of the late 1970s, for instance, spent all of the eighties recommending that people sell stocks and buy gold and his newsletter was ranked the worst, in terms of performance, for the decade.

Market Timing Strategies

If you can time markets, how can you take advantage of this skill? There are at least four ways you can do this, with varying degrees of risk associated with each. The first way is to adjust your mix of assets, allocating more than you normally would (given your time horizon and risk preferences) to markets that you believe are under valued and less than you normally would to markets that are overvalued. The second approach is to switch investment styles and strategies to reflect expected market performance. The third is to shift your funds within the equity market from sector to sector, depending upon your expectations of future economic and market growth. The fourth and most risky way to time markets is to speculate on market direction, using either financial leverage (debt) or derivatives to magnify profits.

---

[14] One of the best books on Livermore is the classic "Reminiscences of a Stock Market Operator" by Edwin LeFevre.

Asset Allocation

The simplest way of incorporating market timing into investment strategies is to alter the mix of assets – stocks, cash, bonds and other assets – in your portfolio. In fact, we judged the capacity of mutual fund managers and investment newsletters to time the market by looking at whether changes that they recommended in the asset allocation mix were useful predictors or future market movements. The limitation of this strategy is that you will shift part or all of your funds out of equity markets if you believe that they are over valued and can pay a significant price if the stock market goes up. If you adopt an all or nothing strategy, shifting 100% into equity if you believe that the market is under valued and 100% into cash if you believe that it is overvalued, you increase the cost of being wrong.

Style Switching

There are some investment strategies that do well in bull markets and others that do better in bear markets. If you can identify when markets are overvalued or undervalued, you could shift from one strategy to another or even from one investment philosophy to another just in time for a market shift.

For instance, in our discussion of growth versus value strategies in chapter 9, we noted the research done by Richard Bernstein which showed that growth investing does better that value investing when earnings growth is low for the entire market and that value investing beats growth investing when earnings growth is high. Bernstein also notes that growth investing tends to do much better when the yield curve is flat or downward sloping. In a related result, Pradhuman presents evidence that small cap investing yields higher returns than value investing when inflation is high and bond default spreads are low. You could take advantage of your market timing skills to shift from growth to value investing if you believe that markets are over valued and headed for a correction, or from value to growth investing if you consider them undervalued and likely to increase. In a paper that examines the payoff to style timing, Kao and Shumaker estimate the returns an investor would have made if she had switched with perfect foresight from 1979 to 1997 from value to growth stocks and back for both small cap and large cap stocks.[15] The annual returns from a perfect foresight strategy each year would have been 20.86% for large cap stocks and 27.30% for small cap stocks. In contrast, the annual return across all stocks was only 10.33% over the period.

---

[15] See “Equity Style Timing” by Kao and Shumaker, Financial Analysts Journal, Jan/Feb 1999, v55,37-48.

While this strategy looks promising, there may less to it than meets the high. In addition to the higher transactions costs and taxes that come with switching from one investment style to another, you also have the problem that most switches occur after the fact, reflecting not market timing skills but reaction to market performance. Thus, value investors seem to switch to growth investing after a market slowdown has occurred and not in advance of a slow down and growth investors switch to value investing well into a bull market. If, in fact, you do have skills as a market timer that make you confident enough to switch investment styles, you could argue that you would get a much bigger payoff by speculating, using index futures or options.

Sector Rotation

There are some investors who believe that staying out of the market, because of their views on the market is either too costly (because of the possibility that they could be wrong) or not feasible (because they are required to invest in the market). They may be able to parlay their market timing skills into superior returns by switching across sectors of the market as their views of the market changes. Thus, if they believe that the market will increase in the coming periods, due to stronger than expected real economic growth, they may switch into cyclical sectors. Alternatively, if their view is that interest rates will go up in the coming year and that this will cause the market to drop, they may switch out of financial stocks into companies that are less sensitive to interest rates (consumer products).

Best performing and worst performing sectors: Take a look at the 5 best performing and worst performing sectors of the last year.

While there are undoubtedly differences across sector rotation models, Stovall provides an excellent summary of the conventional wisdom on which sectors do best at each stage of the market in his book on sector rotation. Figure 12.9, which is extracted from this book, is presented below:

Figure 12.9: Sector Rotation - Sector Emphasis as a function of Market Cycle

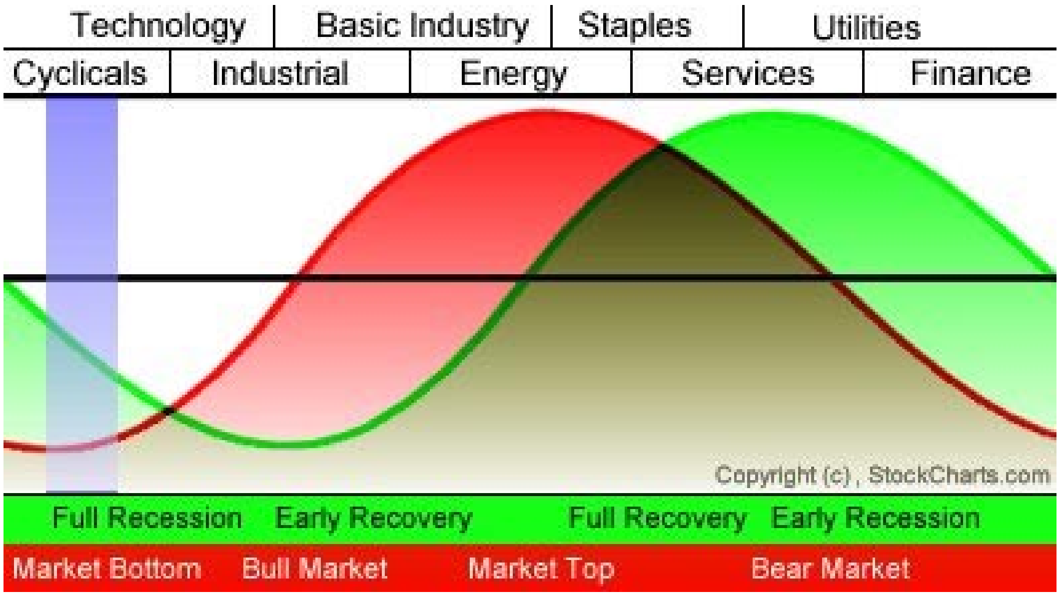

Source: Stovall

Note that the market lead is captured by the fact that the market both bottoms out and peaks before the economy. Your sector bets reflect this leading effect. You invest in cyclicals as the economy enters a recession (and the market hits bottom) and you shift into industrial and energy companies as the economy improves. If you can pick the right sectors to invest in each period, you would undoubtedly earn very high returns. For instance, a strategy where you would have invested in only the best performing sectors each year from 1970 to 1977 instead of the S&P 500 would have generated excess returns of 289%.[16] While this may not be feasible, investing in sectors that have done well in recent periods seems to provide at least short-term excess returns to investors. [17]

Building on the last point, sector rotation is not always based upon market timing views. There are some investors who use the stock selection approaches described in earlier chapters to pick sectors to invest in. For instance, investors who believe in price momentum may invest in sectors that have done well in the recent past, whereas those who are contrarians may invest in the sectors that delivered the worst performance in prior periods.

---

[16] See James Farreell, Jr., " Homogeneous Stock Groupings: Implications for Portfolio Management" in Guide to Portfolio Management.

[17] Sorensen and Burke (1986) report superior returns for at least two quarters from investing in industry groups that have done well in recent periods.

Speculation

The most direct way to take advantage of your market timing abilities is to buy assets in a market that you believe is under valued and sell assets in one that you believe is over valued. In the last decade, this is the strategy that has been adopted by market timing hedge funds to trade across equity, bond and currency markets to take advantage of what they see as potential mispricing, with varying degrees of success. Success can generate large returns because your have relatively little equity invested and because you benefit from both sides of the transactions – the under valued markets increasing and the over valued markets decreasing.

What can go wrong? The high leverage implicit in strategies where you buy some investments and sell short others exaggerates the effects of both success and failure. Thus, while the payoff to predicting markets is very large, the cost of failure is also very large. Whether you should adopt a speculative strategy based upon market timing is entirely dependent upon how confident you are about your predictions. The more confidence you have in your market timing abilities, the more leverage you can use in your strategies. Reviewing the empirical evidence on the performance of market timing strategies, it is quite clear that there are very few strategies that yield high success rates in equity markets. In contrast, though, there seem to be strategies that work a high percentage of the time in the currency and commodity markets, It should come as no surprise, therefore, that some of the biggest successes of hedge funds have come in these markets. Even in these markets, though, as the number of hedge funds increases, the potential for excess returns decreases.

Connecting Market Timing to Security Selection

Can you be a market timer and a security selector? We don't see why not, since they are not mutually exclusive philosophies. In fact, the same beliefs about markets that led you to become a security selector may also lead you to become a market timer. For example, if you believe that markets over react to new information, you may buy stocks after big negative earnings surprises but you may also buy the entire market, after negative economic or employment reports. In fact, there are many investors who combine asset allocation and security selection in a coherent investment strategu.

There are, however, two caveats to an investment philosophy that includes this combination. First, to the extent that you have differing skills as a market timer and as a security selector, you have to gauge where your differential advantage lies, since you have limited time and resources to direct towards your task of building a portfolio. Second, you may find that your attempts at market timing are under cutting your asset selection and that your overall returns suffer as a consequence. If this is the case, you should abandon market timing and focus exclusively on security selection.

Conclusion

Everyone wants to time markets, and it is not difficult to see the reasons for the allure. A successful market timer can deliver very high returns, with relatively little effort. The cost of market timing, though, is high both in terms of transactions costs (higher turnover ratios and tax bills) and opportunity costs (staying out of the market in years in which the market goes up). In fact, you need to be right about two-thirds of the time for market timing to pay off.

If you do decide to time markets, you have a wide range of market timing tools. Some are non-financial and range from the spurious like the Super Bowl indicator (whose correlation with the market is pure chance) to feel-good indicators (that measure the mood of people and thus the level of the market) to hype indicators (such as cocktail party chatter). Some market timing is centered around the macroeconomic variables that affect stocks prices – interest rates and economic growth – with the intuitive argument that you buy stocks when interest rates are low and in advance of robust economic growth. While the intuition may be impeccable, markets are tough to time because they are based upon predictions of these variables. Thus, high economic growth, by itself, may not lead to higher stock prices, if the growth was less than anticipated. One way to incorporate forecasted growth and risk into the analysis is to estimate the intrinsic value of the market – i.e., value the market as the present value of the expected cashflows you would get from investing in it. While this may yield good long term predictions, a better way of getting short term predictions may be an assessment of the value of the market, relative to its own standing in prior years and to other markets.

While the menu may be varied when it comes market timing strategies, there is little evidence of actual market timing success, even when we focus on those who claim to have the most expertise at it. Collectively, mutual funds seem to exhibit reverse market timing skills, at worst, and neutral market timing skills, at best, switching out of stocks (and into cash) just before big up movements in the markets and doing the reverse before stock price declines. Even those mutual funds that market themselves as market timers– the asset allocation funds – do not add any value from market timing. The asset allocation advice that comes from investment newsletters and market strategists also seems to suffer from the same problem of no payoffs.

If you believe that are the exception to this general rule of failure and that you can time markets, you can do it with varying degrees of gusto. The simplest strategy is to alter your asset allocation mix to reflect your market views, but this may require you to be out of stocks for extended periods. If you want to be fully invested in equities, you can try to switch investment styles ahead of market moves, moving from value investing (in periods of high earnings growth) to growth investing (if growth levels off) or shift your money across sectors of the market. Finally, if you have enough faith in your market timing abilities to pull it off, you can buy under valued and sell over valued markets, and make significant profits when they converge. The risk, of course, is that they will diverge and that you will see your portfolio suffer as a consequence.

Lessons for investors

To be a successful market timer, you have to

1. Be right about two thirds of the time: The payoff to timing markets correctly is high, but the cost of getting it wrong is also high. The payoff comes from staying out of the market in bad years, but the cost is that you may stay out of the market in good years.
2. Find an indicator that works consistently: There are dozens of indicators that are used to time markets but few of them seem to work consistently over long periods. Even those that do give you a sense of direction (up or down) but not of magnitude (how much up or how much down).
3. Recognize that you do not have many successful role models: Attempts at market timing on the part of professionals – money managers, investment newsletters and market strategists – have generally failed. Most of them tend to follow markets rather than lead them.

CHAPTER 13

READY TO GIVE UP? THE ALLURE OF INDEXING

Even those who believe that markets can be beaten will admit that it is difficult to do. Many investors begin life as active investors, convinced that they can beat the market, and end up conceding failure in this quest. For such investors and for those who never believed that they had a chance of beating the market, the most practical alternative is investing in an index fund. An index fund is designed to mimic an index and generate returns that are equal to the index (and not beat it). While we give up the chance of ever beating the market with an index fund, we do gain some significant benefits. First, there are almost no trading costs since there is little portfolio turnover, other than adding new stocks to the index and shedding those that leave. Second, their low turnover makes them more tax efficient than other funds.

Index funds were created in the early 1970s and have swiftly caught on as an investment alternative. Ironically, the performance of active investors and money managers provides the best endorsement for index funds. The returns earned by active money managers have trailed the return on the S&P 500 for much of the last two decades. In addition, the market timing abilities of money managers seems to be limited. Under such circumstances, you can and should ask legitimate questions about the economics of paying active money managers and investors to lose money for you.

The Mechanics of Indexing

How do you go about creating an index fund, and once created, how do you maintain it? Note than for an index fund to mimic an index, you not only need to hold every stock in the index, but the proportion invested in each stock also has to match up the weighting of the stock in the index. In this section, we will first explore how you would create an index fund that has these characteristics and then look at why some index funds may choose not to fully replicate the index that they purport to follow.

A Fully Indexed Fund

A fully indexed fund is relatively simple to create once the index to be replicate has been identified. There are three steps in the process. The first step is to obtain information on what stocks (or assets) go into the index and their weights in the index. Note that weighting schemes vary widely across indices. Some indices weight stocks based upon their market capitalizations (S&P 500), some have weights that reflect the history of the index (Dow 30) and some are equally weighted (NYSE composite). There are even indices where the weights are based upon trading volume on the exchange. The second step is to invest the fund in each of the stocks in the index, in exactly the same proportions that they are weighted in the index. An index fund has to be fully invested in the index, since cash holdings can cause the fund's returns to deviate from that of the index.

Once the index fund is created, there are two maintenance requirements. First, indices change over time, as some firms are removed from the index due to acquisitions, bankruptcies or because they no longer meet the criteria for the index and new firms are added. As the index changes, the index fund will have to change as well – the deleted firms will have to be sold and the added firms will have to be acquired. While changes are relatively rare for well established and mature indices such as the S&P 500, they can be much more frequent for indices that track changing or growing markets (as is the case with technology or emerging market indices). The second requirement is that the weights in the fund will have to be monitored and adjusted to reflect changes in the weights of the ingredients of the index. One of the advantages of indexing a market-capitalization weighted fund is that it is self-adjusting. In other words, if one stock doubles and another stock halves, the changes in their weighting in the index fund will mirror changes in the weighting in the index.

A Sampled Index Fund

In some cases, it may not be practical to construct a fully indexed fund. If you are replicating an index that contains thousands of stocks such as the Wilshire 5000 index or an index of illiquid assets or stocks, for instance, the transactions cost of creating and maintaining the index might be very large. One way around this problem is to create a fund that looks very much like the index but does not quite replicate this. You can accomplish this objective by sampling the index and buying some of the stocks listed in the index. Vanguard, which has long been the leader in the index fund business, uses sampling for its Total Market Fund, that attempts to replicate the performance of all stocks traded in the United States, as well as its small stock and European index funds.

How close can you get to the index with a sampled index fund? One way to measure the closeness is to measure the correlation between an index fund and the index that it tracks. A fully indexed fund should have a correlation of 100%, but a sampled index fund can get very close even with a much smaller number of stocks. Figure 13.1 presents the correlation between a sampled index fund and the S&P 500 as a function of the number of stocks in the index fund.

Figure 13.1: Correlation with the S&P 500 index

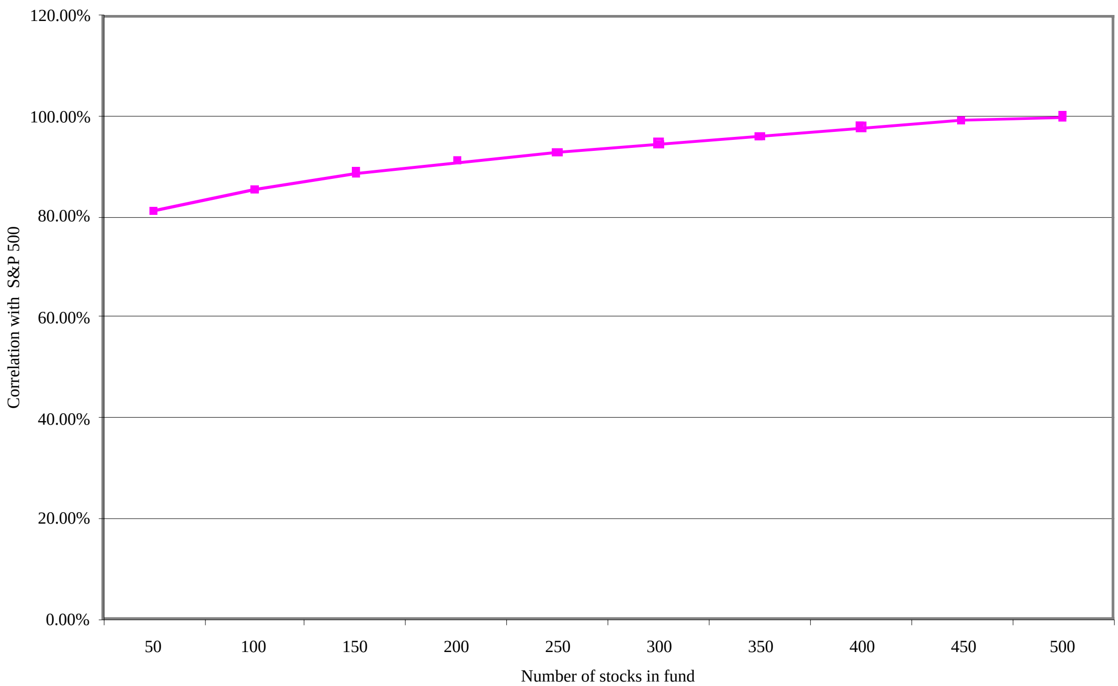

Note that the correlation is fairly high even with 50 stocks and very quickly converges on one.

A History of Indexing

In many ways, index funds owe their existence to academic research in financial economics. Prior to the 1960s, the conventional wisdom was that you entrusted your savings to money managers on Wall Street and they delivered high returns in exchange for management fees and transactions cost. No individual investor, it was believed, could compete against these professionals. As financial economics came into being as a discipline in the 1950s and 1960s, researchers increasingly found evidence tot he contrary. In fact, the efficient market hypothesis, which argued that market prices were the best estimates of value and that all of the time and resources spent by money managers trying to pick stocks was in vain, acquired a strong following. While most practitioners remained resistant to the notion that markets were efficient, there were some who saw opportunities in these academic findings. If diversification was the primary source of portfolio gains and picking stocks was not a fruitful exercise, why not create funds that were centered on diversification and minimize transactions costs? While indexing was initially made available to institutional investors in 1971, individual investors were able to play the indexing game when the Vanguard 500 Index fund made its debut in 1976. Over the last three decades has grown to become the second largest equity fund in the world, investing tens of billions of dollars each year in the S&P 500 Index.

In the last three decades, index funds have gradually increased their share of the overall market for not only individual investor's savings but for institutional funds such as pension and insurance money. In 2002, more than $ 120 billion of individual investor savings were invested in indexed mutual funds, and institutional investors contributed several hundred billion more to institutional index funds. This is not to suggest that index funds have not had their ebbs and flows. The initial impetus for index funds came from the abysmal returns delivered by portfolio managers in the 1970s. While many of these portfolio managers cannot be faulted for not seeing the long running bear market and high inflation of the period, the sub-par returns did draw attention to their management fees. Ever since, the growth of index funds seems to slow in boom periods (the early 1980s and the later 1990s), when it looks like active money management can pay off and picks up again when markets stagnate or decline. Figure 13.2 graphs out the growth of the best known index fund – the Vanguard 500 index fund from 1980 to 2001.

Figure 13.2: Growth of the Vanguard 500 index fund

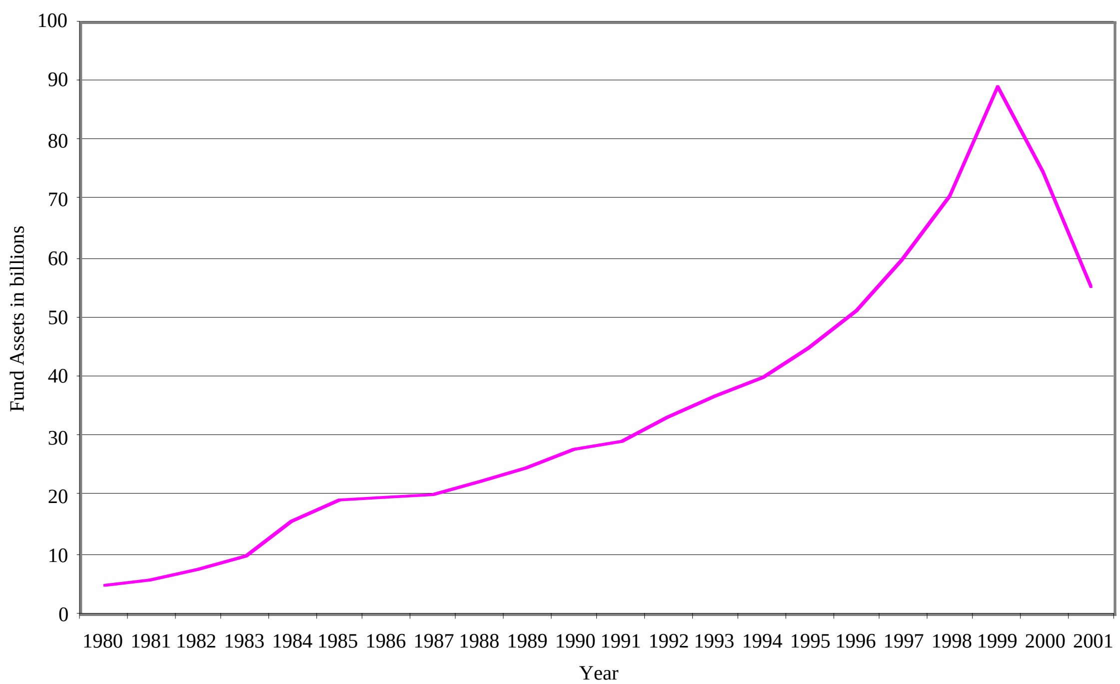

As index funds have grown, the choices have also proliferated. While the first few funds all indexed themselves to the S&P 500, you now see funds indexed to almost every conceivable index. For instance, Vanguard offered 15 index fund choices in March 2002 and they are summarized in table 13.1:

Table 13. 1: Index Fund Choices

Index Fund	*Stocks held*
Domestic Stock Funds	
500 Index Fund Inv	Seeks to track the investment returns of the S&P 500 Index
Calvert Social Index Fund	Seeks to track the Calvert Social Index, an index of large and midsize companies that have been screened for certain social and environmental criteria by its sponsor, the Calvert Group™. This fund may hold up to 40% of its assets in the stocks of a single industry sector.
Growth Index Fund Inv	Seeks to track the investment returns of the S&P 500/BARRA Growth Index, which measures the performance of the stocks in the S&P 500 Index with the highest price-to-book value ratio. This fund provides a convenient way to match the performance of a substantial portion of the nation's largest growth stocks.
Total Stock Mkt Idx Inv	Seeks to track the investment returns of the Wilshire 5000 Total Market Index, which includes virtually all regularly traded U.S. stocks.
Value Index Fund Inv	Seeks to track the investment returns of the S&P 500/BARRA Value Index, which measures the performance of the stocks in the S&P 500 Index with the lowest price-to-book value ratio.
Extended Mkt Index Inv	Seeks to track the investment returns of the Wilshire 4500 Completion Index, which measures the performance of small and midsize stocks.
Mid-Cap Index Fund Inv	Seeks to track the investment returns of the S&P 400 MidCap Index, which measures the performance of the stocks of all regularly traded midsize companies. This fund provides a convenient way to match the performance of companies with market capitalizations of $1.5 billion to $13 billion.
Small-Cap Growth Index	Seeks to track the investment returns of the S&P SmallCap 600/BARRA Growth Index, which measures the performance of the stocks in the S&P SmallCap 600 Index with the highest price-to-book value ratio.
Small-Cap Index Fund Inv	Seeks to track the investment returns of the Russell 2000 Index, which measures the performance of the 2,000 smallest companies out of the 3,000 largest U.S. companies.
Small-Cap Value Index	Seeks to track the investment returns of the S&P SmallCap 600/BARRA Value Index, which measures the performance of the stocks in the S&P SmallCap 600 Index with the lowest price-to-book value ratio.
REIT Index Fund Inv	Seeks to track the investment returns of the Morgan Stanley REIT Index, a barometer of the nation's largest publicly traded equity REITs (real estate investment trusts). The REITs held by the fund own commercial real estate, including shopping centers, offices, apartment buildings, and warehouses.
International/Global Stock	
Developed Markets Index	Seeks to track the investment returns of the Morgan Stanley Capital International Europe, Australasia, Far East (EAFE) Index, which measures the performance of 20 developed markets in Europe and the Pacific Basin. This fund invests in two Vanguard funds—European Stock Index Fund and Pacific Stock Index Fund—whose benchmarks together comprise the EAFE Index.
Emerging Mkts Stock Index	Seeks to track the investment returns of the Select Emerging Markets Free Index. This index measures the performance of 13 less-developed countries worldwide: Argentina, Brazil, the Czech Republic, Hungary, Indonesia, Israel, Korea, Mexico, the Philippines, Poland, South Africa, Thailand, and Turkey.
European Stock Index Inv	Seeks to track the investment returns of the Morgan Stanley Capital International Europe Index. This index measures the performance of 16 developed European countries: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.
Pacific Stock Index Inv	Seeks to track the investment returns of the Morgan Stanley Capital International Pacific Free Index. This index measures the performance of the Pacific Basin's

	established markets: Japan (the major index component), Australia, Hong Kong, New Zealand, and Singapore.
Total Int'l Stock Index	Seeks to track the investment returns of the Total International Composite Index, which measures the performance of nearly all the world's stock markets outside the United States and Canada. This fund invests in three Vanguard funds: European Stock Index Fund, Pacific Stock Index Fund, and Emerging Markets Stock Index Fund.
Balanced Funds	
Balanced Index Fund Inv	Seeks current income and long-term capital growth by investing 60% of its assets in stocks tracking the Wilshire 5000 Total Market Index (which represents the entire U.S. stock market) and 40% in bonds tracking the Lehman Aggregate Bond Index (which represents the entire U.S. bond market). This fund maintains a fixed asset mix regardless of market conditions.
Bond Funds	
Short-Term Bond Index Inv	Seeks current income by investing in U.S. government and high-quality corporate bonds with an average maturity of 1 to 5 years. This fund seeks to provide a return consistent with the performance of the U.S. short-term bond market by tracking the Lehman 1–5 Year Government/Credit Index.
Inter-Term Bond Index Inv	Seeks current income by investing in U.S. government and high-quality corporate bonds with an average maturity of 5 to 10 years. This fund seeks to provide a return consistent with the performance of the U.S. intermediate-term bond market by tracking the Lehman 5–10 Year Government/Credit Index.
Total Bond Mkt Index Inv	Seeks current income by investing in a mix of bonds—corporate, government, and mortgage-backed—that represents the total universe of public investment-grade bonds in the U.S. that have maturities over 1 year. This fund, which maintains an average maturity of 5 to 10 years, tracks the Lehman Aggregate Bond Index.
Long-Term Bond Index	Seeks current income by investing in U.S. government and high-quality corporate bonds with an average maturity of 20 to 30 years. This fund seeks to provide a return consistent with the performance of the U.S. long-term bond market by tracking the Lehman Long Government/Credit Index.

Most of these funds are sampled funds rather than fully indexed funds and some have restrictions on withdrawals. While we have used the Vanguard funds to illustrate the diversity of choices when it comes to index funds today, there are other fund families that also have started competing for this very large market. It is striking, however, how often critics of index funds still act as if the only index fund that exists is the S&P 500 index fund.

Are all index funds alike?

You would expect index funds that replicate the same index to be perfect substitutes for each other, but they are not. In other words, differences seem to remain, albeit small, between two index funds that replicate the same index. The first difference is in transactions costs, with some funds reporting much lower transactions costs than others. A perusal of all S&P 500 index funds listed on the Morningstar database in early 2002 indicates that annual returns on these funds lag the index by margins ranging from 10 basis points to 40 basis points; this is a simple measure of the transactions costs for these funds. The cost differences can be traced to economies of scale – larger index funds should be able to spread their fixed costs across more investors – and to execution efficiencies – Vanguard, with its long experience in the index fund business, has much lower execution costs. The second difference is in how closely the funds track the index. You can measure this by looking at the R-squared of the funds. While a perfect index fund should have an R-squared of 100%, most index funds fall short. Some do because they use sampling for larger indices, and the more inefficient their sampling the system, the lower the R-squared will be. Others that try to replicate the entire index still fail because of execution lags and because remittances and withdrawals prevent them from being fully invested in the index. Using the Morningstar database again, the R-squared for S&P 500 index funds range from a low of 96% to a high of 100%. As an investor looking at index funds on a specific index, you want to pick the fund with the lowest expenses and the highest R squared.

The Case for Indexing

The case for indexing ironically is best made by active investors, who try to beat the market and often fail badly. This seems to be true not just for individual investors but also for the professional money managers. In this section, we consider the depressing evidence that notwithstanding the numerous inefficiencies that academics and practitioners claim to find in markets, usually in paper portfolios, converting these paper profits to real profits seems to be very difficult to do.

Individual Investors

There are thousands of individual investors who attempt every day to pick stocks that they believe will do better than the market. The systems they use for stock picking run the gamut from the naïve to the sophisticated. Some base their stock picks on tips from friends – insider trading with six degrees of separation – whereas others use rigorous quantitative analysis. Aided by easier access to data, more powerful personal computers and online (and cheaper) trading, individual investors have narrowed the gap between themselves and those on Wall Street. But is this a good thing? Are investors earning higher returns than they used to?

These questions have been partially answered by researchers who gained access to the brokerage accounts of 78000 clients of a large discount brokerage service. Barber and Odean examined the trading records of individuals who used this brokerage service in a series of papers and came to several interesting conclusions:

- The average individual investor does not beat the market, after netting out trading costs. Between 1991 and 1996, for instance, the annual net (of transactions costs) return on an S&P 500 index fund was 17.8% whereas the average investor trading at the brokerage house had a net return of 16.4%.[1]

- The more individual investors trade, the lower their returns tend to be. In fact, the returns before transactions costs are accounted for are lower for more active traders than they are for less active traders. After transactions costs are accounted for, the returns to active trading get worse.
- Pooling the talent and strengths of individual investors into investment clubs does not result in better returns. Barber and Odean examined the performance of 166 randomly selected investment clubs that used the discount brokerage house. Between 1991 and 1996, these investment clubs had a net annual return of 14.1%, underperforming the S&P 500 (17.8%) and individual investors (16.4%).

Professional Money Managers

There are many who would view the evidence in the last section as predictable. After all, individual investors are amateurs without the access to the information and trading resources available to mutual fund and pension fund managers, i.e., professional money managers. Professional money managers are supposed to be better informed, smarter, have lower transactions costs and be better investors overall than individual investors. In fact, based on this belief, you trust money managers with your savings, pay large fees for money manager expertise and tolerate large transactions costs and taxes. In return, you would expect actively run funds to do better than index funds that mimic the market.

Does the average mutual fund beat the market?

Until the 1960s, the conventional wisdom that professional money managers did much better than individual investors was widely accepted but not really tested, partly because the data was not available and partly because the tools for testing the proposition were not developed. The development of the capital asset pricing model, in conjunction with access to data and statistical packages, allowed Michael Jensen to conduct one of the first studies of mutual funds in 1968. He examined the returns earned by mutual funds from 1955 to 1964 and compared them to what you would have expected them to earn, given their risk exposure. The expected returns for each fund were calculated using the beta that estimated for the fund and the capital asset pricing model. In fact, the difference between the actual return and the expected return from the CAPM is still called Jensen's alpha, reflecting the influence this study had on empirical finance for the next three decades. His

---

[1] See Barber, B.M. and T. Odean, Too many cooks spoil the profits, Financial Analysts Journal, January/February 2000.

findings, summarized in Figure 13.3, as excess returns on mutual funds, were that the average portfolio manager underperformed the market between 1955 and 1964.

Figure 13.3: Mutual Fund Performance: 1955-64 - The Jensen Study

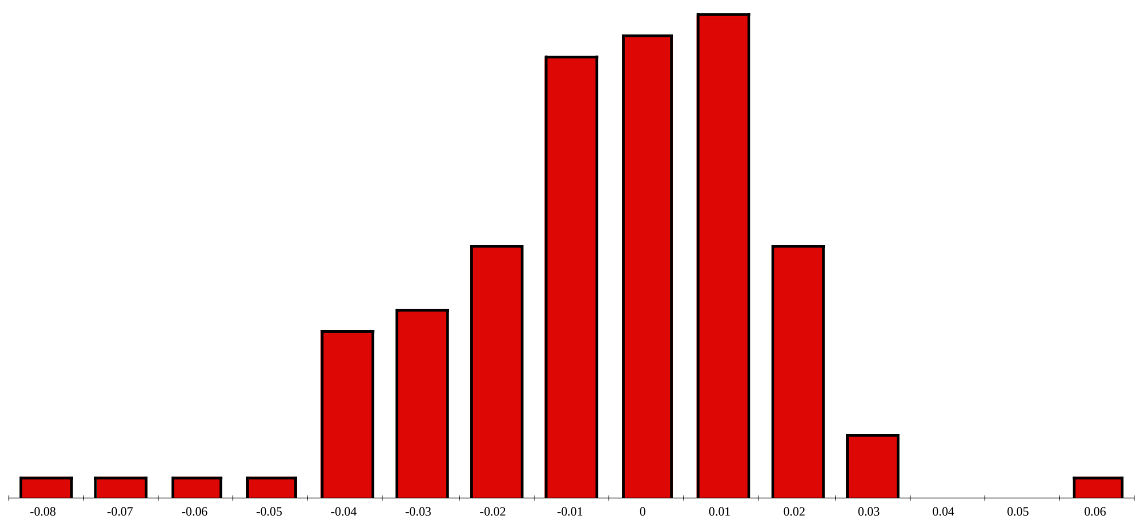

Source: Jensen

Note that an excess return less than 0 indicates that a mutual fund underperformed the market. By this measure, more than 60% of the mutual funds underperformed the market and the average fund delivered returns that were about 1-2% less than expected. These results when published created a controversy that has never quite died out. On the one hand, you have many academics and some practitioners arguing that there is no value added by active money managers. On the other, active money managers have come up with a number of problems they have with the Jensen study, and argue that fixing these problems will reveal to the world the extent of the excess returns they generate. Malkiel updated this study to look at mutual funds from 1972 to 1991. His conclusion is slightly more positive, in the sense that he finds that prior to management expenses, the average mutual fund matches the market. After expenses, though, the average mutual fund underperforms the market by about 1%.[2]

Does active money management work better in the bond markets? Figure 13.4 compares the performance of actively managed bond fund to the Lehman Index (a widely used index of bonds) between 1983 and 1992.

---

[2] This is based upon using the Wilshire 5000 as the index. The underperformance is much greater (about 3.20%) when the S&P 500 is used as the index.

Figure 13.4: Actively Managed Bond Funds versus Bond Index

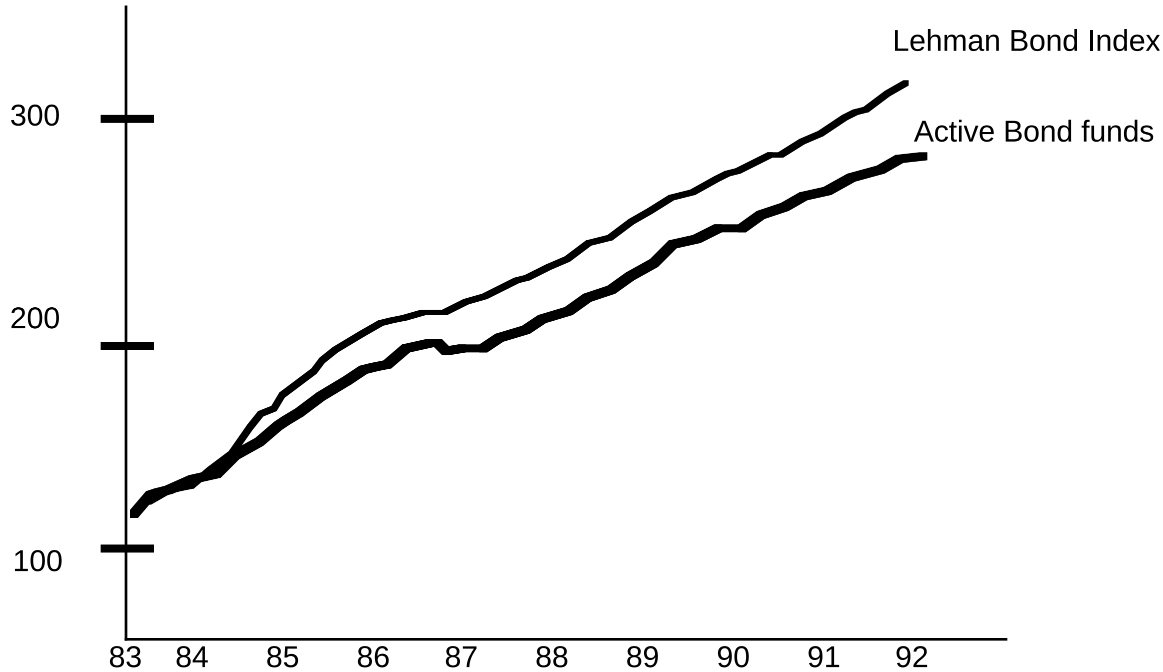

Source: Bogle

Over this 10-year period, the average actively managed bond fund under performed the bond index by about 1.5% a year - $ 100 invested in the bond index would have grown to $ 303 by the end of the period whereas it would have an ending value of $263 if managed by the actively run bond funds. This result has been replicated by other studies.

Measurement Issues

The Jensen study was revolutionary for introducing the notion of risk-adjusted returns to empirical research but it was also limited on a number of dimensions. First, its focus on betas and the capital asset pricing model allowed money managers to claim that the results reflected failures in the model rather than in their performance. Second, the sample used in the Jensen study was a fairly small one and it did not factor in returns from funds that may have been in existence in 1955 but may have failed before 1964. This problem is referred to as survivor bias. Note that this is likely to make the results more negative rather than less negative since it is the poorest performing funds that usually fail. In this section, we will look at studies that have tried to use alternative measures of risk and control for survivor bias.

Excess Return Measurement

In the last two decades, evidence has steadily accumulated on the limitations of the capital asset pricing model and its failure to explain the returns of stocks, especially small cap and low PE ratio stocks. This has provided ammunition for money managers who have wanted to debunk the evidence in Jensen's study. However, there seems to be little evidence that using alternative measures of excess returns generate results that are more positive for money managers. Let us consider a few:

- *Comparison to the Market*: Since any risk and return model is likely to come under assault for one weakness or another, you could revert back to a much simpler comparison of returns on mutual funds to a broad index such as the S&P 500. John Bogle, who pioneered passive investing when he led Vanguard, noted in his book on investing that 166 of 205 mutual funds under performed the Wilshire 5000 between 1983 and 1992. In figure 13.5, we report on the percentage of active mutual funds that are outperformed by the S&P 500 each year from 1971 to 2000. Note that more than 50% of the active funds beat the S&P 500 in only seven out of the 30 years.

Figure 13.5: Percent of Money Managers who beat the S&P 500

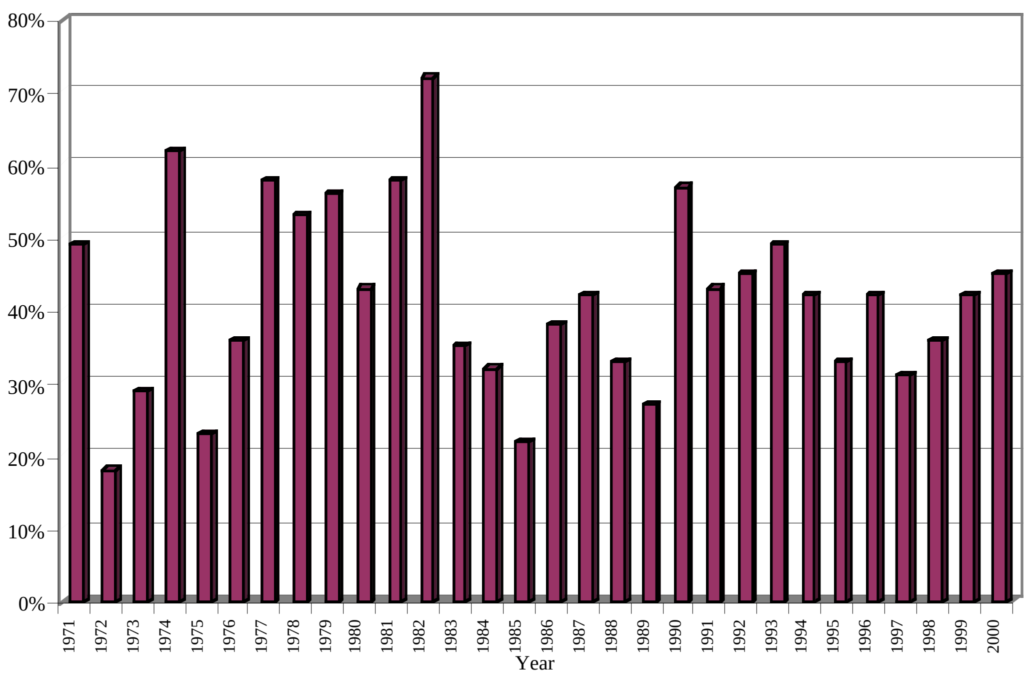

Defenders of mutual fund managers will undoubtedly argue that this measure is biased against mutual funds because they hold cash. Since many of these funds choose to hold cash as a market-timing device, we feel little sympathy for that point of view.

- *Other Risk Measures:* The Sharpe ratio, which is computed by dividing the excess return on a portfolio by its standard deviation, the Treynor measure, which divides the excess return by the beta and the appraisal ratio which divides the alpha from the regression by the standard deviation can be considered close relatives of Jensen's alpha. Studies using all three of these alternative measures conclude that mutual funds continue to under perform the market.[3] In a study that examined the sensitivity of the conclusion to alternative risk and return models, Lehmann and Modest computed the abnormal return earned by mutual funds using the arbitrage pricing model for 130 mutual funds from 1969 to 1982. While the magnitude of the abnormal returns earned is sensitive to alternative specifications of the model, every specification of the model yields negative abnormal returns.
- *Expanded Proxy Models:* In chapter 2, we referenced the study by Fama and French that found that small capitalization stocks with low price to book value ratios earned much higher returns than you would have predicted with the capital asset pricing model. In chapter 7, we also noted the evidence that has accumulated on price momentum – stocks that have done well recently are more likely to continue to do well, at least in the short term. If you do not control for these well known empirical irregularities, you are likely to find negative excess returns in mutual funds that invest in large capitalization stocks with high price to book ratios and negative recent stock price performance. In 1997, Carhart used a four-factor model, including beta, market capitalization, price to book ratios and price momentum as factors, and concluded that the average mutual fund still under performed the market by about 1.80% a year. In other words, you cannot blame empirical irregularities for the under performance of mutual funds.

Given this evidence, it seems safe to conclude that the poor performance attributed to mutual funds cannot be blamed on researchers using the wrong benchmarks for comparison.

Survivor Bias

One of the limitations of many studies of mutual funds is that they use only mutual funds that have data available for a sample period and are in existence at the end of the sample period. Many databases of mutual funds, such as Morningstar, report only on live funds and remove funds that cease operations. Since the funds that fail are likely to be the poorest performers, there is likely to be a bias introduced in the returns that we compute for

---

[3] Sharpe used the ratio in 1966 to evaluate 34 mutual funds to conclude that they under performed the market. Treynor used his index to come to same conclusion a few years later.

funds. In particular, we are likely to over estimate the returns earned by mutual funds by focusing only on the survivors. While this topic has been studied by numerous researchers, the most comprehensive study of survivor bias is by Carhart who examined all equity mutual funds (including failed funds) from January 1962 to December 1995. Over that period, approximately 3.6% of the funds in existence failed each year and they tend to be smaller and riskier than the average fund in the sample. In addition, and this is important for the survivor bias issue, about 80% of the non-surviving funds under perform other mutual funds in the 5 years preceding their failure. Ignoring them as many studies do when computing the average annual return from holding mutual funds results in annual returns being overstated by 0.17% with a one-year sample period to more than 1% with 20-year time horizons. In practical terms, this would mean that if we found mutual funds to have underperformed the market by 1% a year over the last 20 years and we ignored the failed funds, the real underperformance would be closer to 2% a year.

Performance by sub-categories

It may not surprise some that the average fund does not do much better than the market. After all, you could argue that when institutional investors account for 60-65% of the overall market, it will be difficult for them to beat the market. In fact, Charles Ellis has a provocative treatise on money management that is titled “The Loser’s game” that makes exactly this point and should be required reading for anyone embarking on a money management career. You may still believe that there are subsets of funds or superior fund managers that consistently beat the market. In this section, we consider a number of subsets of funds and examine their performance, relative to the market.

By Style

Mutual funds vary when it comes to investment styles and objectives. Some funds label themselves as growth funds and invest in stocks with high expected growth rates and PE ratios. Others are value funds, specializing in stocks trading at low multiples of earnings and book value. There are also yield funds (that concentrate on stocks paying high dividends), diversified funds and small cap funds. The managers of funds in each of these style classes will probably claim that their group outperforms passive investors and it is investors in the other groups that are responsible for the overall under performance. In a paper examining the money management industry, Lakonishok, Shleifer and Vishny classified pension funds into growth, value, yield and other and examined the annual return and percentage of money managers in each group that beat the S&P 500 between 1983 and 1990. Their results are summarized in figure 13.6:

Figure 13.6: Performance by Fund Style - 1983-1990

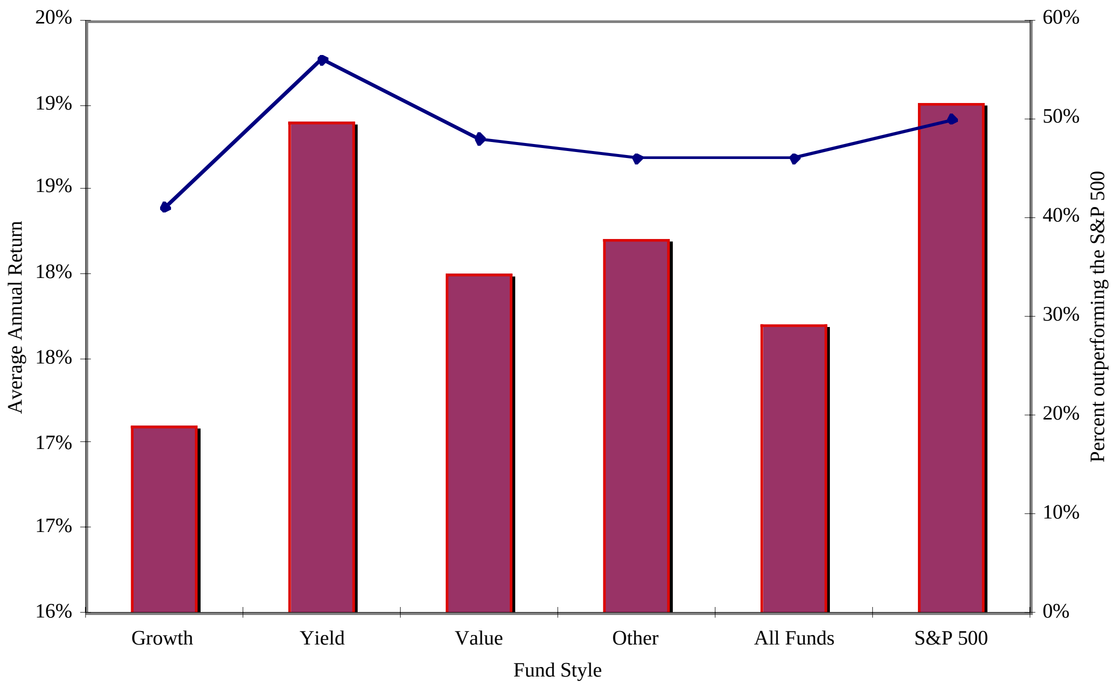

For every style class, other than yield, more than 50% of the managers underperformed the S&P 500. In addition, the returns on the S&P 500 exceeded the annual returns earned by funds in every class.

Growth and value fund investors may take issue with this study because of the comparison to the S& P 500, arguing instead that the comparison should be to a growth and a value index respectively. While this does seem self-serving (since both groups present themselves to investors as the better overall investment), growth funds emerge looking better from this comparison. The average value fund investor underperforms a value index by about 1.2% more than an average growth fund underperforms a growth index.[4] Figure 13.7, drawn from a book by Bernstein on investment styles, presents the comparison of growth (value) funds to a growth (value) index between 1987 and 1993:

---

[4] See Chan, Chan and Lakonishok (1999).

Figure 13.7: Returns on Growth and Value Funds

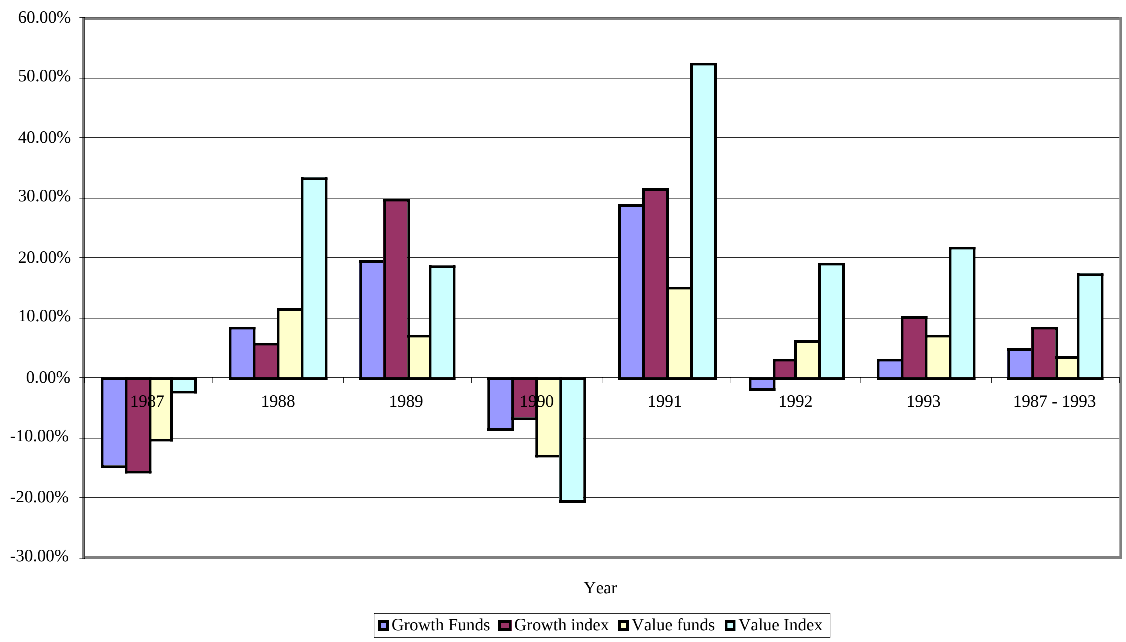

While this comparison was made using only a small sample of value and growth funds, it adds some basis to the notion that growth investors, on average, may have an easier time beating their passive counterparts.

A related issue is whether small cap mutual funds outperform large cap mutual funds. On a raw return basis, the answer seems to be yes. On an excess return basis, the answer depends upon how you adjust for risk. Table 13.2 summarizes the excess returns (using the capital asset pricing model) for small cap, mid cap and large cap funds, categorized by investment style from 1997 to 2001.

Table 13.2: Excess Returns by Market Capitalization and Style

	Value	*Blend*	*Growth*
Large Cap	-0.80%	-2.92%	-4.00%
Mid Cap	-0.56%	-1.17%	-1.84%
Small Cap	0.32%	-0.56%	-1.65%

Small cap funds do better than large cap funds, and value funds outperform growth funds, but the results should not be surprising since the capital asset pricing model consistently yields lower expected returns for small cap and low PE stocks. When a three factor or four factor model that corrects for this bias is used to compute excess returns, these differences either narrow or disappear.

One of the limitations of categorizing mutual funds based upon style is that they often make investments that are at variance with their purported style. Thus, you often find value funds that buy growth stocks and growth funds investing in mature value companies.

Emerging Market and International Funds

While active investing may not have much of a payoff in a mature market with wide access to information like the United States, intuitively you would expect the payoff to be much larger in emerging markets, where information is still not widely disseminated, or even in some European markets, where information tends to be tightly controlled by companies. You would therefore expect active mutual funds in these markets to do much better than they do in the United States, relative to passive indices. Ahmed, Gangopadhyay and Nanda examined 172 emerging market funds listed on Morningstar between 1980 to 2000 and computed the excess returns for these funds. Figure 13.8 summarizes their results:

Figure 13.8: Emerging Market Funds versus Indices

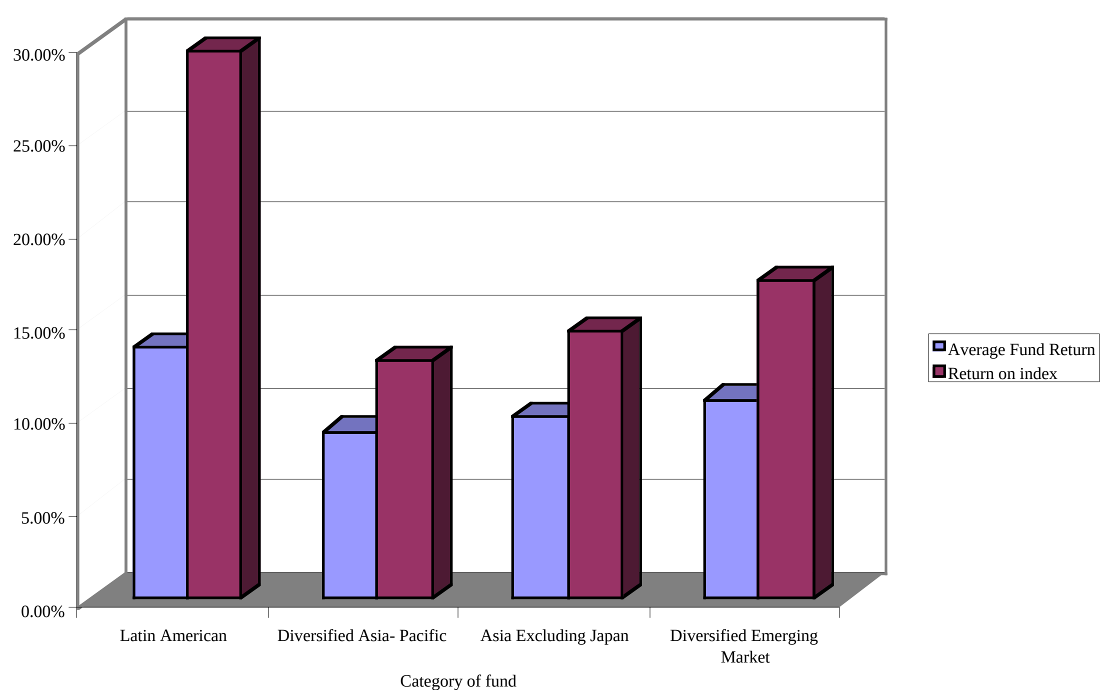

In each of the groupings, the actively managed funds underperformed the index. These results mirror those found in earlier studies of emerging market funds and suggest that active money management does not necessarily pay off in terms of excess returns, even in markets where money managers have information advantages. While this may seem surprising, transactions costs are also higher in these markets and whatever is gained by picking better stocks may very well be lost in trading costs.

What about funds in other developed markets? Actively managed Japanese funds under perform the index by even more than their U.S, counterparts. A study in 1997 by Cai, Chan and Yamada concluded that the average rate of return on 800 actively managed Japanese mutual funds between 1981 and 1992 was only 1.74% a year while the Japanese equity market increased by 9.28% a year during that same period.[5] In one of few bright spots for active money management, Otten and Bams examine 508 actively managed European funds in 2000 and find some evidence of excess returns, especially in small-cap funds.

Other Categorizations

There are a number of other ways in which mutual funds can be categorized and while we will not dedicate entire sections to each, we will summarize the results below:

<u>Load versus No-load funds</u>: Some mutual funds charge an up-front fee, usually a percent of the money invested in the fund, from investors. These fees are called loads and can range from 2 to 5% of the investment. These funds justify these up-front costs by arguing that they will deliver much higher returns than funds that do not charge these fees. Again, the evidence does not back up these claims. Morey (1998) compares the performance of load and no-load funds, both before and after the adjustment of the loads. Using a sample of 301 load and 334 no-load funds from 1993, he tracks performance in the next 5 years, incorporating the effects of funds that cease to exist. Figure 13.9 summarizes his findings -

[5] While this is a truly mind-boggling difference, it should be noted that the net asset values of Japanese mutual funds are adjusted for tax liabilities. In fact, Brown, Goetzmann, Hiraki, Otsuki and Shirashi (2001) argue that much of these negative excesss returns can be explained by tax effects.

Figure 13.9: Jensen's Alpha: Load versus No-load Funds

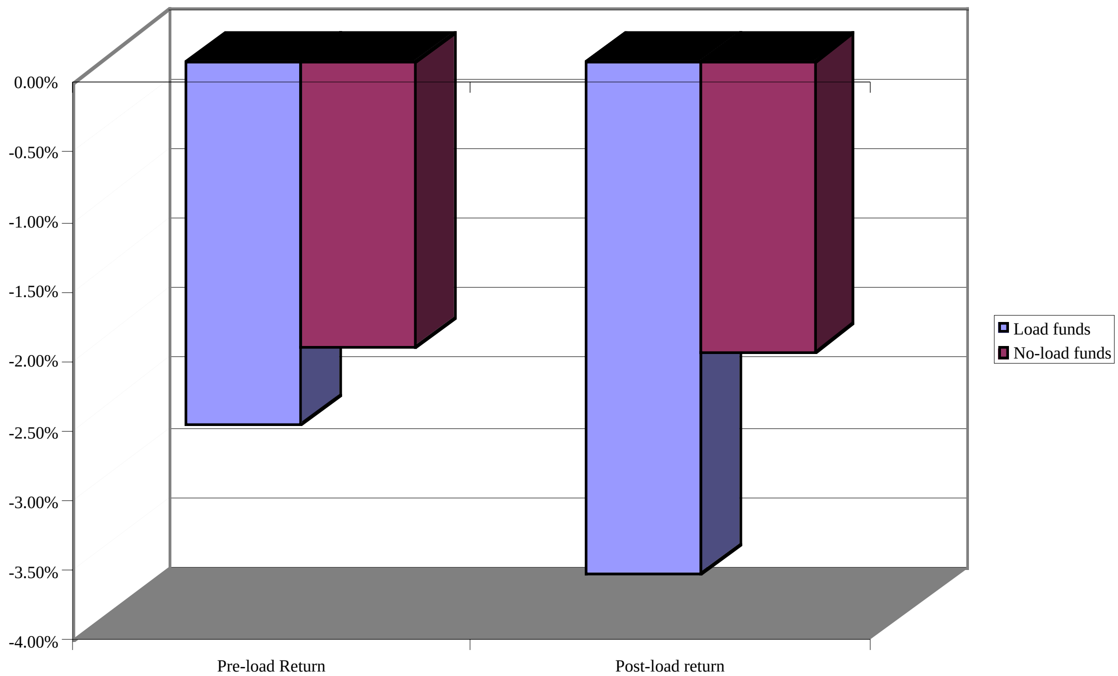

The results are clearly not favorable to load funds. Not only do they fall short of no-load funds, when we consider the load-adjusted returns, but they fall short even when we look at pre-load returns. Given this evidence, it is perplexing that the proportion of funds that are no-load has been dropping over the last decade.

<u>Age and Size of fund</u>: Are funds that have been around for longer (more seasoned funds) better or worse investments than newer funds? Morey, in his study of load and no-load funds, attempted to answer this question as well by categorizing funds into seasoned (more than 10 years old), middle aged (5 to 10 years) and young funds (less than 5 years) in 1993 and examining returns over the subsequent five year period. Figure 13.10 presents his conclusions-

Figure 13.10: Excess Returns by Fund Age

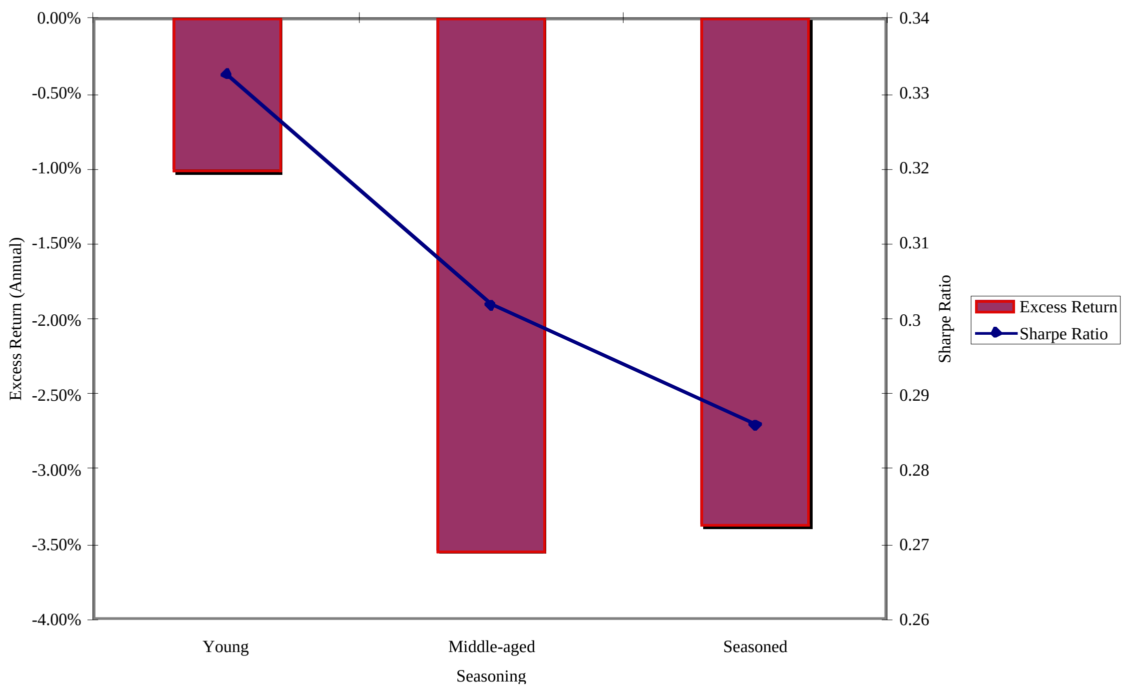

Younger funds seem to do much better than older funds, both in terms of excess returns and delivering higher returns per unit of risk (Sharpe Ratio). When funds are categorized by size, you find similar results, with smaller funds delivering marginally better performance than larger funds, though both lag the indices. Indro, Jiang, Hu and Lee (1999) examined the relationship between fund size and returns by categorizing funds into ten size classes from largest to smallest. While the funds that are in the bottom two deciles (the smallest funds) earn lower returns than other funds, largely because of higher costs, the economies of scale quickly decline and funds that exceed an optimal size (the top 10% of funds in terms of size) also have lower returns.

Fund Manager Characteristics: Does experience make fund managers better? Are older fund managers more likely to deliver high returns than younger fund managers? When funds are categorized based upon the age and experience of their managers, younger managers are more likely to generate positive excess returns than older managers. Younger managers are also more likely to exhibit herd behavior than older managers and to be fired after poor years (which may explain why they exhibit herd behavior in the first place).[6] One

---

[6] See Chevalier and Ellison (1999). They look at funds between 1988 and 1994 and correlate performance to age, SAT scores and status of undergraduate institution. They find at the managers with higher SAT scores

study even looked for differences between male and female money managers and found no significant differences in returns.[7]

Retail versus Institutional Funds: There are some funds that cater exclusively to institutional and very wealthy individuals. They have minimum investment requirements of $100,000 or greater. Some of these funds are stand-alone offerings and some are offered by fund families that have retail mates. In figure 13.11 we report the annual excess returns earned by these funds from 1995 to 1999 categorized by whether they cater to retail or institutional investors, and categorize the latter by minimum investment requirements - big if the minimum investment is greater than $ 500,000, small if the minimum is between $100,000 and $ 500,000- and by whether they are stand-alone or have retail mates.[8]

Figure 13.11: Institutional versus Retail Funds: Annualized Excess Returns

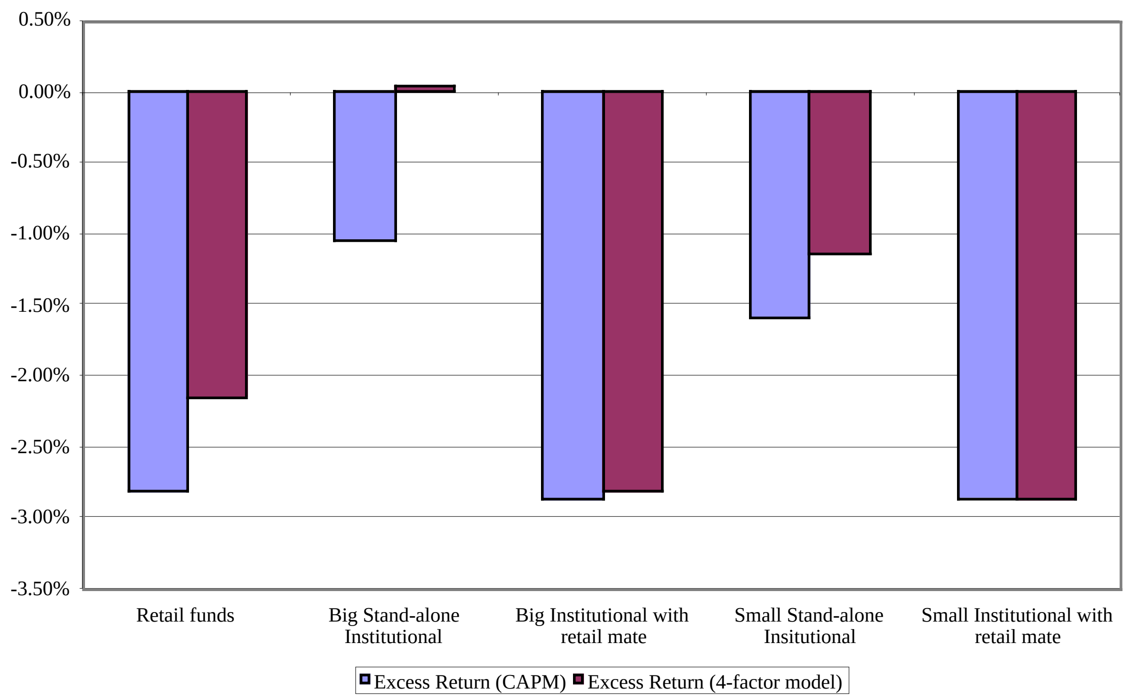

Note that the only funds that marginally beat the market are big institutional funds that have no retail mates.

---

who went to more prestigious undergraduate institutions have slightly more positive returns than other managers.

[7] See Atkinson, Baird and Frye (2001). They also find that net asset flows into funds managed by females are lower than for males, especially for the manager's initial year managing the fund.

[8] See Captured Money? Differences in the Performance Characteristics of Retail and Institutional Mutual Funds, James and Karceski (2002).

<u>Socially responsible funds</u>: In the last decade, a large number of funds have been created to cater to investors who want to avoid companies that they deem socially irresponsible. While the definition of social responsibility varies from fund to fund, the managers of these funds all argue that investing in "ethical" companies will generate higher returns in the long term. Arrayed against them are others who believe that constraining your investment choices will result in lower returns, not higher. In a finding that is bound to leave both groups dissatisfied, Bauer, Koedijk and Otten examined 103 ethical funds in the United States, UK and Germany from 1990 to 2001 and found no significant differences in excess returns between these funds and conventional funds.

Performance Continuity

When confronted with the evidence that the average actively managed fund under performs the market, the reaction of some active money managers is that the average return is brought down by the laggards at the bottom. A profession, they argue, should be judged based upon how those who are best do, rather than by the average. If they are correct, the best money managers should show both consistency and continuity in performance and earn much higher returns than the market.

Transition Likelihood

Perhaps the simplest way to check for continuity is to rank money managers, based upon performance, in one period and then look at the rankings in the next period. Lakonishok, Vishny and Shleifer categorized pension fund money managers from 1983 to 1989 into quartiles, based upon performance each year, and looked at the likelihood of repeat performance. Their results are summarized in table 13.3:

Table 13.3: Continuity of Performance for Pension Funds: 1983- 1990

	Quartile ranking next period			
Quartile ranking this period	1	2	3	4
1	26%	24%	23%	27%
2	20%	26%	29%	25%
3	22%	28%	26%	24%
4	32%	22%	22%	24%

Note that you would have 25% in each box, if performance rankings were completely random - a manager is the first quartile this year will have an equal chance of being in any of the four quartiles next year. The actual percentages are not significantly different from zero, with one exception. A manager who is in the lowest quartile this year has a higher chance of being in the highest quartile next year than in any other quartile. This should not be surprising, since this is exactly what you would expect from mutual funds that take considerable risk and make big bets on a few stocks. If the bets pay off, they move to the top of the rankings and they do not, they drop to the bottom.

Third Party Rankings and Ratings

The rankings in the table above were based entirely on returns and can be faulted for not considering other qualitative factors. There are services like Morningstar that rate mutual funds, and rankings of mutual funds are also provided by the financial news media (the Wall Street Journal, Forbes and Business Week, for example). Do funds that score high on these rankings repeat in future periods? More generally, are these rankings that are often used by investors as the basis for picking funds useful at predicting future performance?

Blake and Morey (2000) examine these questions, using the Morningstar ratings. Morningstar, which maintains one of the most comprehensive databases on mutual funds, assigns ratings ranging from one star (poor) to five stars (outstanding) to funds, based upon both past returns and consistency. The influence of these ratings is illustrated by one study that found that 97% of the money flowed into funds with four or five star ratings.[9] To test whether ratings provide any predictive power, Blake and Morey created a weighted score based upon the 3-year, 5-year and 10-year ratings (with 20%, 30% and 50% weights respectively) for each fund and ranked the funds into ten deciles based upon the weighted score. They the computed the excess returns on funds in each decile between 1994 and 1997 and the results are summarized in figure 13.12:

[9] This statistic was quoted in a Wall Street Journal article by Karen Damato. Karen Damato, "Morningstar Edges Toward One-Year Ratings." Wall Street Journal, April 5th, 1996.

Figure 13.12: Annualized Return based on Morningstar Ratings- 1994-1997

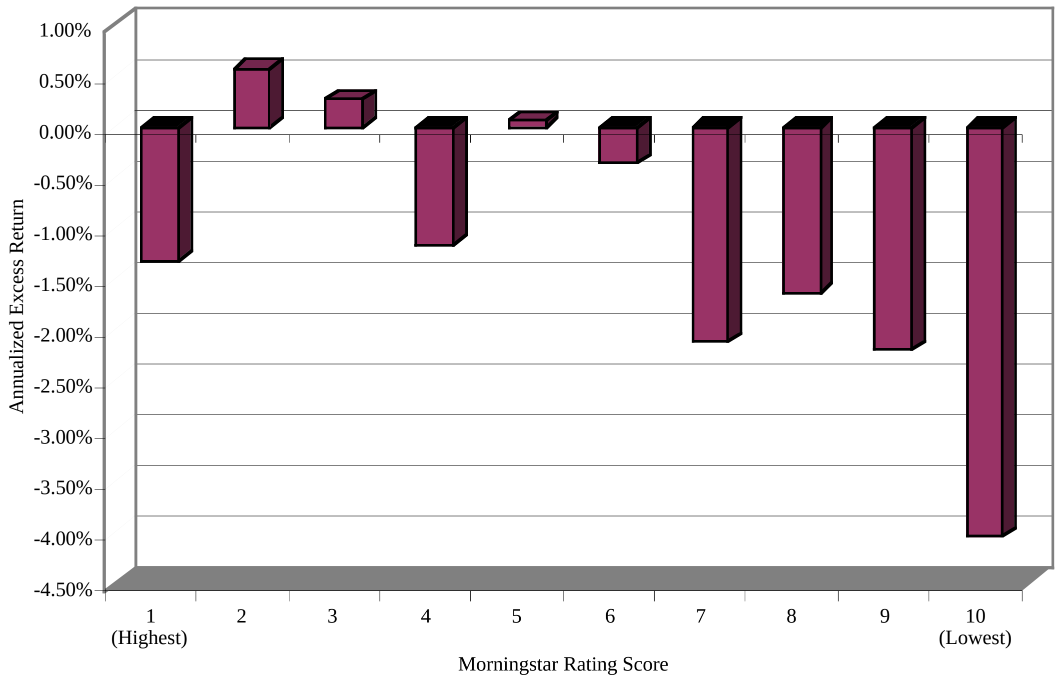

Morningstar ratings seem to have little or no predictive power except for those funds in the lowest ratings. These poorly rated funds tend to do much worse than other funds in the following year. The highest rated funds do worse or no better than the funds in the average ratings.

A study of rankings in the financial magazines of 757 funds between 1993 and 1995 by Detzler finds that while the funds that are ranked highly tend to have high pre-ranking performance (which should be no surprise since the rankings are heavily influenced by recent performance), their performance in the post-ranking period is no different from funds that are ranked lower. About 54% of highly ranked funds under perform the market in the post-ranking period and about 65% of these funds have much poorer post-ranking returns than pre-ranking returns.

The Hot Hands Phenomenon

While much of the evidence that we have presented so far suggests that there is little continuity of performance, there is some contradictory evidence that has accumulated about the very top ranked mutual funds. A number of studies[10] seem to indicate that mutual funds that earn above-average returns in one period will continue to earn above-average returns in

---

[10] See Grinblatt and Titman (1992), Goetzmann and Ibbotson (1994) and Hendricks, Patel and Zeckhauser (1995).

the next period. Malkiel (1995) tests for this "hot hands" phenomenon by looking at the percentage of winners each year who repeat the next year in the 1970s and 1980s. His results are summarized in table 13.4 below -

Table 13.4: Repeat Winners by year – 1971- 1990

Year	*Percent of repeat winners*	*Year*	*Percent of repeat winners*
1971	64.80%	1980	36.50%
1972	50.00%	1081	62.30%
1973	62.60%	1982	56.60%
1974	52.10%	1983	56.10%
1975	74.40%	1984	53.90%
1976	68.40%	1985	59.50%
1977	70.80%	1986	60.40%
1978	69.70%	1987	39.30%
1979	71.80%	1988	41.00%
1971-79	65.10%	1989	59.60%
		1990	49.40%
		1980-90	51.70%

Best performing mutual funds: Take a look at best perfoming mutual funds over the last year.

This table tells a surprising story. The percent of repeat winners clearly is much higher than dictated by chance (50%) in the 1970s. However, the percent of repeat winners during the 1980s looks close to random. Is this because mutual fund rankings became more ubiquitous during the 1980s? Maybe. It is also possible that what you are seeing are the effects of overall market performance. In the 1970s, when the equity markets had a string of negative years, mutual funds that held more cash consistently moved to the top of the rankings. Malkiel also compares the returns you would have earned on a strategy of buying the top funds (looking at the top 10, top 20, top 30 and top 40 funds) from each year and holding it for the next year. The returns are summarized in figure 13.13:

Figure 13.13: Returns on top-ranked funds: 1973- 1991

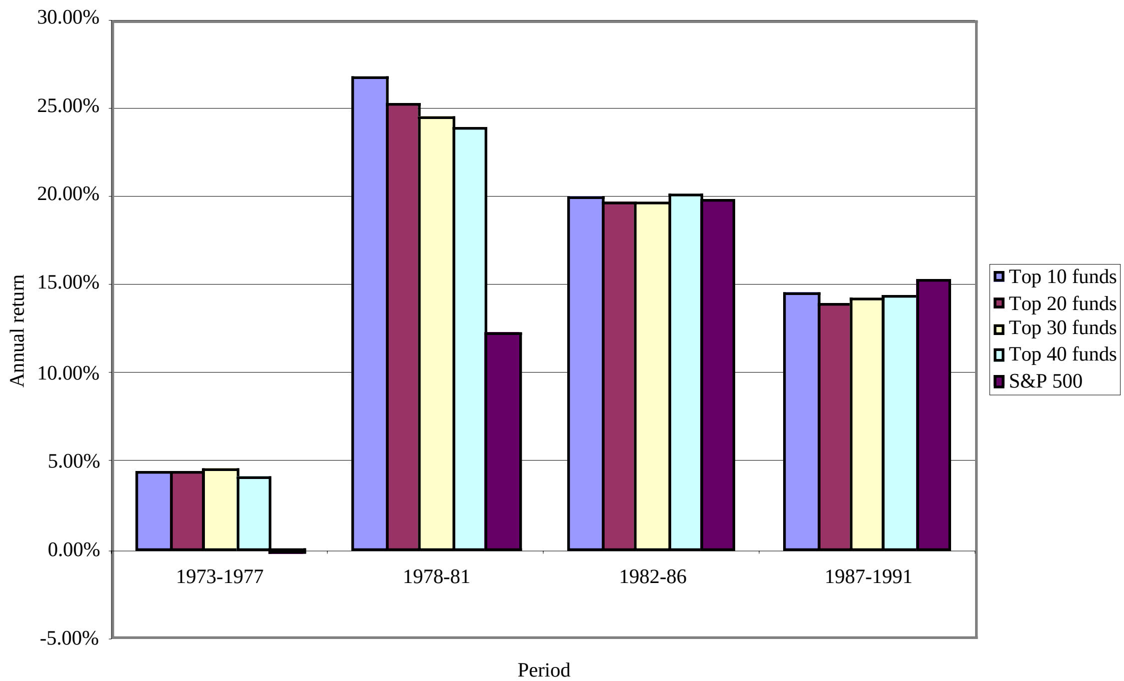

Again, the contrast is striking. While the top funds outperformed the S&P 500 in 1973-77 and 1978-81 time periods, they matched the index from 1982 to 1986 and underperformed the index from 1987 to 1991.

In Summary

Looking at the evidence on mutual fund performance, we are struck by how negative the conclusions are. In general, mutual funds under perform the market and this under performance cannot be explained away by critiquing the risk and return models used by researchers. The under performance is also pervasive and seems to affect funds in every style category. In fact the only mildly positive result is that small cap and high growth funds do less badly against their indices than large cap and value funds do against their indices. If the argument being mounted by active money managers is that we should focus on the more consistent winners, the results reveal little evidence of consistency. Funds that have done well and are rated highly by services may increase their management fees and loads to reflect higher demand, but they do not deliver high returns in the subsequent periods.

Funds and fund managers- Who delivers the excess returns?

Most of the studies of mutual funds that we have reviewed here look at the funds themselves rather the fund managers. But who really is responsible for the excess returns on a fund? Is it the fund itself, because of competitive advantages it has built up over other funds, or is it the fund manager, because of his or her special skills at picking stocks? Consider the question in the context of Fidelity Magellan. Was its success in the early years of its existence due to the fund family – Fidelity – or was it attributable to Peter Lynch's special skills at picking growth companies? The question is important not only for purposes of attribution but also because it may shed some light on the findings in the previous pages. If success at a fund is due to super star managers rather than the fund's own qualities and these managers tend to shift from one fund to another or set up their own funds, you would not expect to see excess returns persisting at funds. You should, however, observe continuity in performance if you track individual money managers.

While it is more difficult to track money managers rather than funds, there are some studies that have attempted to do so. There studies indicate that performance continuity is just as difficult to find with money managers as it is with funds. Though there are a few high profile examples of success (such as Lynch), there are far more examples of managers who have followed successful tenures at one fund with failures at their own funds or at a different fund.

Why do active investors not perform better?

Based upon the evidence in the last section, we cannot avoid drawing the conclusion that active investors under perform the index (and by extension, index funds). In this section, we consider some of the reasons for the under performance of active funds. While much of the evidence comes from looking at mutual funds, you could extend it to cover individual investors.

Transactions Costs

Funds with highest trading costs: Take a look at the 20 funds with the highest and lowest trading costs as % of assets.

The simplest explanation for the difference in returns between actively managed mutual funds and index funds (or passive investing) is transactions cost. Index funds are inexpensive to create - there are no costs to collecting information and no analyst expenses - and inexpensive to run - minimal transactions costs and management fees. For instance, the Vanguard 500 index fund has transactions costs and management fees that amount to 0.17% of the fund. In contrast, the transactions costs and fees at actively managed funds can easily exceed 2%. Figure 13.14 presents the expense ratios in 2001 of actively managed equity funds in the United States.

Figure 13.14: Expense Ratios at Equity Mutual Funds

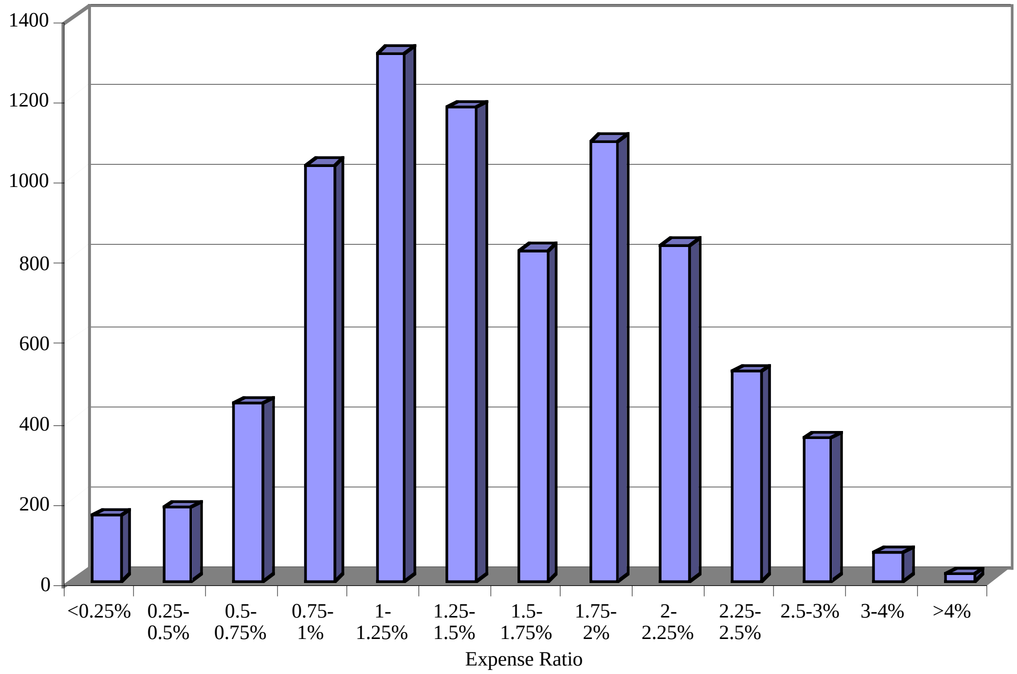

The average expense ratio in 2001 for these equity funds is about 1.78%. Since these are annual, recurring costs, an actively managed fund has to generate an excess return of this amount on its stock picks to cover its expenses.

The key variable determining expenses is turnover. Firms that trade more will usually generate higher transactions costs, but the effect of turnover will be much greater if the fund trades smaller and less liquid companies. In fact, Chalmers, Edelen and Kadlec (1999) find that the relationship between turnover and excess returns is fairly weak, as evidenced in figure 13.15:

Figure 13.15: Turnover Ratios and Returns: Mutual Funds

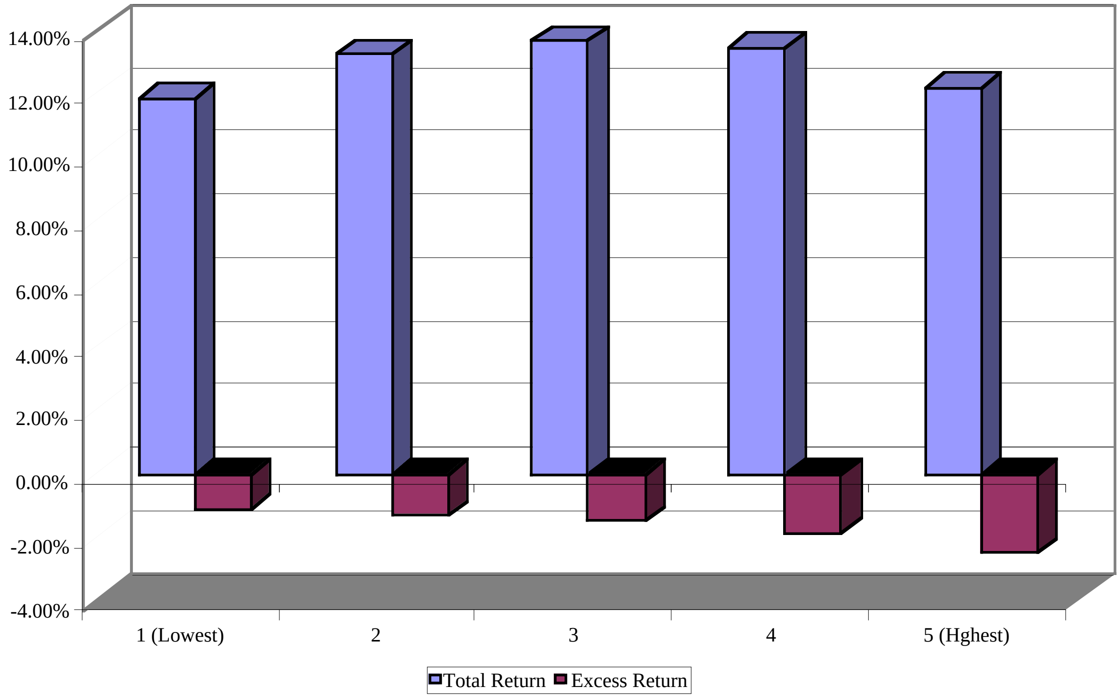

Note that while the funds with low turnover have slightly less negative excess returns than funds with high turnover, they also earn lower total returns. When the study looked at the relationship between total costs (including trading costs and other expenses), the results are much stronger [see figure 13.16]

Figure 13.16: Trading Costs and Returns: Mutual Funds

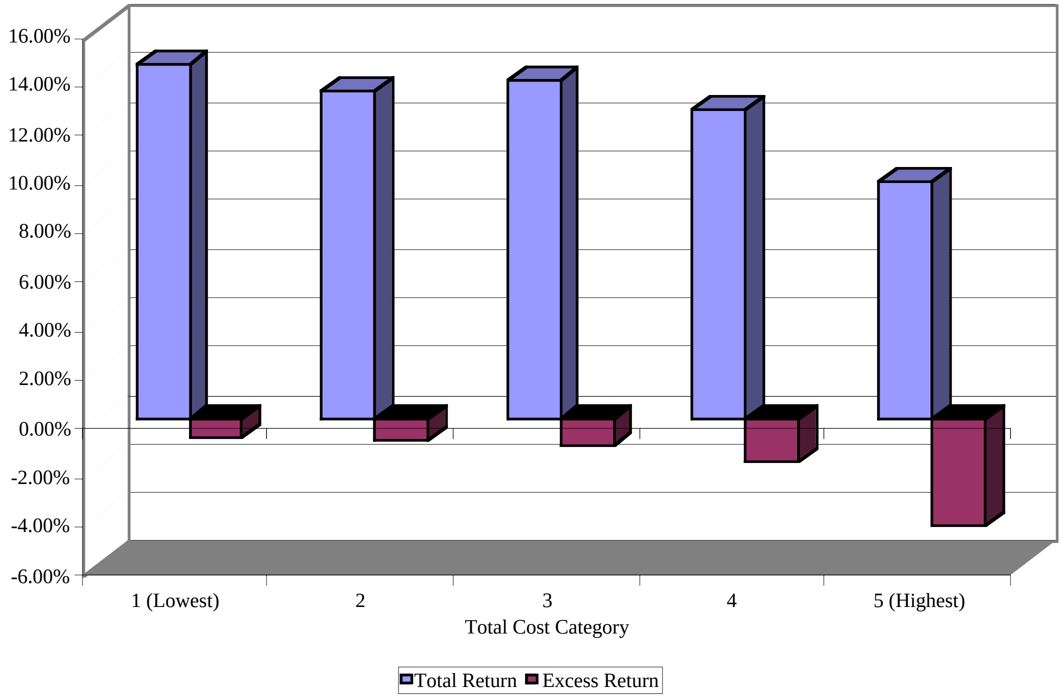

Funds with higher total costs have much lower total returns and much more negative expense ratios than funds with lower total costs.

High Taxes

In chapter 6, we considered the interrelationship between taxes and transactions costs. Mutual funds that trade a lot also create much larger tax bills for their investors and this can best seen by contrasting the pre-tax and after-tax returns at funds. In figure 13.17, we consider the difference between pre-tax and after-tax returns at the five largest index funds and contrast them with pre-tax and after-tax returns at the five largest actively managed funds.

Figure 13.17: Tax Effects at Index and Actively Managed Funds

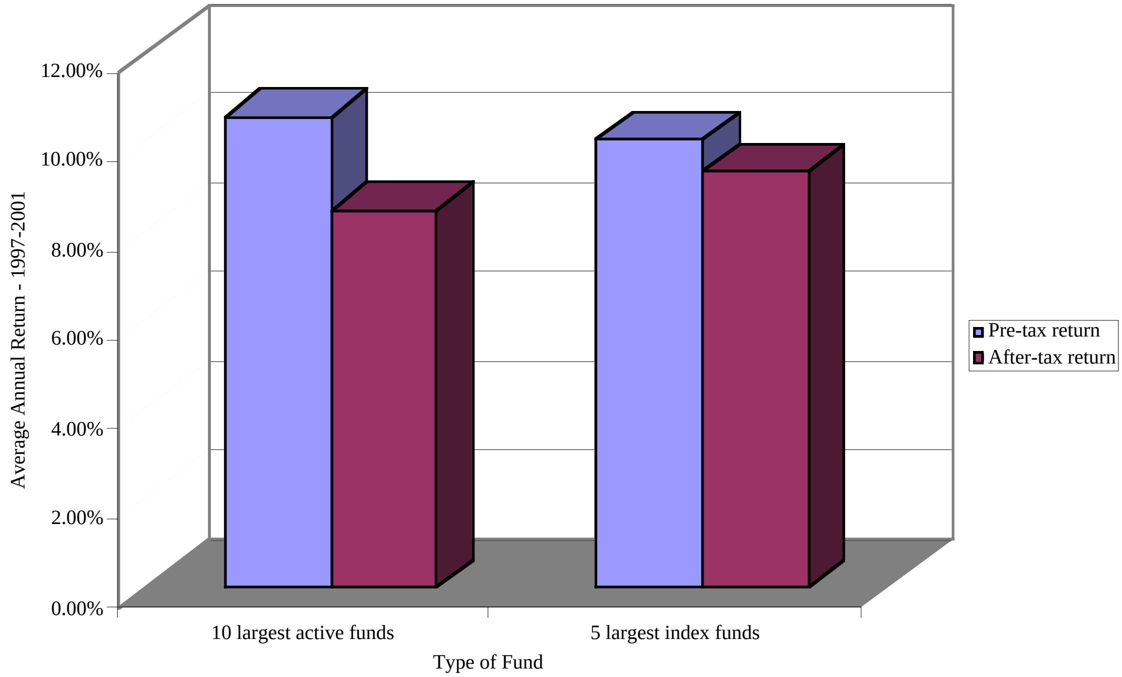

Note that the after-tax return at an active fund is almost 40% lower than the pre-tax return. In contrast, the difference between pre-tax and after-tax returns is much smaller at index funds.

Too much Activity

We invest in actively managed funds because we want them to be actively seeking out under valued stocks. But does this activity pay off? In the study by Lakonishok, Shleifer and Vishny that looked at pension funds, they contrasted the return that funds would have made if they had frozen their portfolios at the beginning of each year with the return that they actually made, after trading through the year. You can consider the difference between the returns as the payoff to active money management and figure 13.18 presents the results for funds in different style classes.

Figure 13.18: Payoff to Active Money Management

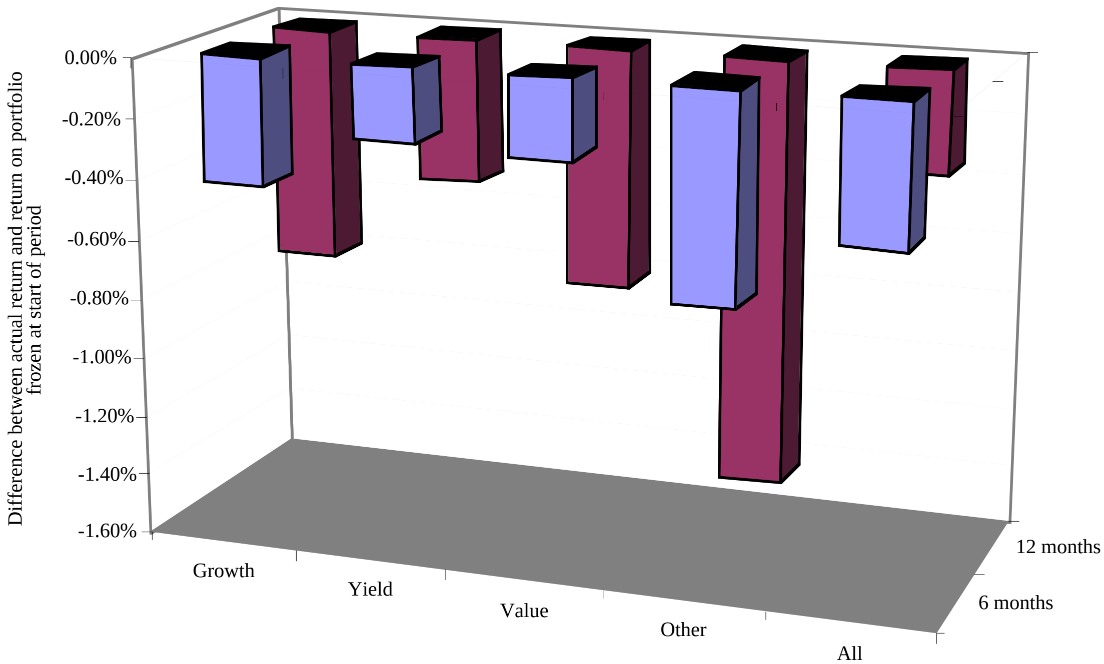

The results indicate that far from adding value, activity reduces returns by between 0.5% (for yield funds) to 1.4% (for other funds). In other words, these funds would have done better if they had sent everyone home at the start of the year and not traded over the course of the year.

Chen, Jegadeesh and Wermers paint a more favorable picture of the stock picking skills of mutual fund managers.[11] While they find no evidence that stocks that are widely held by mutual funds do any better than other stocks, they do find that stocks that are bought by mutual funds earn higher returns in subsequent periods than stocks that are sold. They also conclude that growth oriented funds have better stock selection skills that funds in other categories.

Failure to stay fully invested in equities – Delusions of market timing

Actively managed funds often hold far more cash than they need to meet normal needs and these cash holdings often reflect the market timing views of managers, increasing when managers are bearish and decreasing when they are bullish. As a consequence, mutual funds will tend to under perform the equity index in periods where the equity index increases by more than the riskless rate. Active money managers concede this drain on

---

[11] See "The Value of Active Mutual Fund Management: An Examination of the Stockholdings and Trades of Fund Managers", Journal of Financial and Quantitative Analysis, September 2000.

returns, but argue that their investors are more than compensated by the payoff from market timing. In particular, mutual fund managers contend that they keep investors out of bear markets and thus reduce their downside. In figure 13.19, we compare the performance of actively managed funds and the S&P 500 in six market downturns.

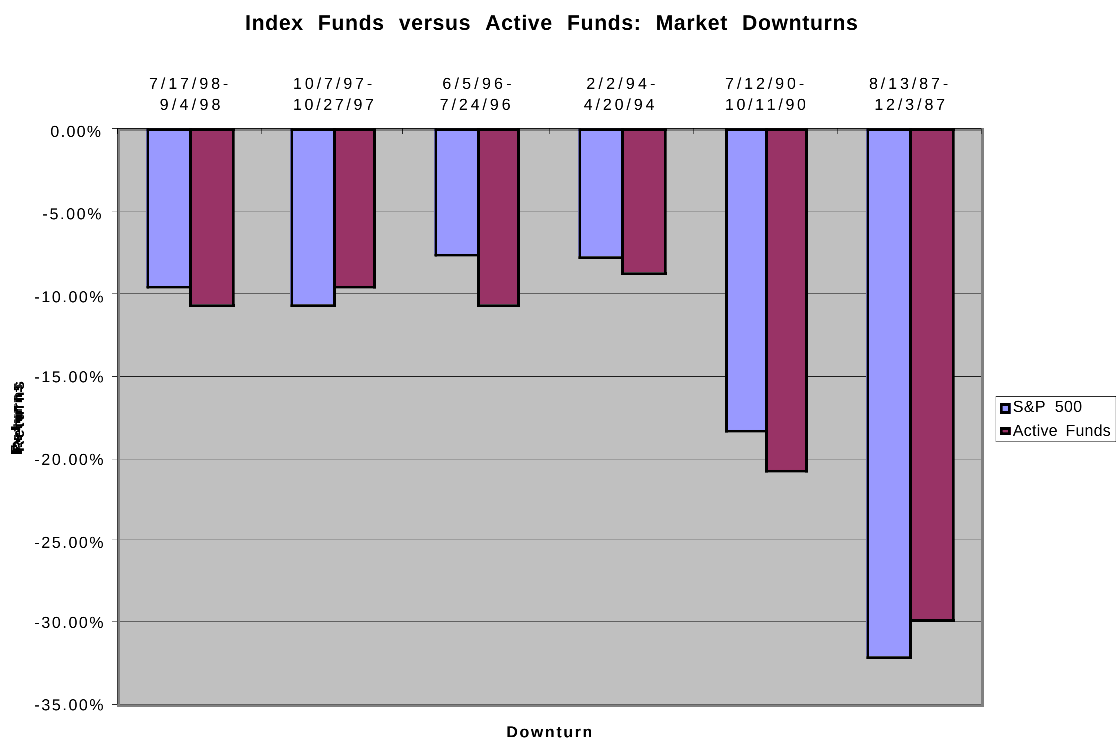

The results don't indicate much market timing ability, since active funds did even worse than the S&P 500 in four of the six downturns. There is another cost, as well. Active money managers often seem to shift into cash during bear markets and they tend to stay in cash too long. Table 13.5 reports on the returns on active funds and the index in three bear markets in the 1970s and 1980s and in the twelve months each bear market. Whatever additional returns may have been earned by active funds during the bear market are effectively wiped out in the 12 months following.

Performance of General Equity Managers During 12 Months Following Bear Markets

	Standard & Poor's 500	Lipper General Equity Average	Index Out-Performance
October, 1974 - September, 1975	38.11%	35.36%	+2.75%
August, 1982 - July, 1983	60.57	66.10	-5.53
December, 1987 - November, 1988	25.40	21.95	+3.45
Cumulative 12 Months Following Bear Markets	178.09%	174.18%	+3.91%
Cumulative Bear Plus Subsequent 12 Month Periods	-8.48%	-8.49%	+.01%

In summary, then, the cash holdings of active money management seem to cost them more in opportunities lost than any potential gain from market timing.

Behavioral Factors

There are three other aspects of mutual fund behavior that seem to contribute to their poor performance. One is the lack of consistency when it comes to investing style/strategy, the second is the tendency to indulge in herd behavior and the third is the practice of making portfolios look better after the fact (window dressing).

- Lack of Consistency: As we noted in an earlier section, funds all too often invest in assets that do not match their stated objectives and philosophy. In fact, explicitly or implicitly, managers switch from one investment style to another. Studies seem to indicate that there is substantial switching of styles from period to period, usually in reaction to market performance in the last period. In a study in 2002, Brown and Van Harlow examined several thousand mutual funds from 1991 to 2000 and categorized them based upon style consistency. They noted that funds that switch styles had much higher expense ratios and much lower returns than funds that maintain more consistent styles.[12]

[12] The style consistency is measured using the R squared of a fund with an index that reflects its stated objective. A fund that has a lower R squared is deviating more from its stated style.

- Herd Behavior: One of the striking aspects of institutional investing is the degree to which institutions tend to buy or sell the same investments at the same time. Thus, you find the institutional holdings in a company drop off dramatically after a poor year and increase in sectors that outperform the market. With emerging markets, you often notice the phenomenon of institutional flight from an emerging market after a severe market decline. Borensztein and Gelos examine emerging market funds in Asia, Latin America, Europe and the Middle East/Africa and report that there is significant herding behavior in each of these regions. There are two negative consequences to funds from herding. The first is that collective selling can make a retreat into a rout, and a small price drop on an investment into a big one. Similarly collective buying can push stock prices up for all buyers. The second is that herd behavior wreaks havoc on investment strategies. A portfolio manager who sells a low price earnings ratio stock, after a price drop, may be under cutting her own long term potential for returns.
- Window Dressing: It is a well documented fact that portfolio managers try to rearrange their portfolios just prior to reporting dates, selling their losers and buying winners (after the fact). This process is called window dressing and it is based upon the premise that investors will look at what is in the portfolio on the reporting date and ignore the actual return earned by the portfolio. Whatever the rationale for window dressing, it creates additional transactions costs for portfolios. O'Neal, in a paper in 2001, presents evidence that window dressing is most prevalent in December and that it does impose a significant cost on mutual funds.[13]

Finding the Right Fund

What are the lessons that individual investors can draw from past studies on mutual funds? These are a few:

1. Pick a fund that best fits your philosophy and needs: Before you pick a fund, develop a view on markets. In other words, choose an investment philosophy first. If you have trouble developing a philosophy, go with an index fund.
2. Do not invest in a load fund: The up-front load creates a mountain that is too tall for even a good portfolio manager to overcome.

[13] He estimates that window dressing costs mutual funds about $ 1 billion each year in transactions costs and price impact.

3. Avoid funds with high turnover ratios: Funds that trade a lot tend to have high trading costs (which eat into your pre-tax returns) and create large tax bills (which reduce your after-tax returns).
4. Avoid funds that do not stay style consistent: A fund manager who keeps shifting styles not only trades more (see item 3 above) but also lacks a core philosophy on investing.
5. Avoid equity funds that have large cash holdings: While funds may claim that they use cash holdings to time markets, there is no evidence that they can. You can make your own decisions on how much cash to hold.
6. Look at the rankings but don't them let them determine your fund choices: Fund rankings may create some short-term momentum, but in the long term, they tell you nothing unless they are abysmal. (If they are abysmal, avoid the fund since it may not be in existence for much longer.)

Alternative Paths to Indexing

If you decide to be a passive investor, there are alternatives to buying an index fund. Exchange traded funds are among the fastest growing and most liquid securities traded in the United States and represent, in the views of their proponents, a much more efficient way of investing in an index. In the last few years, we have also seen a variety of derivatives being created on indices that may allow us to create the equivalent of an index fund at a much lower cost. Finally, we have also seen the development of a class of funds that categorize themselves as enhanced index funds. They claim to provide all of the benefits of indexing – diversification, low transactions costs and tax efficiency – while delivering the excess returns that are a product of active investing.

Exchange Traded Funds

In 1993, the American Stock Exchange began trading depository receipts on the S&P 500 (SPDR, pronounced Spider). As an individual investor, you can buy SPDRs just as you buy other stocks and trade the index with relatively small amounts of money and trivial transactions costs. Not surprisingly, SPDRs have become the most heavily traded instruments on the AMEX. The success of SPDRs has opened the door for a large number of other exchange traded funds – DIAMONDS (indexed to the Dow Jones Industrial Average), Nasdaq-100 shares and iFT-SE (indexed to the FTSE Index in the U.K).

There are many who argue that exchange traded funds are a much more efficient and cheaper way of replicating indices than buying index funds for several reasons:

- They are liquid and can be bought and sold all through the trading day. In addition, you can put restrictions such as limit orders and stop loss rules on the trade, just as you would on an individual stock.
- Unlike index funds, you can sell short on exchange traded funds. This provides you with a way of taking advantage of your market timing skills, and also allows you more flexibility in terms of creating composite positions.
- Index funds sometimes deviate from the index, either because of sampling or because of execution problems. An exchange traded fund always replicates the index.

Does this mean that exchange traded funds will drive index funds out of existence? Not necessarily. Elton, Gruber, Comer and Li take a close look at SPDRs and conclude that there are a few hidden costs. The first is that there is a management fee of about 0.18% charged every year for maintaining the securities and the returns are also reduced by the transactions costs incurred in replicating the index. In other words, SPDRs bear the same costs that an index fund bears in replicating the portfolio and have an additional management fee assessed on top of that. The second cost is that dividends received on the underlying stock cannot be reinvested but have to be held in a non-interest bearing account. As a result of this, they find that SPDRs under perform the S&P 500. Table 13.6 presents the returns on SPDRs and the S&P 500, and measures the shortfall each year from 1993 to 1998:

Table 13.6: SPDRs versus S&P 500

	1993	*1994*	*1995*	*1996*	*1997*	*1998*	*1993-98*
SPDR NAV	8.92%	1.15%	37.20%	22.72%	33.06%	28.28%	21.90%
S & P 500	9.19%	1.32%	37.56%	22.97%	33.40%	28.57%	22.17%
Shortfall	-0.27%	-0.17%	-0.36%	-0.25%	-0.34%	-0.29%	-0.28%

Of the 28 basis point shortfall, they estimate that 18.45 basis points comes from the management fee and the balance is caused by the non-investment of dividends. In contrast, the Vanguard 500 institutional index fund underperformed the index by only 10 basis points a year and the individual index fund underperformed the index by 17 basis points.[14]

There may be tax advantages associated with investing in exchange traded fund. When exchange traded funds are redeemed, the trustee has the option of distributing the securities that comprise the index rather than cash. If the securities distributed on

---

[14] The Vanguard Index fund which is open to individuals has slightly higher costs and underperforms the index by about 17 basis points a year.

redemption are those with substantial capital gains, exchange traded funds may be able to reduce their ultimate tax bills, relative to index funds.

In closing, investors get the advantage of immediate liquidity with exchange traded funds but they do pay a price for the immediacy both in terms of transactions cost (from buying and selling SPDRs) and slightly lower returns. These costs may decrease over time as management fees come down, but for the moment, index funds still have a cost advantage and may be more efficient investments for investors with no need for liquidity and long time horizons.

Index Futures and Options

There are both futures and options traded on the S&P and other equity indices, and they can be used to replicate the index. For instance, a passive investor with $ 10 million to invest can generate the equivalent of an index fund by buying a futures contract on the index and investing the cash in treasury bills. If futures contract are priced at their arbitrage value (as defined in chapter 11), the return on such a strategy should be equal to the return on the S&P 500 less the transactions cost of buying a futures contract. For a large institutional investor, this cost should be very low and the strategy may yield higher returns than investing in an index fund. Derivatives based strategies do not make sense for individual investors, since they tend to have much higher transactions costs. For these investors, index funds will continue to dominate and generate higher returns than alternative strategies.

Enhanced Index Funds

Enhanced index funds claim to have all of the advantages of index funds, while delivering the excess returns associated with actively managed funds. There are no free lunches in investing and enhanced index funds are no exception. In fact, enhanced indexing is, in our view, an oxymoron. You are either be an index fund or an actively managed fund but you cannot be both. Enhanced indexing is really active money management with constraints on activity.

Mechanics

The idea behind enhanced indexing is simple. Staying as close as you can to the index, you try to find pockets of mispricing that allow you to deliver slightly higher returns than the index. Broadly speaking, there are three classes of enhanced indexing strategies.

- In synthetic enhancement strategies, you build on the derivatives strategies that we described in the last section. Using the whole range of derivatives – futures, options and swaps- that may be available at any time on an index, you look for mispricing that you can use to replicate the index and generate additional returns. With a futures-based strategy, Elton Gruber, Comer and Li note that since the implicit spot price [estimated from the arbitrage relationship] in S&P futures contracts is generally slightly lower than the actual spot price (by about 2.7 basis points), a strategy of investing in futures may generate returns that exceed the returns on the index.

- In stock-based enhancement strategies, you adopt a more conventional active strategy using either stock selection or allocation to generate the excess returns. To see how the first would work, consider a strategy of holding all but the 20 most overvalued stocks in the S&P 500. If you could correctly identify these 20 stocks, you would have a portfolio that closely mimicked the index while delivering higher returns. In the second approach, you would hold all the stocks in the index but you would overweight those stocks or sectors that you believed to be under valued and under weight those that you believed to be over valued. How would you identify these? You could use any of the strategies described in earlier chapters, ranging from screening to intrinsic valuation, to come up with the valuations.
- In quantitative enhancement strategies, you use the mean-variance framework that is the foundation of modern portfolio theory to determine the optimal portfolio in terms of the trade-off between risk and return. Thus, if you have the expected return and standard deviation of each stock in the S&P 500 index, and the correlations between every pair of stocks, you can find the portfolio that generates the best trade off between return and risk.[15]

Especially with the second strategy, notice that the only real difference between an enhanced indexing strategy and any other active investment strategy is the constraint that the portfolio you create stay close to the index.

Tracking Error and Information Ratio

Since enhanced index funds claim that they are index funds that deliver extra returns, their risk is measured by looking at how much the returns on these funds deviate, period by period, from the returns on the index. As we noted in chapter 6, the measure used for this is tracking error and is computed by taking the squared deviations of the returns on the fund from the returns on the index. A fund that perfectly tracks the index will have zero tracking error. An enhanced indexed fund will always have a tracking error that exceeds zero, but if it delivers on its promise, it will also have a return greater than the index. The

---

[15] This is not a new idea. Harry Markowitz laid the foundations by capital market theory by explaining portfolio optimization in the 1950s.

performance measure used therefore to evaluate enhanced index funds looks at the ratio of returns in excess of the S&P 500 to the tracking error.

$$\text{Information Ratio} = \frac{\text{(Return on enhanced fund - Returns on Index)}}{\text{Tracking Error}}$$

The nature of the constraint imposed on enhanced index funds becomes clear when we look at the tracking error. Since it is measured as the deviation between the fund's return and the index return in any period, and positive deviations count as much as negative deviations, it operates as a limit on investing outside the index or letting allocations within the fund deviate too much from the index allocations. In other words, if you were a portfolio manager who is judged based upon tracking error, and you had to choose between 2 undervalued stocks, one undervalued by 10% and within the index and one undervalued by 25% and outside the index, you may very well go with the first because the latter will create a much larger tracking error.

Why might investors put this constraint on active money managers and why would managers seek this constraint? The answer lies in our findings on active money management earlier in the chapter. Active money managers tend to deviate significantly from their stated objectives and trade too much, with mostly negative consequences. By imposing a tracking error constraint, investors hope to rein in these excesses and money managers have no choice but to live with this constraint.

Performance

While the promise of enhanced index funds is index fund risk with enhanced returns, it remains an open question as to whether it is delivered. In a study of enhanced index funds in 1998, Riepe examined ten enhanced index funds, two of which used synthetic enhancement strategies and eight of which used stock-based or quantitative strategies from 1991 to 1997. He compared the annual returns on each fund to the returns on the S& P 500 index and his findings are summarized in figure 13.20:

Figure 13.20: Enhanced Index funds vs. S&P 500

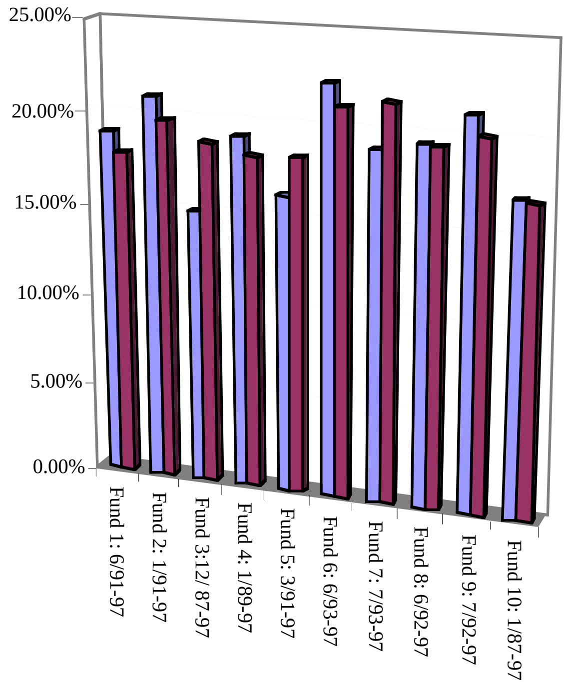

While the names of the funds were not revealed, the period for which each fund had been in existence was provided and is reported next to the fund. Note that three of the ten funds delivered returns lower than the S&P 500, which clearly puts the “enhancement” part of the strategy into question. When the standard deviations in returns on the funds were compared to the S&P 500 (in figure 13.21), seven of the ten funds reported standard deviations that exceeded that of the index, raising doubts about the “indexing” part of the strategy.

Figure 13.21: Enhanced Index funds: Standard deviation vs. S&P 500

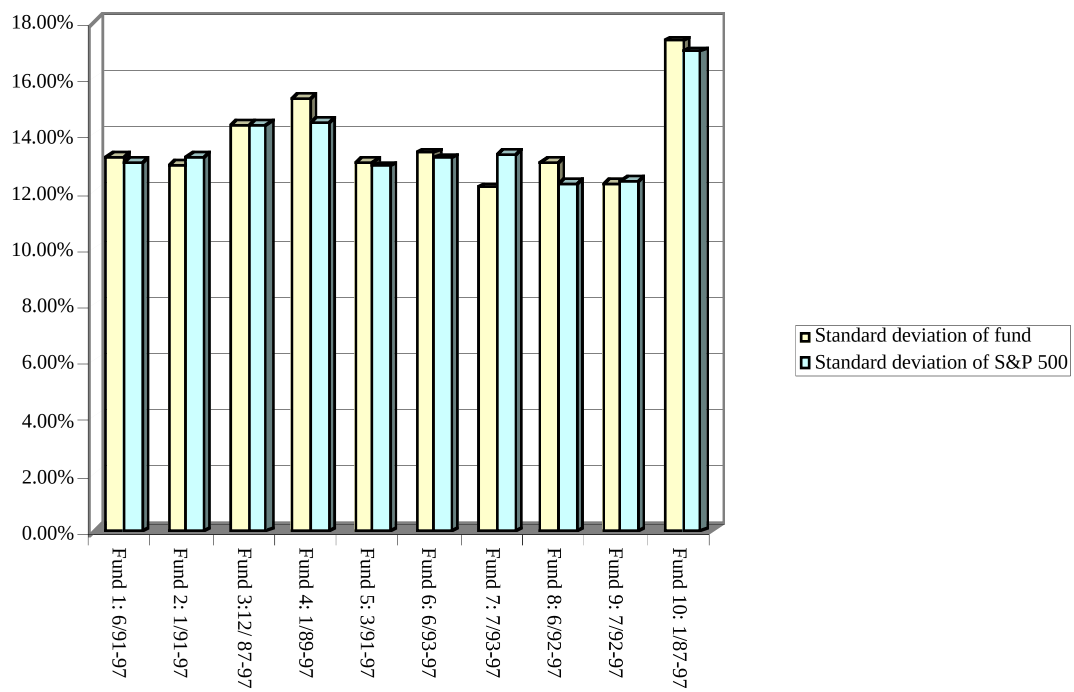

These findings do not surprise us. The only rationale that can be offered for why a more constrained strategy (like enhanced indexing) may outperform a less constrained strategy (like active money management) is that it protects fund managers from their own excesses. The constraints on staying within the index may reduce the style shifting and resulting turnover that is so destructive in conventional mutual funds. If finding bargains in a market is difficult to begin with, handicapping yourself to stay not only within a specified index and constraining yourself to have returns that also track the index returns will make it doubly difficult.

Active/Passive Allocation

We have presented the choice between index funds and actively managed funds as an all or nothing proposition. You either invest all your money actively or all your money passively. There are some investors who may be more comfortable splitting the funds that they have to invest between passive and active management, depending both on their views of the payoffs to active money management and where they are in the investment cycle. Thus, you may allocate 60% of your funds to actively managed funds and 40% to index funds in one year and reverse the allocations in the next.

Another way you can combine passive and active investing is to use specialized index funds to exploit your market timing skills. If you have a good track record on forecasting which style of investing – value or growth – will dominate in the next period, you may buy switch from value index funds to growth index funds to exploit this skill. Similarly, you can use sector-based index funds to take advantage of sector rotation in the market.

Conclusion

In the preceding six chapters, we presented substantial evidence of irregularities in market behavior, related to systematic factors such as size, price-earnings ratios and price book value ratios. We noted that there are individual investment strategies that seem to make excess returns on paper and that there may even be arbitrage opportunities. If you can trade on information in a timely fashion, you may be able to augment these returns. You can consider this chapter to be both a cautionary note and a reality check.

While there may be pockets of inefficiency in the market and substantial evidence that markets make mistakes, there is also the sobering evidence that professional money managers, who are in a position to best exploit these inefficiencies, have a very difficult time consistently beating financial markets. Read together, the persistence of the irregularities and the inability of money managers to beat the market is testimony to the gap between tests on paper portfolios and real world money management. It should also remind us of the fact that investors and portfolio managers are human and succumb to some very human frailties – hubris, insecurity and herd behavior, to name a few. The performance of active money managers provides the best evidence yet that indexing may be the best strategy for many investors.

Lessons for Investors

To invest successfully in actively managed funds, you need to:

1. Understand that the odd are against you: The average active mutual fund under performs the average index fund on a pre-tax basis. On an after-tax basis, actively run funds under perform index funds by even more. The under perfomance shows up in funds in every style class.
2. Recognize why active funds under perform: The average active mutual fund is much too active, trades too much and holds too much cash, all of which drag down returns. You need to find funds that stay consistent in their investment styles and keep their expenses (both management and trading) under control.
3. Not rely too heavily on past performance: While there may be some short-term persistence in fund performance, it is weak and it is very likely to be overwhelmed by the additional fees and expenses that funds add on after a good year.
4. Pick a fund that is close to your investment philosophy: In other words, if your instincts lead you to be a value investor, you should pick a value fund.
5. Keep open the possibility that an index fund may still be your best choice: You should measure up active funds not against each other but against index funds in their category. If they consistently fail to measure up, notwithstanding your best efforts to pick the right funds, you should switch to index funds.

CHAPTER 14

A ROAD MAP TO CHOOSING AN INVESTMENT PHILOSOPHY

If the purpose of this book is to provide you with the tools to pick an investment philosophy, you may very well feel that it has failed. After all, there seems to be both good and bad in every philosophy and no one philosophy seems to dominate over time and yield consistent winners. What purpose has been served you may wonder from this examination of diverse and contradictory views of how markets work and fail to work? In this chapter, we hope to not only wrap up loose ends but also to bring the process of picking a philosophy back to you as an investor.

A Self Assessment

As we noted in chapter 1, there is no one investment philosophy that is best suited for all investors, and much of what we have said in the intervening chapters reinforces this point. A strategy that works for an investor who is patient and has substantial capital to invest may not work for an investor with unpredictable cash needs and a smaller portfolio. In this section, we will consider three aspects that will help determine your investment philosophy – your personal characteristics as an individual, your financial standing and the beliefs you have formed about markets.

Personal Characteristics

Investors who pick investment philosophies that do not fit their personalities are destined to abandon them sooner rather than latter, weighed down not just by the fact that they do not work for them but by personal discomfort with the vagaries of their portfolios. While some of the factors in this section may come perilously close to being on a psychiatrist's couch, you cannot be a successful investor if you do not have a clear eyed view of your own strengths and weaknesses.

a. Patience: Some investment strategies require a great deal of patience, a virtue that many of us lack. Much as you may plead with the powers that be for more patience, you have to accept the reality that you may not suited to an investment strategy that requires ten years of waiting for rewards. If impatient by nature, you should consider adopting an investment philosophy that provides payoffs in the short term.

b. Risk Aversion: Your willingness to bear risk should play a key role in what investment philosophy or strategy you choose for yourself. If you are risk averse, adopting a strategy that entails a great deal of risk – trading on earnings announcements, for instance – will not be a strategy that works for you in the long term.

c. Individual or Group Thinker: Some investment strategies require you to go along with the crowd and some against it. Which one will be better suited for you may well depend upon whether you are more comfortable going along with the conventional wisdom or whether you are a loner. If you are easily subject to peer pressure, odds are high that you will be uncomfortable with contrarian philosophies. If on the other hand, you are comfortable going against the crowd, the fact that most investors are betting against you may bother you little or not at all.

d. Time you are willing to spend on investing: Some investment strategies are much more time and resource intensive than others. Generally, short-term strategies that are based upon pricing patterns or on trading on information are more time and information intensive than long-term buy and hold strategies.

e. Age: If you are an individual investor, your age clearly will make a difference in your choice of investment philosophy. To begin with, as you age, you may find that your willingness to take risk, especially with your retirement savings, decreases. Investment philosophies that you found attractive when you were younger may no longer be attractive or appropriate vehicles for you. With age, they say, also comes wisdom, though we are not sure that this adage applies to investing. It is true, though, that even as a successful investor, you will have learnt lessons from prior investment experiences that will both constrain and guide your choice of investment philosophy.

In summary, your choice of investment philosophy is only partially under your control. Even if you are a patient investor who is willing to go against the crowd, you may find that as you age and become more risk averse, your philosophy and the strategies that go with it have to be modified.

Signs of a misfit

1. You lie awake at night thinking about your portfolio. Investors who choose investment strategies that expose them to more risk than they are comfortable taking will find themselves facing this plight. It is true that your expected returns will be lower with low risk strategies, but the cost of taking on too much risk is even greater.
2. Day to day movements in your portfolio lead to reassessments of your future: While long term movements of your portfolio should affect your plans on when you will retire and what you will do with your future, day-to-day movements should not. It is common

in every market downturn to read about older investors, on the verge or retirement, having to put off retiring because of the damage created to their portfolios. While some of them may have no choice when it comes to where they invest, most investors do have the choice of shifting into low-risk investments (bonds) as they approach retirement.

3. Second guessing your investment decisions: If you find yourself second guessing your investment choices every time you read a contrary opinion, you should reconsider your strategy.

Financial Characteristics

Your choice of investment philosophy will also be affected by your financial characteristics – your job security, the funds you have to invest, your cash needs and your tax status. Since these characteristics change over time, you may have to modify your investment choices to reflect these changes.

Job Security and Earning Capacity

One of the interesting characteristics that we see with financial markets is that investors become more risk averse as the economy weakens - you see this in the widening of default spreads on bonds and in the increase in equity risk premiums during recessions. While we can present a macro economic story for why this happens, we suspect that a great deal of what we see reflects personal insecurity. In the midst of a recession, even those with jobs worry more about their investments and demand larger risk premiums for investing in assets. The flight to quality and, at the limit, to riskless investments is exacerbated by natural and financial crises.

Your investment philosophy will also be heavily influenced by what you perceive your earning capacity to be. If you expect to earn a high income that more than covers your expenses, you have far more degrees of freedom when it comes to picking an investment philosophy. If, on the other hand, your income barely covers your expenses or worse still, falls short, your investment portfolio will have to be tailored to meet your cash needs.

How will this affect your choice of investment philosophy? If you are lucky enough to have a high and predictable income, you can adopt an investment strategy that yields little in the short term but has large payoffs in the long term. If the lessons about risk and return that we have drawn in investing apply to human capital, high-income jobs will probably come with less security, and you will have to invest accordingly. Ultimately, your willingness to bear risk and your time horizon will be heavily influenced by both the level and predictability of your earnings.

Investment Funds

Your choices in terms of investment philosophy expand as the funds at your disposal increase. It may be unfair, but if you have a few thousand dollars to invest, you have little choice but to invest in an index fund. If, on the other hand, you have several hundred thousand dollars to invest, most of the investment philosophies in this book become viable. When considering the investment funds at your disposal, you should look at not only your savings but also money that you have accumulated in pension funds, IRAs and insurance savings accounts. While you are sometimes restricted in your investment choices on some of these funds, you have more choice now than you used to and odds are that your choices will continue to increase over time.

Cash Needs

One of the perils we face both as individual investors and portfolio managers is unpredictable demands for cash withdrawals. For individual investors, this may occur as the result of a personal crisis – a sickness that is not covered by health insurance or the unanticipated loss of income. For professional money managers, it arises because clients can change their minds and demand their money back. If this occurs, you may have to liquidate your investments and lose any long-term return potential that you may have in them.

If your cash demands are unpredictable, what can you do? While you may not be able to forecast when cash withdrawals may need to occur, you can still consider the probabilities when you choose your investment philosophy. If you are a salesman and you make the bulk of your income as a commission, you should expect more volatility in your income and a greater likelihood that you will have cash withdrawals. If you are a portfolio manager of a small, technology fund, you should also assume that your investors are much more likely to shift their savings out of your fund, if you have a bad year. In either case, the expected need for cash shortens your time horizon and may ultimately require you to adopt an investment philosophy with a shorter payoff period.

Tax Status

Much as you wish otherwise, you have to pay taxes and it would be imprudent to pick an investment strategy without considering your tax status. Investors who face high taxes on income should choose investment strategies that reduce their tax liabilities or at least defers taxes into the future. What makes the interplay between investment philosophy and taxes complicated is the fact that different portions of the same individual's income can be subject to different tax treatment. Thus, an investor, when deciding what to buy with her pension fund, where income is tax exempt, may adopt a strategy that generates large amounts of current income, but when investing her personal savings, has to be more careful about tax liabilities.

Market Beliefs

This is perhaps the most difficult component for investors to wrestle with for several reasons. The first is that so much of what we believe about markets comes from anecdotal evidence – from friends, relatives and experts in the field. It is to provide a counter balance that we have looked at the prevailing empirical evidence and disagreements among researchers on what works and does not in financial markets. Needless to say, our work is not done since new research continues to be done.

The second problem is that your views about market behavior and the performance of investment strategies will undoubtedly change over time, but all you can do is make your choices based upon what you know today. In fact, while staying consistent to an investment philosophy and core market beliefs may be central to success in investing, it would be foolhardy to stay consistent as the evidence accumulates against the philosophy.

Finding an Investment Philosophy

We have looked at a variety of investment philosophies in the course of this book and provided some evidence on when they work best and when they fail. We have also categorized these strategies based upon a number of different dimensions – time horizon, funds needed for success and beliefs about market behavior. We will begin with a summary of these findings and then look at the matching an investment philosophy to your financial and personal characteristics.

The Choices

Considering again all of the choices in terms of investment philosophy laid out in this book, we can categorize them based upon the three criteria noted above. While some of these categorizations are hazy – a strategy may be more medium term than long term or more opportunistic than contrarian –they are useful nevertheless.

Time Horizon

The time horizon required to succeed at an investment philosophy and the strategies that flow from it runs the gamut. At one extreme are the long-term strategies such as investing in loser stocks (which have gone down the most over the last six months or year). These require you to invest for five or more years and success is by no means guaranteed even then. At the other extreme are strategies where the time horizon is measured in hours or days, which is where we would categorize trading on earnings announcements and pure arbitrage strategies. In the middle lie strategies that need several months to unfold – buying stocks on relative strength is one example – to a few years.

Capital Requirements

The funds that you need to invest to be successful at investing also vary across strategies. Some strategies require very large portfolios and the benefits that flow from them – low transactions costs, large positions in individual companies and diversification. This is true, for instance, with activist value investing and activist growth investing. Other strategies may be feasible even to investors with small portfolios – a style switching strategy where you switch from value to growth mutual funds and vice versa, depending upon your expectations of earnings growth in future periods would be a good example. Finally, there are some strategies where you need to be large enough to be able to have low execution costs and access to debt but not so large that you create a price impact every time you trade. This is the case with several near or pseudo arbitrage strategies.

Market Beliefs

We can categorize all of the investment strategies described in this book into three groups based upon the underlying market beliefs that drive them. The first set of strategies can be categorized as momentum strategies – i.e. they are based upon the assumption that what has happened in the recent past is likely to continue to happen in the future. Here, we would include most of the technical momentum indicators, such as trend lines and relative strength, as well as some passive growth investing strategies that are based upon momentum in earnings growth. The second set of strategies are contrarian strategies, where you assume that there is a tendency for all aspects of firm behavior – earnings growth, stock returns and multiples such as price earnings – to revert back to historical averages over time. Value investing strategies, where you buy stocks whose prices have hit lows or after substantial bad news, is a good example, as are market timing strategies based upon normalized PE and interest rates. The third set of strategies are opportunistic, where you assume that markets make mistakes but that these mistakes can sometimes lead prices to overshoot (which is what contrarians assume) and sometimes to undershoot (which is what momemtum investors assume). Most arbitrage strategies and some technical indicators (such as price patterns and cycles) can be categorized in this group. Note that there is a fourth group here that we have not recognized explicitly. If we assume that markets are efficient – mistakes are random, cut both ways and are unlikely to be uncovered by investors searching for them – the appropriate strategy is indexing.

In table 14.1, we have categorized most of the strategies described in this book, based upon time horizon and market beliefs. We have also highlighted the strategies that we believe are not feasible for small investors.

Table 14.1: Categorizing Investment Philosophies

	Momentum	*Contrarian*	*Opportunisitic*
Short term (days to a few weeks)	• Technical momentum indicators – Buy stocks based upon trend lines and high trading volume. • Information trading: Buying after positive news (earnings and dividend announcements, acquisition announcements)	• Technical contrarian indicators – mutual fund holdings, short interest. These can be for individual stocks or for overall market.	• *Pure arbitrage in derivatives and fixed income markets.* • Tehnical demand indicators – Patterns in prices such as head and shoulders.
Medium term (few months to a couple of years)	• Relative strength: Buy stocks that have gone up in the last few months. • Information trading: Buy small cap stocks with substantial insider buying.	• Market timing, based upon normal PE or normal range of interest rates. • Information trading: Buying after bad news (buying a week after bad earnings reports and holding for a few months)	• Near arbitrage opportunities: Buying discounted closed end funds • *Speculative arbitrage opportunities: Buying paired stocks and merger arbitrage.*
Long Term (several years)	• Passive growth investing: Buying stocks where growth trades at a reasonable price (PEG ratios).	• Passive value investing: Buy stocks with low PE, PBV or PS ratios. • Contrarian value investing: Buying losers or stocks with lots of bad news.	• *Active growth investing: Take stakes in small, growth companies (private equity and venture capital investing)* • *Activist value investing: Buy stocks in poorly managed companies and push for change.*

Italics: Strategies that are not feasible for small investors

The Right Investment Philosophy

Once you have an inventory of your personal needs and preferences, finding an investment philosophy that is most appropriate for you should be a simple exercise, but you have two choices:

- Single Best Strategy: You can choose the one strategy that best suits you. Thus, if you are a long-term investor who believes that markets overreact, you may adopt a passive value investing strategy.
- Combination of strategies: You can adopt a combination of strategies to maximize your returns. For instance, you may mix in a basic long-term strategy of passive growth investing with a medium term strategy of buying stocks with high relative strength. Obviously, you are hoping to augment your returns from the first strategy with your returns on the second. In creating this combined strategy, you should keep in mind the following caveats:
    - You should not mix strategies that make contradictory assumptions about market behavior over the same periods. Thus, a strategy of buying on relative strength would not be compatible with a strategy of buying stocks after very negative earnings announcements. The first strategy is based upon the assumption that markets learn slowly whereas the latter is conditioned on market overreaction.
    - When you mix strategies, you should separate the dominant strategy from the secondary strategies. Thus, if you have to make choices in terms of investments, you know which strategy will dominate.

Review

Investing is a continuous process and we learn (or should learn) all the time from both our successes and failures. We learn about how other investors in the market behave and we learn about ourselves. Our circumstances in terms of job security (or at least the perception of it) and income also change. As a consequence, we may have to revisit the decision of which investment philosophy is best suited repeatedly, and in some cases, change to reflect both what we have learnt in the recent past and our current status. In making these changes, though, we should fight the temptation to go with the strategy that worked best in the recent past or expert advice. Financial experts, for the most part, are no better at forecasting the future than the rest of us.

Conclusion

Choosing an investment philosophy is at the heart of successful investing. To make the choice, though, you need to look within before you look outside. The best strategy for you is one that matches both your personality and your needs. If you are patient by nature, have a secure income stream and little or no need for cash withdrawals, you can choose a strategy that may be risky in the short term but has a good chance of paying off in the long term. If, on the other hand, you cannot wait for extended periods and have immediate cash needs, you may have to settle for a shorter-term strategy.

Your choice of philosophy will also be affected by what you believe about markets and investors and how they work (or do not). You may come to the conclusion that markets overreact to news (in which case you would migrate towards contrarian strategies), markets learn slowly (leading to momentum strategies), markets make mistakes in both directions (yielding opportunistic strategies) or that market mistakes are random. Since your beliefs are likely to be affected by your experiences, they will evolve over time and your investment strategies have to follow suit.

References for Investment Philosophy

Chapter 2

Amihud, Y., B. Christensen and H. Mendelson, 1992*, Further Evidence on the Risk-Return Relationship*, Working Paper, New York University.

Bernstein, P., 1992, *Capital Ideas*, The Free Press, New York.

Bernstein, P., 1996, *Against the Gods*, John Wiley & Sons, New York.

Chan, L.K. and J. Lakonsihok, 1993, *Are the reports of Beta's death premature?,* Journal of Portfolio Management, v19, 51-62.

Chen, N., R. Roll and S.A. Ross, 1986, *Economic Forces and the Stock Market*, Journal of Business, 1986, v59, 383-404.

Elton, E.J. and M.J. Gruber, 1995, *Modern Portfolio Theory and Investment Management*, John Wiley & Sons, New York.

Fama, E.F. and K.R. French, 1992, *The Cross-Section of Expected Returns*, Journal of Finance, v47, 427-466.

Jensen, M.C, 1969, *Risk, the Pricing of Capital Assets, and the Evaluation of Investment Portfolios,* Journal of Business, v42, pp 167-247.

Kothari, S.P. and J. Shanken, 1995, *In Defense of Beta*, Journal of Applied Corporate Finance, v8(1), 53-58.

Lintner, J., 1965, *The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets*, Review of Economics and Statistics, v47, 13-37.

Markowitz, Harry M., *Foundations Of Portfolio Theory*, Journal of Finance, 1991, v46(2), 469-478.

Roll, R., 1977, *A Critique of the Asset Pricing Theory's Tests: Part I: On Past and Potential Testability of Theory*, Journal of Financial Economics, v4, 129-176.

Ross, Stephen A., 1976, *The Arbitrage Theory Of Capital Asset Pricing*, Journal of Economic Theory, v13(3), 341-360.

Sharpe, W.F*,* 1964, *Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk*, Journal of Finance, v19, 425-442.

Weston, J.F. and T.E. Copeland, 1992, *Managerial Finance*, Dryden Press.

Chapter 3

Choi, F.D.S. and R.M. Levich, 1990*, The Capital Market Effects of International Accounting Diversity*, Dow Jones Irwin, New York.

Stickney, C.P., 1993, *Financial Statement Analysis*, Dryden, Fort Worth

White, G.I, A. Sondhi and D. Fried, 1997, *The Analysis and Use of Financial Statements*, Wiley, New York.

Williams, J. R., 1998, *GAAP guide*, Harcourt Brace, New York.

Chapter 4

Black, F. and M. Scholes, 1972, *The Valuation of Option Contracts and a Test of Market Efficiency*, Journal of Finance, Vol 27, 399-417.

Damodaran, A., 2002, *Investment Valuation*, Second Edition, John Wiley & Sons. New York.

Chapter 5
Amihud, Y. and H. Mendelson, 1986, *Asset Pricing and the Bid-Ask Spread*, Journal of Financial Economics, v17, 223-249.
Barclay, M.., 1997, *Bid-Ask Spreads and the Avoidance of Odd-Eighth Quotes on Nasdaq: An Examination of Exchange Listings*, *Journal of Financial Economics*, 45, 35-60
Christie, W., and P. Schultz. 1999, *The Initiation and Withdrawal of Odd-Eighth Quotes among Nasdaq Stocks: An Empirical Analysis*, *Journal of Financial Economics*, 52, 409-442.
Christie, W., and P. Schultz. 1994, *Why Do Nasdaq Market Makers Avoid Odd-Eighth Quotes? Journal of Finance*, 49 , 1813-1840.
Chung, K., B. Van Ness, and R. Van Ness, 2001, *Can the Treatment of Limit Orders Reconcile the Differences in Trading Costs between NYSE and Nasdaq Issues?*, Working Paper, Social Sciences Research Network.
Dann, L., D. Mayers and R. Raab, 1977, *Trading Rules, Large Blocks and the Speed of Adjustment*, v2, 3-22.
DeBondt, W.F.M. & R. Thaler, 1985, *Does the Stock Market Overreact?*, Journal of Finance, v40, pp 793-805.
Dey, M.K.and B. Radhakrishna , 2001, *Institutional Trading, Trading volume and spread*, Working paper, Social Science Research Network.
Hasbrouck, J., 1991a, *Measuring the Information Content of Stock Trades*, *Journal of Finance*, 66, 179-207.
Hasbrouck, J., 1991b, *The Summary Informativeness of Stock Trades: an Econometric Analysis*, Review of Financial Studies, 4, pp. 571-595.
Heflin. Shaw and Wild, 2001, *Disclosure Quality and Market Liquidity*, Working paper, Social Sciences Research Network.
Holthausen, R. W., R. W. Leftwich, and D. Mayers, 1990, *Large-Block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects*, Journal of Financial Economics, v26, 71-95.
Huang, R.D. and H.R. Stoll, 1996, *Dealer versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE*, Journal of Financial Economics, v 41, pg 313-357.
Keim, D. B., and A. Madhavan, 1995, *Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Trades*, Journal of Financial Economics, 37, 371-398.
Keim, D.B. and A. Madhavan , 1998, *The Cost of Institutional Equity Trades*, Financial Analysts Journal, July/August 1998.
Kothare, M. and P.A. Laux , 1995, *Trading Costs and the Trading Systems for NASDAQ stocks*, Financial Analysts Journal, March/April 1995.
Leinweber, D.J., 1996, *Using Information from Trading in Trading and Portfolio Management*, in Execution Techniques, True Trading Costs and Microstructure of Markets, AIMR.
New York Stock Exchange Fact Book, 1996, New York Stock Exchange.
Rose, J.D. and D.C. Cushing , 1996, *Making the Best Use of Trading Alternatives*, in Execution Techniques, True Trading Costs and the Microstructure of Markets, AIMR.
Saar, G., 2001, *Price Impact Asymmetry of Block Trades: An Institutional Trading Explanation*, Review of Financial Studies, v15, 1153-1181.

Silber, W.L., 1991, *Discounts on Restricted Stock: The Impact of Illiquidity on Stock Prices*, Financial Analysts Journal, 60-64.
Spierdijk, L. , T. Nijman, and A.H.O. van Soest (2002), *The Price Impact of Trades in Illiquid Stocks in Periods of High and Low Market Activity*, Working paper, Tilburg University.
Stoll, H. 1978, *The Pricing of Security Dealer Services: An Empirical Study of Nasdaq Stocks*. Journal of Finance, 33, 1153-1172.
Thomas Loeb, 1983, *Trading Costs: The Critical Link Between Investment Information and Results*, Financial Analysts Journal, May/June 1983.
Tinic, S. and R. West , 1972*, Competition and the Pricing of Dealer Service in the Over-the-Counter Market*, Journal of Financial and Quantitative Analysis, v7, pg 1707-27.
Treynor, J., 1981, *What does it take to win the trading game?*, Financial Analysts Journal, January-February.

Chapter 6
Conrad, J., 1989, *The Price Effect of Option Introduction*, Journal of Finance, v44, pp 487-498.
Fama, E.F. and K.R. French, 1992, *The Cross-Section of Expected Returns*, Journal of Finance, v47, 427-466.
Fama, E.F., 1970, *Efficient Capital Markets: A Review of Theory and Empirical Work*, Journal of Finance, v25, pp 383-417.
Fama, E.F., 1972, *Components of Investment Performance*, Journal of Finance, v27, 551-567.
Markowitz, Harry M., *Foundations Of Portfolio Theory*, Journal of Finance, 1991, v46(2), 469-478.
Sharpe, W.F., 1965, *Mutual fund Performance*, Journal of Finance, vv39, 119-138.on Sharpe Ratio
Treynor, J.L., 1965, *How to rate management of mutual funds*, Harvard Business Review, 43, 63-70.

Chapter 7
Abraham, A. and D.L.Ikenberry,1994, *The Individual Investor and the Weekend Effect*, Journal of Financial and Quantitative Analysis, v29, 263-277.
Alexander, S.S., 1964, *Price Movements in Speculative Markets: Trends or Random Walks*, in The Random Character of Stock Market Prices, MIT Press.
Ariel, R.A., 1990, *High Stock Returns before holidays*, Journal of Finance, v45, 1611-1626.
Asch, S., 1952, *Social Psychology*, Prentice Hall, Englewood Cliffs.
Bekaert,, G., C.B. Erb, C.R. Harvey and T.E. Viskanta. 1997, What matters for emerging market equity investments, Emerging Markets Quarterly (Summer 1997), 17-46.
Brown, S.J., W. N. Goetzmann and A. Kumar , 1998, *The Dow Theory: William Peter Hamilton's Track Record Reconsidered*, Working paper, NYU.
Chancellor, E., 2000, *Devil takes the Hindmost*, Plume.
Conrad, J.S., A. Hameed and C. Niden, 1994, *Volume and Autocovariances in Short-Horizon Individual Security Returns*, Journal of Finance, v49, 1305-1330.
Cootner, P. H. 1961, *Common Elements In Futures Markets For Commodities And Bonds*, American Economic Review, 1961, v51(2), 173-183.

Cootner, P.H., 1962, *Stock Prices: Random versus Systematic Changes*, Industrial Management Review, v3, 24-45.
Datar, V., N. Naik and R. Radcliffe, 1998, *Liquidity and Asset Returns: An alternative test*, Journal of Financial Markets, v____
Dooley, M.P. and R. Shafer, 1983, *Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981*, in Exchange Rate and Trade Instability, Causes, Consequences and Remedies, 1983, Ballinger.
Edwards, R. and J. Magee, 2001, *Technical Analysis of Stock Trends*, St. Lucie Press.
Fama, E.F. and K.R. French, 1988, *Permanent and Temporary Components of Stock Prices*, Journal of Political Economy, v96, 246-273.
Fama, E.F., 1965, *The Behavior of Stock Market Prices*, Journal of Business, v38, pp 34-105.
Fama, E.F., *Market Efficiency, Long Term Returns and Behavioral Finance*, Journal of Financial Economics, v 49, pg 283-306.
Fischhhof, B., P. Slovic and S. Lichtenstein, 1977, *Knowing with Uncertainty: The Appropriateness of Extreme Confidence*" Journal of Experimental Psychology, v3, 522-564.
Frost, A.J. and E. Prechter, 1998, *The Elliott Wave Principle: Key to Market Behavior*,, New Classics Library.
Garber, P.M., 1990, *Who put the tulip in tulipmania?* in "Crashes and Panics: The lessons of History", Dow Jones Irwin.
Gehm, F., 1983, Who is R.N. Elliott and why is he making waves? Financial Analysts Journal, January-February, 51-58.
Grinblatt, M., S. Titman and R. Wermers, 1995, *Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior*, American Economic Review, v85, 1088-1105.
Hamilton, W., 1922, *The Stock Market Barometer*, Barrons, New York.
Haugen, R. and J. Lakonishok, 1988, *The Incredible January Effect*, Dow-Jones Irwin, Homewood, Ill.
Jegadeesh, N. and S. Titman, 1993, *Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,* Journal-of-Finance; 48(1), 65-91.
Jegadeesh, N. and S. Titman, 2001, *Profitability of Momentum Strategies: An Evaluation of Alternative Explanations,* Journal-of-Finance; 56(2), 699-720.
Jennergren, L.P., and P.E. Korsvold, 1974, *Price Formation in the Norwegian and Swedish Stock Markets - Some Random Walk Tests*, Swedish Journal of Economics, 76, 171-185.
Kahnemann, D. and A. Tversky, 1979, *Prospect Theory: An analysis of decisions under risk, Econometrica*, v46, 171-185.
Kho, B.C., *Time Varying Risk Premia, Volatility and Technical Trading Rules*, Journal of Financial Economics, v41, 246-290.
Lee, C.M.C and B. Swaminathan, 1998, *Price Momentum and Trading Volume*, Working Paper, Social Science Research Network.
Liss, D., 2001,*A Conspiracy of Paper*, Ballantine Books.
Lo, Wang and Mamaysky, 2000*, Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation*, Journal of Finance, v55, 1705-1765.

Mackay, C., 1852, *Extraordinary Popular Delusions and the Madness of Crowds*, Reprinted by John Wiley & Sons, New York.
Odean, T., 1997, *Are investors reluctant to realize their losses?* Working paper, University of California, Davis.
Osler,C.L. and P.H.K. Chang , *Head and Shoulders: Not a flaky pattern*, Staff Paper, 1995, Federal Reserve Bank of New York.
Robert Schiller, 2000, *Irrational Exuberance*, Princeton Press, Princeton.
Rouwenhorst, G.K., 1998, *International Momentum Strategies*, Journal of Finance, v53, 267-284.
Santoni and Dwyer 1990, *Bubbles or Fundamentals: New Evidence from the Great Bull Markets*, in "Crashes and Panics: The lessons of History", Dow Jones Irwin.
Senchack, A.J. and L.T. Starks, 1993, *Short-Sale Restrictions and Market Reaction to Short-Interest Announcements*, v28, 177-194.
Shafir, E., I. Simonson and A. Tversky, 1997, *Money Illusion*, Quarterly Journal of Economics, v112, 341-374.
Shefrin, H. and M. Statman, 1985, *The disposition to sell winners too early and ride losers too long: Theory and Evidence*, Journal of Finance, v40, p777-790.
Shiller, R., 1990, Market Volatility, MIT Press.
Stickel and Verecchia, 1994, *Evidence that trading volume sustains stock price changes*, Financial Analysts Journal, Nov-Dec, 57-67.

Chapter 8

Berger, Philip G., and Eli Ofek, 1995, *Diversification's effect on firm value*, Journal of Financial Economics 37, 39–65.
Bhide, A., 1989, *The Causes and Consequences of Hostile Takeovers*, Journal of Applied Corporate Finance, v2, 36-59.
Bradley, M., A. Desai and E.H. Kim, 1983*, The Rationale behind Interfirm Tender Offers*, Journal of Financial Economics, v11, 183-206.
Bradley, M., A. Desai and E.H. Kim, 1988, *Synergistic Gains from Corporate Acquisitions and their Division between the Stockholders of Target and Acquiring Firms*, Journal of Financial Economics, v21, 3-40.
Bradley, M., G.A. Jarrell, and E.H. Kim, 1984, On the Existence of an Optimal Capital Structure: Theory and Evidence, Journal of Finance, v39, 857-878.
Capaul, C., I. Rowley and W.F. Sharpe, 1993, *International Value and Growth Stock Returns*, Financial Analysts Journal, 27-36.
Caton, G.L., J. Goh and J. Donaldson, 2001, *The Effectiveness of Institutional Activism*, Financial Analysts Journal, July/August, 21-26.
Chan, L.K., Y. Hamao, and J. Lakonishok, 1991, *Fundamentals and Stock Returns in Japan*, Journal of Finance. v46. 1739-1789.
Clayman, M., 1994, *Excellence revisited*, Financial Analysts Journal, May/June 1994, pg 61-66.
Cusatis, P.J, J.A. Miles and J.R. Woolridge, 1993, *Restructuring Through Spin Offs: The Stock Market Evidence*, Journal of Financial Economics, v33, 293-311.
Damodaran, 2001, *Investment Valuation (Second Edition)*, John Wiley & Sons, New York

DeBondt , W.F.M. & R. Thaler, 1987, *Further Evidence on Investor Overreaction and Stock Market Seasonality*, Journal of Finance, v42, pp 557-581.
DeBondt, W.F.M. & R. Thaler, 1985, *Does the Stock Market Overreact?*, Journal of Finance, v40, pp 793-805.
Furtado, E.P.H., and V. Karan, 1990. *Causes, consequences, and shareholder wealth effects of management turnover: A review of the empirical evidence*. Financial Management 19, 60-75.
Graham, B. and D. Dodd, 1934, *Security Analysis*. McGraw Hill.
Hagstrom, R.G., 1994, *The Warren Buffett Way*, John Wiley and Sons, New York.
Hirschey, M., 2000, *The "Dogs of the Dow" Myth*, Financial Review, v35, 1-15.
Jacobs, B.I. and K.N. Levy, 1988a, *Disentangling Equity Return Irregularities: New Insights and Investment Opportunities*, Financial Analysts Journal, Vol 44, 18-44.
Jacobs, B.I. and K.N. Levy, 1988b, *On the Value of 'Value*, Financial Analysts Journal, Vol 44, 47-62.
Jegadeesh, N. and S. Titman, 1993, *Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,* Journal-of-Finance; 48(1), 65-91.
Jegadeesh, N. and S. Titman, 2001, *Profitability of Momentum Strategies: An Evaluation of Alternative Explanations,* Journal-of-Finance; 56(2), 699-720.
Karpoff, J.M., 1998, *The Impact of Shareholder Activism on Target Companies: A Survery of Empirical Findings*, Working Paper, University of Washington.
Karpoff, J.M., P.H. Malatesta and P.A. Walkling, 1996, *Corporate Governance and Shareholder Initiatives: Empirical Evidence*, Journal of Financial Economics, v42, 365-395.
Klein, A., 1986, *The Timing and Substance of Divestiture Announcements: Individual, Simultaneous and Cumulative Effects*, Journal of Finance, v41, 685-696.
Lang, L.H., R.M. Stulz and R.A. Walkling, 1989, *Managerial Performance, Tobin's Q, and the Gains from Successful Tender Offers*, Journal of Finance, v24, 137-154.
Lang, Larry H.P., and René M. Stulz, 1994, *Tobin's q, corporate diversification, and firm performance*, Journal of Political Economy 102, 1248–1280.
Linn, Scott C. and Michael S. Rozeff. 1985, *The Effect Of Voluntary Spin-Offs On Stock Prices: The Anergy Hypothesi*s, Advances in Financial Planning and Forecasting, v1(1), 265-292.
Lowe, J.C., 1994, *Benjamin Graham on Value Investing: Lessons from the Dean of Wall Street*, Dearborn Financial, Chicago.
Lowenstein, R., 1996, *Buffett: The Making of an American Capitalist*, Doubleday.
Miles, J. and J.R. Woolridge., 1999, *Spin-Offs & Equity Carve-Outs,* Financial Executives Research Foundatio*n*.
McQueen, G., K. Shields and S.R. Thorley, 1997, Does the Dow-10 Investment Strategy beat the Dow statistically and economically? Financial Analysts Journal, July/August, 66-72.
Modigliani, F. and M. Miller, 1958, *The Cost of Capital, Corporation Finance and the Theory of Investment*, American Economic Review, v48, 261-297.
Oppenheimer, H R. 1984, *A Test of Ben Graham's Stock Selection Criteria*, (September/October): vol. 40, no. 5 , 68-74.
Peters, T.. 1988, *In Search of Excellence: Lessons form America's Best Run Companies*, Warner Books.

Rosenberg, B., K. Reid, and R. Lanstein, 1985, *Persuasive Evidence of Market Inefficiency*, Journal of Portfolio Management, v11, 9-17.
Schipper, K. and A. Smith, 1983, *Effects of Recontracting on Shareholder Wealth: The Case of Voluntary Spin-Offs*, Journal of Financial Economics, Vol 12, 437-468.
Schipper, K. and A. Smith, 1986, *A Comparison of Equity Carve-Outs and Seasoned Equity Offerings: Share Price Effects and Corporate Restructuring*, Journal of Financial
Hite, G.L. and J.E. Owers, 1983, *Security Price Reactions around Corporate Spin-off Announcements*, Journal of Financial Economics, Vol 12, 409-436.
Senchack, A.J., Jr., and J.D. Martin, 1987,*The Relative Performance of the PSR and PER Investment Strategies*, Financial Analysts Journal, Vol 43, 46-56.
Smith, C.W., 1986, *Investment Banking and the Capital Acquisition Process*, Journal of Financial Economics, v15, 3-29.
Wahal, S., 1996, *Pension Fund Activism and Firm Performance*, Journal of Financial and Quantitative Analysis, v31, 1-24.
Zarowin, P., 1990, *Size, Seasonality and Stock Market Overreaction*, Journal of Financial and Quantitative Analysis, v25, 113-125.

Chapter 9

Banz, R., 1981, *The Relationship between Return and Market Value of Common Stocks*, Journal of Financial Economics, v9.
Basu, S., 1977, *The Investment Performance of Common Stocks in Relation to their Price-Earnings: A Test of the Efficient Market Hypothesis*, Journal of Finance, v32, 663-682.
Basu, S., 1983, *The Relationship between Earnings Yield, Market Value and Return for NYSE Common Stocks: Further Evidence*, Journal of Financial Economics, v12.
Bernstein, 1995, *Style Investing*, John Wiley & Sons.
Bernstein, R., 1995, *Style Investing*, John Wiley and Sons.
Black, B.S. and R.J. Gilson, 1998, *Venture Capital and the Structure of Capital Markets: Banks versus Stock Markets*, Journal of Financial Markets, v47, 243-277.
Chan, L.K., Y. Hamao, and J. Lakonishok, 1991, *Fundamentals and Stock Returns in Japan*, Journal of Finance. v46. 1739-1789.
Damodaran, 2001, *Investment Valuation (Second Edition)*, John Wiley & Sons, New York.
Dimson, E. and P.R. Marsh, 1986, *Event Studies and the Size Effect: The Case of UK Press Recommendations*, Journal of Financial Economics, v17, 113-142.
Dreman, D. and E. Lufkin, 1997, *Do contrarian strategies work within industries?* Journal of Investing, (Fall), 7-29.
Dreman, D. and E. Lufkin, 2000, *Investor Overreaction: Evidence that its basis is psychological*, Journal of Psychology and Financial Markets, v1.
Fama, E.F. and K.R. French, 1998, *Value versus Growth: The International Evidence*, Journal of Finance, v53, 1975-1999.
Haugen, R.A. and Lakonishok, J., 1988, *The Incredible January Effect*, Homewood Ill,, Dow Jones-Irwin.
Lee, I., S. Lockhead, J.R. Ritter and Q. Zhao, 1996, *The Costs of Raising Capital*, Journal of Financial Research, v19, 59-74.
Little, I.M.D., 1962, *Higgledy Piggledy Growth*, Institute of Statistics, Oxford.

Loughran, T. and J.R. Ritter, 1995, *The New Issues Puzzle*, Journal of Finance, v50, 23-51.
Lynch, P., 1997, *How to invest a million*, Worth Magazine, March issue.
Malkiel, B.G., 1995, *Returns from Investing in Equity Mutual Funds 1971 to 1991*, Journal of Finance, v50, 549-572.
Pradhuman, S., 2000, *Small Cap Dynamics*, Bloomberg Press.
Ritter, J.R., 1998, *Initial Public Offerings*, Contemporary Finance Digest, v2, 5-31.
Siegel, J., 1998, *Stocks for the Long Run*, McGraw Hill, New York.

Chapter 10

Aharony, J. and I. Swary, 1981, *Quarterly Dividends and Earnings Announcements and Stockholders' Returns: An Empirical Analysis*, Journal of Finance, Vol 36, 1-12.
Bettis, J., Vickrey, D., and Donn Vickrey, 1997, *Mimickers of Corporate Insiders Who Make Large Volume Trades*, Financial Analyst Journal 53, 57-66.
Bettis, J.C., J.M. Bizjak and M.L. Lemmon, 2002, *Insider Trading in Derivative Securities: An Empirical Investigation of Zero cost collars and Equity Swaps by Corporate Insiders*, Working Paper, Social Sciences Reseach Network.
Bhide, A., 1989, *The Causes and Consequences of Hostile Takeovers*, Journal of Applied Corporate Finance, v2, 36-59.
Bhide, A., 1993, *Reversing Corporate Diversification*, in The New Corporate Finance- Where Theory meets Practice, ed. D.H. Chew Jr., McGraw Hill.
Bradley, M., A. Desai and E.H. Kim, 1983*, The Rationale behind Interfirm Tender Offers*, Journal of Financial Economics, v11, 183-206.
Bradley, M., A. Desai and E.H. Kim, 1988, *Synergistic Gains from Corporate Acquisitions and their Division between the Stockholders of Target and Acquiring Firms*, Journal of Financial Economics, v21, 3-40.
Brown, K.C., W.V. Harlow and S.M. Tinic, 1988, *Risk Aversion, Uncertain Information, and Market Efficiency*, Journal of Financial Economics, v22, pg 355-385.
Brown, L.D. and M.S. Rozeff, 1980, *Analysts can forecast accurately*!, Journal of Portfolio Management, v6, 31-34.
Capstaff, J. , K.Paudyal and W. Rees, 2000, *Revisions of Earnings Forecasts and Security Returns: Evidence from Three Countries*, Working Paper, SSRN.
Chan, K., L.K.C. Chan, N. Jegadeesh and J. Lakonishok, 2001*, Earnings Quality and Stock Returns*, Working Paper, SSRN.
Charest, G., 1978, *Split Information, Stock Returns and Market Efficiency-I*, Journal of Financial Economics, v6, 265-296.
Chopra, V.K., 1998, *Why so much error in analyst forecasts?* Financial Analysts Journal, Nov-Dec, pg 35-42.
Collins, W. and W. Hopwood, 1980, *A Multivariate Analysis of Annual Earnings Forecasts generated from Quarterly Forecasts of Financial Analysts and Univariate Time Series Models*, Journal of Accounting Research, v18, 390-406.
Cooper, R.A., T.E. Day and C.M. Lewis, 1999, *Following the Leader: A Study of Individual Analysts Earnings Forecasts*, Working Paper, SSRN.
Copeland, T. E. *Liquidity Changes Following Stock Splits*, Journal of Finance, 1979, v34(1), 115-141.
Cragg, J.G., and B.G. Malkiel, 1968, *The Consensus and Accuracy of Predictions of the Growth of Corporate Earnings*, Journal of Finance, v23, 67-84.

Crichfield, T., T. Dyckman and J. Lakonishok, 1978, *An Evaluation of Security Analysts Forecasts*, Accounting Review.
Damodaran, A., 1989, *The Weekend Effect in Information Releases: A Study of Earnings and Dividend Announcements*, Review of Financial Studies, v2, 607-623.
Dennis and McConnell, 1986, *Corporate Mergers and Security Returns*, Journal of Financial Economics, v16, 143-188.
Dreman, D.N. and M. Berry, 1995, *Analyst Forecasting Errors and their Implications for Security Analysis*, Financial Analysts Journal, May/June, pg 30-41
Finnerty, J.E, 1976, *Insiders and Market Efficiency*, Journal of Finance, v31, 1141-1148.
Fuller, R.J., L.C. Huberts and M. Levinson, 1992, *It's not Higgledy-Piggledy Growth!* Journal of Portfolio Management, 38-46.
Givoly. D. and J. Lakonishok, 1984, *The Quality of Analysts' Forecasts of Earnings*, Financial Analysts Journal, v40, 40-47.
Grinblatt, M.S., R.W. Masulis and S. Titman, 1984, *The Valuation Effects of Stock Splits and Stock Dividends*, Journal of Financial Economics, v13, 461-490.
Healy, P.M., K.G. Palepu and R.S. Ruback, 1992, *Does Corporate Performance improve after Mergers?*, Journal of Financial Economics, v31, 135-176.
Higgins, H.N., 1998, Analyst Froecasting Performance in Seven Countries, Financial Analysts Journal, May/June, v54, 58-62.
Huang, R.D. and R. Walkling, 1987, *Acquisition Announcements and Abnormal Returns*. Journal of Financial Economics, v19, 329-350.
Ikenberry, D.L., G. Rankine and E.K., Stice. 1996. What Do Stock Splits Really Signal?, Journal of Financial and Quantitative Analysis, v31, 357-375.
Jaffe, J., 1974, *Special Information and Insider Trading*, Journal of Business, v47, pp 410-428.
Jarrell, G.A., J.A. Brickley and J.M. Netter, 1988, *The Market for Corporate Control: The Empirical Evidence since 1980*, Journal of Economic Perspectives, v2, 49-68.
Jensen, M.C. and R.S. Ruback, 1983, *The Market for Corporate Control*, Journal of Financial Economics, v11, 5-50.
Kaplan, S. and M.S. Weisbach, 1992, *The Success of Acquisitions: The Evidence from Divestitures*, Journal of Finance, v47, 107-138.
Ke, B., S. Huddart and K. Petroni, 2002, What insiders know about future earnings and how they use it: evidence from insider trades, Working Paper, Social Sciences Research Network.
KPMG, 1999, *Unlocking Shareholder Value: The Keys to Success*, KPMG Global Research Report.
La Porta, R., J. Lakonishok, A. Shleifer and R. Vishny, 1995, *Good News for Value Stocks: Further Evidence of Market Inefficiency*, NBER Working Paper.
Lakonishok, J. and I. Lee, 1998, *Are insiders' trades informative*? Working Paper, Social Sciences Research Network.
Lang, L.H., R.M. Stulz and R.A. Walkling, 1989, *Managerial Performance, Tobin's Q, and the Gains from Successful Tender Offers*, Journal of Finance, v24, 137-154.
Michael, R, R.H. Thaler and K.L. Womack, 1995, *Price Reactions to Dividend Initiations and Omissions: Overreaction or Drift?* Journal of Finance, v50, 573-608.
Michaely, R. and K.L. Womack, *Conflicts of Interests and the Credibility of Underwriter Analysts Recommendations*, Review of Financial Studies, Winter, 635-686.

Mitchell, M.L. and K. Lehn, 1990, *Do Bad Bidders make Good Targets?*, Journal of Applied Corporate Finance, v3, 60-69.
Nail, L.A. , W.L. Megginson and C. Maquieira, 1998, *Wealth Creation versus Wealth Redistributions in Pure Stock-for-Stock Mergers*, Journal of Financial Economics, v48, 3-
Fama, E. F., L. Fisher, M. C. Jensen and R. Roll. *The Adjustment Of Stock Prices To New Information*, International Economic Review, 1969, v10(1), 1-21.
O'Brien, P., 1988, *Analyst's Forecasts as Earnings Expectations*, Journal of Accounting and Economics.
Parrino, J.D. and R.S. Harris, *Takeovers, Management Replacement and Post-Acquisition Operating Performance: Some Evidence from the 1980s*, Journal of Applied Corporate Finance, v11, 88-97.
Penman, S. H., 1987, *The Distribution Of Earnings News Over Time And Seasonalities In Aggregate Stock Returns*, Journal of Financial Economics, v18(2), 199-228.
Rendleman, R.J., C.P. Jones and H.A. Latene, 1982, *Empirical Anomalies based on Unexpected Earnings and the Importance of Risk Adjustments*, Journal of Financial Economics,
Rozeff, M., and M. Zaman, 1988, Market *Efficiency and Insider Trading: New Evidence, Journal of Business 61*, 25-44.
Seyhun, H.N., 1998, *Investment Intelligence from Insider Trading*, MIT Press, Cambridge.
Seyhun, N., 1986, *Insiders' Profits, Costs of Trading, and Market Efficiency*, Journal of Financial Economics 16, 189-212.
Sirower, M.L., 1996, *The Synergy Trap*, Simon & Schuster.
Vander Weide, J.H., and W.T. Carleton, 1988, *Investor Growth Expectations: Analysts Vs. History*, Journal of Portfolio Management, v14, 78-83.
Womack, K., 1996, *Do brokerage analysts' recommendations have investment value?* Journal of Finance, v51, 137-167.
Woodruff, Catherine S. and A. J. Senchack, Jr., *Intradaily Price-Volume Adjustments Of NYSE Stocks To Unexpected Earnings*, Journal of Finance, 1988, v43(2), 467-491.

Chapter 11

Ackermann, C., R. McEnally and D. Ravenscraft, 1999, *The Performance of Hedge Funds: Risk, Return and Incentives*, Journal of Finance, v54, 833-874.
Alejandro Balbás and Susana López, 2001, *Financial innovation and arbitrage in the Spanish bond market*, Working Paper, SSRN.
Black, F. and M. Scholes, 1972*, The Valuation of Option Contracts and a Test of Market Efficiency*, Journal of Finance, Vol 27, 399-417.
Brown, Stephen J., William N. Goetzmann and James Park, 2001, *Careers and Survival: Competition and Risk in the Hedge Fund and CTA Industry*, Journal of Finance, v56, 1869-1886.
Brown, Stephen J., William N. Goetzmann and Roger G. Ibbotson, 1999, *Offshore hedge funds,:survival and performance, 1989 – 1995*, *Journal of Business*, 72(1) 91-119.
Dimson, E. and C. Minio-Kozerzki, 1998, *Closed-end Funds, A Survey*, Working Paper, London Business School.
Garbade, K.D. and W.L. Silber, 1983, *Price Movements and Price Discovery in Futures and Cash Markets*, The Review of Economics and Statistics, v115, 289-297.

Gatev, E.G., W.N.Goetzmann and K.G. Rouwenhorst, 1999, *Pairs Trading, Performance of a Relative Value Arbitrage Rule*, Working Paper, SSRN.
Grinblatt, M. and F.A. Longstaff, 2000, *Financial innovation and the role of derivative securities: An empirical analysis of the U.S. treasury's strips program*, Journal of Finance
Kamara, A. and T.W. Miller, 1995, *Daily and intradaily tests of European put-call parity*. Journal of Financial and Quantitative Analysis 30, 4, 519-541.
Kin, M. A.C. Szakmary and I. Mathur, 2000, *Price Transmission Dynamics between ADRs and Their Underlying Foreign Securities*, Journal of Banking and Finance, v24, 1359-1382.
Klemkosky, R.C. and B.G. Resnick, 1979, Put-*Call Parity and Market Efficiency*, Journal of Finance, v 34, pg 1141-1155.
Lee, Charles M.C., Andrei Shleifer, and Richard H. Thaler, 1991, *Investor Sentiment and the Closed-End Fund Puzzle*, *Journal of Finance* 46, 76-110.
Lee, Charles M.C., Andrei Shleifer, and Richard H. Thaler, 1990, *Anomalies: Closed-End Mutual Funds*, *Journal of Economic Perspectives* 4, 153-164.
Liang, B., 2001, *Hedge Fund Performance: 1990-1999*, Financial Analysts Journal, Jan/Feb 2001.
Lowenstein , R., 2000, *When Genius Failed: The Rise and Fall of Long Term Capital*, Random House.
Minio-Paluello, Carolina, 1998, *The UK Closed-End Fund Discount*, PhD thesis, London Business School
Mitchell, M., and T. Pulvino, 2001. *Characteristics of risk in risk arbitrage*. Journal of Finance, v56, 2135-2175.
Neal, R., 1996, *Direct Tests of Index Arbitrage Models*, Journal of Financial and Quantitative Analysis, v31, 541-562.
Pontiff, Jeffrey, 1996, *Costly Arbitrage: Evidence from Closed-End Funds,* Quarterly Journal of Economics 111, 1135-1151.
Pontiff, Jeffrey, 1997, *Excess Volatility and Closed-End Funds,* American Economic Review 87, 155-169.
Schilling, A.G., 1992, *Market Timing better than a Buy-and-Hold Strategy*, Financial Analysts Journal (March-April), 46-50.
Shleifer, Andrei and Robert W. Vishny, 1997, *The limits of arbitrage*, Journal of Finance, v52, 35-55
Swaicki, J. and J. Hric, 2001, *Arbitrage Opportunities in Parallel Markets: The Case of the Czech Republic*, Working Paper, SSRN.
Thompson, Rex, 1978, *The Information Content of Discounts and Premiums on Closed-End Fund Shares*, Journal of Financial Economics 6, 151-186.

Chapter 12

1998, Becker, Ferson, Myers and Schill (1998)
Abhyankar , A.and P.R. Davies, 2002, *Return Predictability, Market Timing and Volatility: Evidence from the Short Rate Revisited*, Working paper, SSRN.
Ang, A. and G. Bekaert, 2001, *Stock Return Predictability: Is it there?,* Working Paper, Columbia Business School.
Bennett . J.A. and R.W. Sias, 2001, *Can Money Flows predict stock returns?*, Financial Analysts Journal, Nov/Dec.
Bernstein, 1995, *Style Investing*, John Wiley & Sons.

Breen, W., L.R. Glosten and R. Jagannathan, 1989, *Economic Significance of Predictable Variations in Stock Index Returns*, Journal of Finance, v44, 1177-1189.
Brinson, G., B. Singer and G. Beebower, G., 1991, *Determinants of portfolio performance II: an update*, Financial Analysts Journal, May-June, 40-47.
Brinson, G.L. R. Hood, and G. Beebower, 1986, *Determinants of portfolio performance*, Financial Analysts Journal, July-August, 39-44.
Campbell, J. and R.Shiller, 2001, *Valuation and the Long-Run Stock Market Outlook: An Update*, NBER Working Paper 8221, National Bureau of Economic Research.
Campbell, J. and R.Shiller, 1998, *Valuation and the Long-Run Stock Market Outlook*, Journal of Portfolio Management, v24, 11-26.
Chan, K., A. Hameed and W. Tong, 2000, *Profitability of Momentum Strategies in the International Equity Markets*, Journal of Financial and Quantitative Analysis, v35, 153-172.
Chance, D. M., and M.L. Hemler, 2001, *The performance of professional market timers: Daily evidence from executed strategies*, Journal of Financial Economics, v62, 377-411.
Chowdhury, M., J.S. Howe and J.C. Lin, 1993, *The Relation between Aggregate Insider Transactions and Stock Market Returns*, Journal of Financial and Quantitative Analysis, v28, 431-437.
Chua, J. H., R.S. Woodward, and E.C. To. 1987, *Potential Gains From Stock Market Timing in Canada*, Financial Analysts Journal (September/October), vol. 43, no. 5, 50-56.
Fisher, K.L. and M. Statman, 2000, *Investor sentiment and Stock Returns*, Financial Analysts Journal, March/April.
Graham, John R., and Campbell R. Harvey, 1996, *Market timing ability and volatility implied in investment newsletters' asset allocation recommendations*, Journal of Financial Economics 42, 397-421.
Haugen, R.A., E. Talmor and W.N. Torous, 1991, *The Effect of Volatility Changes on the Level of Stock Prices and Subsequent Expected Returns*, Journal of Finance, v46, 985-1007.
Henriksson, Roy D., and Robert C. Merton, 1981, *On market timing and investment performance. II. Statistical procedures for evaluating forecasting skills*, Journal of Business, v54, 513-533.
Hirsch, Y, 1992, *Stock Trader's Almanac*, Probus Publishing Company, Chicago.
Ibbotson, R. and Kaplan, P., 2000, *Does asset allocation explain 40, 90, or 100 per cent of performance?*, Financial Analysts Journal, January-February.
Jagannathan, Ravi, and Robert A. Korajczyk, 1986, *Assessing the market timing of managed portfolio*s, Journal of Business 59, 217-235.
James Farreell, Jr., 1975, *Homogeneous Stock Groupings: Implications for Portfolio Management*, Financial Analysts Journal (May-June), 50-62.
Jeffrey, R., 1984, *The Folly of Stock Market Timing*, Financial Analysts Journal (July-August), 102-110.
Kao, D., and R. D. Shumaker. *Equity Style Timing (corrected)*, Financial Analysts Journal, vol. 55, no. 1 (January/February 1999): 37-48.
Kon, Stanley J., 1983, *The market-timing performance of mutual fund managers,* Journal of Business, v56, 323-347.
Nuttall, J.A. and J. Nuttall, 1998, *Asset Allocation Claims - Truth or Fiction?,* Working Paper.
Pradhuman, S., 2000, *Small Cap Dynamics*, Bloomberg Press.
Sharpe, W. F., 1975, *Are Gains Likely From Market Timing*, Financial Analysts Journal, vol. 31, no. 2 (March/April): 60-69.

Sorensen, E.H. and T. Burke, 1986, *Portfolio Returns from Active Industry Group Rotation*, Financial Analysts Journal (September –October), 43-50.
Stovall, 1996, *Sector Investing*, McGraw Hill.
Treynor, Jack L., and Kay Mazuy, 1966, *Can mutual funds outguess the market? Harvard Business Review* 44, 131-136.

Chapter 13
Ahmed, P., P. Gangopadhyay and S. Nanda, 2001, *Performance of Emerging Market Mutual Funds and U.S. Monetary Policy*, Working paper, SSRN.
Atkinson, S.M., S.B. Baird and M.B. Frye, 2001, *Do female mutual fund managers manage differently?* Working Paper, SSRN.
Barber, B.M. and T. Odean, *Too many cooks spoil the profits*, Financial Analysts Journal, January/February 2000.
Bauer, R., K. Koedijk and R. Otten, 2002, *International Evidence on Ethical Mutual Fund Performance and Investment Style*, Working paper, SSRN.
Bernstein, 1995, *Style Investing*, John Wiley & Sons.
Blake, C.R. and M. M. Morey, 2000, *Morningstar Ratings and Mutual Fund Performance*, Journal of Financial and Quantitative Analysis, v35, 451-483.
Bogle, J.C., 1994, *Bogle on Mutual funds*, Richard D. Irwin.
Borensztein, E.R. and R. G. Gelos, 2001, *A panic-prone pack? The Behavior of Emerging Market Mutual Funds*, Working Paper, IMF Working Paper No. 00/198
Brown, K.C. and K.V. Harlow, 2002, *Staying the Course: The Impact of Investment Style Consistency on Mutual Fund Performance*, Working Paper, SSRN.
Brown, S.J., W.N. Goetzmann, Hiraki, Otsuki and Shirashi, 2001, The Japanese Open-End Fund Puzzle, Journal of Business, v74, 59-77.
Brown, Stephen J., and William N. Goetzmann, 1995, *Performance persistence*, *Journal of Finance* 50, 679-698.
Brown, Stephen J., William N. Goetzmann, Roger G. Ibbotson, and Stephen A. Ross, 1992, *Survivorship bias in performance studies*, *Review of Financial Studies* 5, 553-580.
Cai, J., K.C. Chan and T. Yamada, 1997, *The Performance of Japanese Mutual Funds*, Review of Financial Studies, v10, 237-273.
Carhart, Mark M., 1997, *On persistence in mutual fund performance*, Journal of Finance 52, 57-82.
Chalmers, J.M.R. R.M. Edelen and G.B. Kadlec, 1999, *An Analysis of Mutual Fund Trading Costs*, Working paper, SSRN.
Chan, L.K.C., H.L. Chen and J. Lakonishok, 1999, *On Mutual Fund Investment Styles*, NBER working paper
Chen, J.L., N. Jegadeesh and R. Wermers, 2000, *The Value of Active Mutual Fund Management: An Examination of the Stockholdings and Trades of Fund Managers*, Journal of Financial and Quantitative Analysis, v35, 343-368.
Chevalier and Ellison, 1999, *Are Some Mutual Fund Managers Better Than Others? Cross-Sectional Patterns in Behavior and Performance*, Journal of Finance, v54, 875-899.
Detzler, M.L., 1999, *The Value of Mutual Fund Rankings to Individual Investors*, Working Paper, SSRN.
Ellis, C.D, 1998, *Winning the Loser's Game*, McGraw Hill.

Elton, E.J., M.J. Gruber, G. Comer and K.Li, 2002, *Spiders: Where are the bugs?* In Exchange Traded Funds, NYU Working Ppaers.
Fama, E.F. and K.R. French, 1992, *The Cross-Section of Expected Returns*, Journal of Finance, v47, 427-466.
Grinblatt, M. and S.Titman, 1992, *The persistence of mutual fund performance*, Journal of Finance, v42, 1977-1984.
Goetzmann, W.N. and R. Ibbotson, 1994, *Do winners repeat? Patterns in mutual fund performance*, Journal of Portfolio Management, v20, 9-18.
Hendricks, Patel and Zeckhauser, 1995, *Hot Hands in Mutual Funds: Short run persistence in performance*, 1974-1987, Journal of Finance, v48, 93-130.
Indro, D.C., C.X. Jiang, M.Y. Hu and W.Y. Lee, 1999, *Mutual Fund Performance: Does size matter?,* Financial Analysts Journal, May/June, v55, 74-87.
James, C. and J. Karceskсi, 2002, *Captured Money? Differences in the Performance Characteristics of Retail and Institutional Mutual Funds*, Working Paper, SSRN.
Jensen, M., 1968, *The Performance of Mutual Funds in the period 1945-64*, Journal of Finance, v2, 389-416.
Karen Damato. Karen Damato, *Morningstar Edges Toward One-Year Ratings*, Wall Street Journal, April 5th, 1996.
Lakonishok, Shleifer and Vishny, 1994, *Contrarian Investment, Extrapolation, and Risk*, Journal of Finance, v49, 1541-1578.
Lehmann, B.N. and D.M. Modest, 1987, *Mutual Fund Performance Evaluation: A Comparison of Benchmarks and Benchmark Comparisons*, Journal of Finance, v42, 233-265.
Malkiel, B.G., 1995, *Returns from Investing in Equity Mutual Funds 1971 to 1991*, Journal of Finance, v50, 549-572.
Morey , M.R., 1998, *Should You Carry the Load? A Comprehensive Analysis of Load and No-Load Mutual Fund Out-of-Sample Performance*, Working paper, SSRN.
O'Neal, E.S., 2001, *Window Dressing and Equity Mutual Funds*, Working paper, SSRN.
Odean, T., 1988, *Are investors reluctant to realize their losses*, Journal of Finance, v53, pg 1775-1798.
Otten and Bams, 2002, *European Mutual Fund Performance*, European Financial Management, v8, 75-101.
Riepe, M.W. and J. Zils, 1997, *Are Enhanced Index Mutual Funds Worthy of Their Name?*, Working Paper, Ibbotson Associates.