CHAPTER 13

RISKY ASSETS

In the last chapter we examined a model of individual behavior under uncertainty and the role of two economic institutions for dealing with uncertainty: insurance markets and stock markets. In this chapter we will further explore how stock markets serve to allocate risk. In order to do this, it is convenient to consider a simplified model of behavior under uncertainty.

13.1 Mean-Variance Utility

In the last chapter we examined the expected utility model of choice under uncertainty. Another approach to choice under uncertainty is to describe the probability distributions that are the objects of choice by a few parameters and think of the utility function as being defined over those parameters. The most popular example of this approach is the **mean-variance model**. Instead of thinking that a consumer's preferences depend on the entire probability distribution of his wealth over every possible outcome, we suppose that his preferences can be well described by considering just a few summary statistics about the probability distribution of his wealth.

Let us suppose that a random variable w takes on the values w_s for $s = 1, \ldots, S$ with probability π_s. The **mean** of a probability distribution is simply its average value:

$$\mu_w = \sum_{s=1}^{S} \pi_s w_s.$$

This is the formula for an average: take each outcome w_s, weight it by the probability that it occurs, and sum it up over all outcomes.[1]

The **variance** of a probability distribution is the average value of $(w - \mu_w)^2$:

$$\sigma_w^2 = \sum_{s=1}^{S} \pi_s (w_s - \mu_w)^2.$$

The variance measures the "spread" of the distribution and is a reasonable measure of the riskiness involved. A closely related measure is the **standard deviation**, denoted by σ_w, which is the square root of the variance: $\sigma_w = \sqrt{\sigma_w^2}$.

The mean of a probability distribution measures its average value—what the distribution is centered around. The variance of the distribution measures the "spread" of the distribution—how spread out it is around the mean. See Figure 13.1 for a graphical depiction of probability distributions with different means and variances.

The mean-variance model assumes that the utility of a probability distribution that gives the investor wealth w_s with a probability of π_s can be expressed as a function of the mean and variance of that distribution, $u(\mu_w, \sigma_w^2)$. Or, if it is more convenient, the utility can be expressed as a function of the mean and standard deviation, $u(\mu_w, \sigma_w)$. Since both variance and standard deviation are measures of the riskiness of the wealth distribution, we can think of utility as depending on either one.

This model can be thought of as a simplification of the expected utility model described in the preceding chapter. If the choices that are being made can be completely characterized in terms of their mean and variance, then a utility function for mean and variance will be able to rank choices in the same way that an expected utility function will rank them. Furthermore, even if the probability distributions cannot be completely characterized by their means and variances, the mean-variance model may well serve as a reasonable approximation to the expected utility model.

We will make the natural assumption that a higher expected return is good, other things being equal, and that a higher variance is bad. This is simply another way to state the assumption that people are typically averse to risk.

[1] The Greek letter μ, mu, is pronounced "mew." The Greek letter σ, sigma, is pronounced "sig-ma."

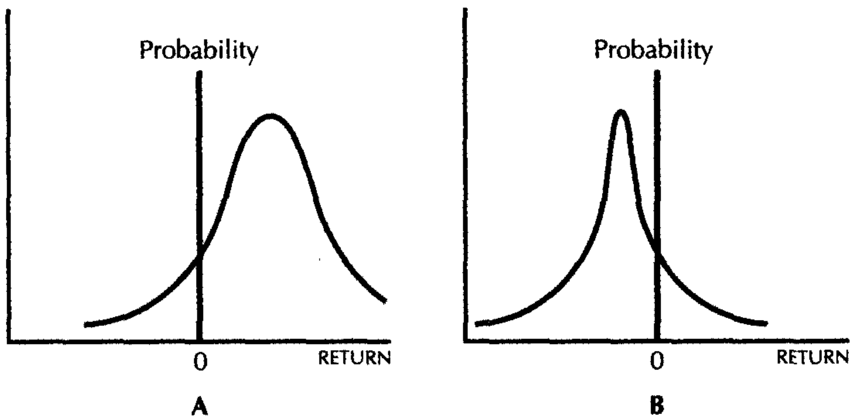

Mean and variance. The probability distribution depicted in panel A has a positive mean, while that depicted in panel B has a negative mean. The distribution in panel A is more "spread out" than the one in panel B, which means that it has a larger variance.

Let us use the mean-variance model to analyze a simple portfolio problem. Suppose that you can invest in two different assets. One of them, the **risk-free asset**, always pays a fixed rate of return, r_f. This would be something like a Treasury bill that pays a fixed rate of interest regardless of what happens.

The other asset is a **risky asset**. Think of this asset as being an investment in a large mutual fund that buys stocks. If the stock market does well, then your investment will do well. If the stock market does poorly, your investment will do poorly. Let m_s be the return on this asset if state s occurs, and let π_s be the probability that state s will occur. We'll use r_m to denote the expected return of the risky asset and σ_m to denote the standard deviation of its return.

Of course you don't have to choose one or the other of these assets; typically you'll be able to divide your wealth between the two. If you hold a fraction of your wealth x in the risky asset, and a fraction $(1-x)$ in the risk-free asset, the expected return on your portfolio will be given by

$$\begin{aligned} r_x &= \sum_{s=1}^{S}(xm_s + (1-x)r_f)\pi_s \\ &= x\sum_{s=1}^{S} m_s\pi_s + (1-x)r_f\sum_{s=1}^{S}\pi_s. \end{aligned}$$

Since $\sum \pi_s = 1$, we have

$$r_x = xr_m + (1-x)r_f.$$

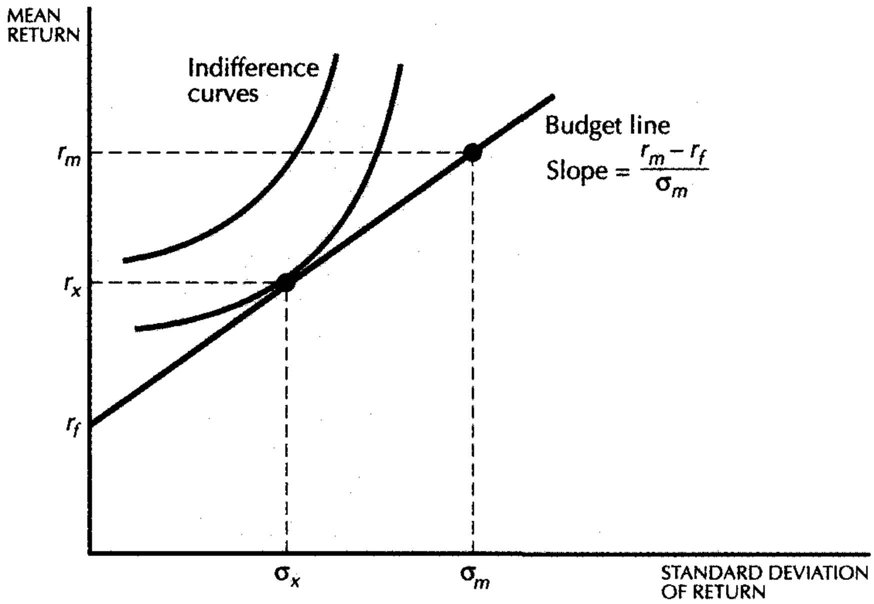

Risk and return. The budget line measures the cost of achieving a larger expected return in terms of the increased standard deviation of the return. At the optimal choice the indifference curve must be tangent to this budget line.

Thus the expected return on the portfolio is a weighted average of the two expected returns.

The variance of your portfolio return will be given by

$$\sigma_x^2 = \sum_{s=1}^{S} (xm_s + (1-x)r_f - r_x)^2 \pi_s.$$

Substituting for r_x, this becomes

$$\begin{aligned}\sigma_x^2 &= \sum_{s=1}^{S} (xm_s - xr_m)^2 \pi_s \\ &= \sum_{s=1}^{S} x^2 (m_s - r_m)^2 \pi_s \\ &= x^2 \sigma_m^2.\end{aligned}$$

Thus the standard deviation of the portfolio return is given by

$$\sigma_x = \sqrt{x^2 \sigma_m^2} = x\sigma_m.$$

It is natural to assume that $r_m > r_f$, since a risk-averse investor would never hold the risky asset if it had a lower expected return than the risk-free asset. It follows that if you choose to devote a higher fraction of your wealth to the risky asset, you will get a higher expected return, but you will also incur higher risk. This is depicted in Figure 13.2.

If you set $x = 1$ you will put all of your money in the risky asset and you will have an expected return and standard deviation of (r_m, σ_m). If you set $x = 0$ you will put all of your wealth in the sure asset and you have an expected return and standard deviation of $(r_f, 0)$. If you set x somewhere between 0 and 1, you will end up somewhere in the middle of the line connecting these two points. This line gives us a budget line describing the market tradeoff between risk and return.

Since we are assuming that people's preferences depend only on the mean and variance of their wealth, we can draw indifference curves that illustrate an individual's preferences for risk and return. If people are risk averse, then a higher expected return makes them better off and a higher standard deviation makes them worse off. This means that standard deviation is a "bad." It follows that the indifference curves will have a positive slope, as shown in Figure 13.2.

At the optimal choice of risk and return the slope of the indifference curve has to equal the slope of the budget line in Figure 13.2. We might call this slope the **price of risk** since it measures how risk and return can be traded off in making portfolio choices. From inspection of Figure 13.2 the price of risk is given by

$$p = \frac{r_m - r_f}{\sigma_m}. \tag{13.1}$$

So our optimal portfolio choice between the sure and the risky asset could be characterized by saying that the marginal rate of substitution between risk and return must be equal to the price of risk:

$$\text{MRS} = -\frac{\Delta U/\Delta\sigma}{\Delta U/\Delta\mu} = \frac{r_m - r_f}{\sigma_m}. \tag{13.2}$$

Now suppose that there are many individuals who are choosing between these two assets. Each one of them has to have his marginal rate of substitution equal to the price of risk. Thus in equilibrium all of the individuals' MRSs will be equal: when people are given sufficient opportunities to trade risks, the equilibrium price of risk will be equal across individuals. Risk is like any other good in this respect.

We can use the ideas that we have developed in earlier chapters to examine how choices change as the parameters of the problem change. All of the framework of normal goods, inferior goods, revealed preference, and so on can be brought to bear on this model. For example, suppose that an individual is offered a choice of a new risky asset y that has a mean return of r_y, say, and a standard deviation of σ_y, as illustrated in Figure 13.3.

If offered the choice between investing in x and investing in y, which will the consumer choose? The original budget set and the new budget set are both depicted in Figure 13.3. Note that every choice of risk and return that was possible in the original budget set is possible with the new budget

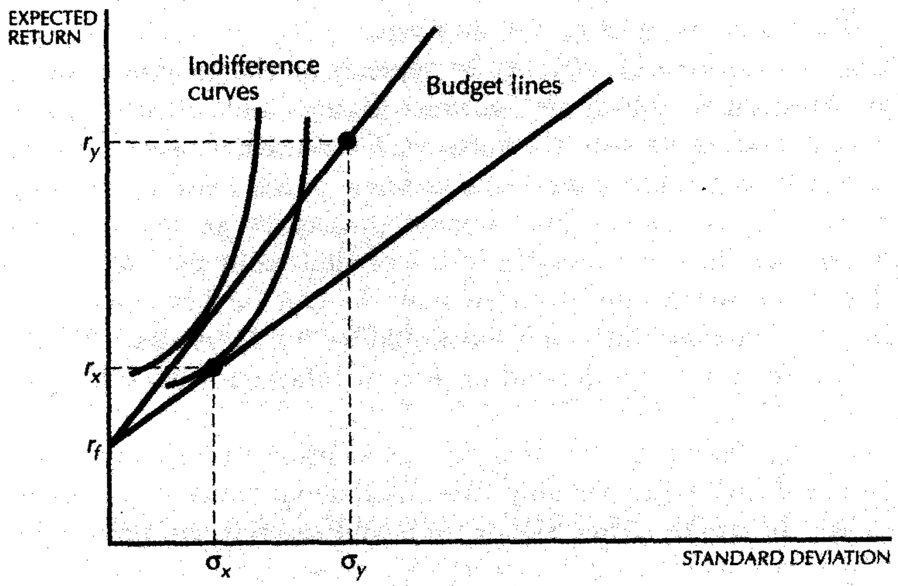

Preferences between risk and return. The asset with risk-return combination y is preferred to the one with combination x.

set since the new budget set contains the old one. Thus investing in the asset y and the risk-free asset is definitely better than investing in x and the risk-free asset, since the consumer can choose a better final portfolio.

The fact that the consumer can choose how much of the risky asset he wants to hold is very important for this argument. If this were an "all or nothing" choice where the consumer was compelled to invest all of his money in either x or y, we would get a very different outcome. In the example depicted in Figure 13.3, the consumer would prefer investing all of his money in x to investing all of his money in y, since x lies on a higher indifference curve than y. But if he can mix the risky asset with the risk-free asset, he would always prefer to mix with y rather than to mix with x.

13.2 Measuring Risk

We have a model above that describes the price of risk ... but how do we measure the *amount* of risk in an asset? The first thing that you would probably think of is the standard deviation of an asset's return. After all, we are assuming that utility depends on the mean and variance of wealth, aren't we?

In the above example, where there is only one risky asset, that is exactly right: the amount of risk in the risky asset is its standard deviation. But if there are many risky assets, the standard deviation is not an appropriate measure for the amount of risk in an asset.

This is because a consumer's utility depends on the mean and variance of total wealth—not the mean and variance of any single asset that he might hold. What matters is how the returns of the various assets a consumer holds *interact* to create a mean and variance of his wealth. As in the rest of economics, it is the marginal impact of a given asset on total utility that determines its value, not the value of that asset held alone. Just as the value of an extra cup of coffee may depend on how much cream is available, the amount that someone would be willing to pay for an extra share of a risky asset will depend on how it interacts with other assets in his portfolio.

Suppose, for example, that you are considering purchasing two assets, and you know that there are only two possible outcomes that can happen. Asset A will be worth either \$10 or −\$5, and asset B will be worth either −\$5 or \$10. But when asset A is worth \$10, asset B will be worth −\$5 and vice versa. In other words the values of the two assets will be *negatively correlated:* when one has a large value, the other will have a small value.

Suppose that the two outcomes are equally likely, so that the average value of each asset will be \$2.50. Then if you don't care about risk at all and you must hold one asset or the other, the most that you would be willing to pay for either one would be \$2.50—the expected value of each asset. If you are averse to risk, you would be willing to pay even less than \$2.50.

But what if you can hold both assets? Then if you hold one share of each asset, you will get \$5 whichever outcome arises. Whenever one asset is worth \$10, the other is worth −\$5. Thus, if you can hold both assets, the amount that you would be willing to pay to purchase *both* assets would be \$5.

This example shows in a vivid way that the value of an asset will depend in general on how it is correlated with other assets. Assets that move in opposite directions—that are negatively correlated with each other—are very valuable because they reduce overall risk. In general the value of an asset tends to depend much more on the correlation of its return with other assets than with its own variation. Thus the amount of risk in an asset depends on its correlation with other assets.

It is convenient to measure the risk in an asset relative to the risk in the stock market as a whole. We call the riskiness of a stock relative to the risk of the market the **beta** of a stock, and denote it by the Greek letter β. Thus, if i represents some particular stock, we write β_i for its riskiness relative to the market as a whole. Roughly speaking:

$$\beta_i = \frac{\text{how risky asset } i \text{ is}}{\text{how risky the stock market is}}.$$

If a stock has a beta of 1, then it is just as risky as the market as a whole;

when the market moves up by 10 percent, this stock will, on the average, move up by 10 percent. If a stock has a beta of less than 1, then when the market moves up by 10 percent, the stock will move up by less than 10 percent. The beta of a stock can be estimated by statistical methods to determine how sensitive the movements of one variable are relative to another, and there are many investment advisory services that can provide you with estimates of the beta of a stock.[2]

13.3 Equilibrium in a Market for Risky Assets

We are now in a position to state the equilibrium condition for a market with risky assets. Recall that in a market with only certain returns, we saw that all assets had to earn the same rate of return. Here we have a similar principle: all assets, after adjusting for risk, have to earn the same rate of return.

The catch is about adjusting for risk. How do we do that? The answer comes from the analysis of optimal choice given earlier. Recall that we considered the choice of an optimal portfolio that contained a riskless asset and a risky asset. The risky asset was interpreted as being a mutual fund—a diversified portfolio including many risky assets. In this section we'll suppose that this portfolio consists of *all* risky assets.

Then we can identify the expected return on this market portfolio of risky assets with the market expected return, r_m, and identify the standard deviation of the market return with the market risk, σ_m. The return on the safe asset is r_f, the risk-free return.

We saw in equation (13.1) that the price of risk, p, is given by

$$p = \frac{r_m - r_f}{\sigma_m}.$$

We said above that the amount of risk in a given asset i relative to the total risk in the market is denoted by β_i. This means that to measure the *total* amount of risk in asset i, we have to multiply by the market risk, σ_m. Thus the total risk in asset i is given by $\beta_i \sigma_m$.

What is the cost of this risk? Just multiply the total amount of risk, $\beta_i \sigma_m$, by the price of risk. This gives us the *risk adjustment:*

$$\begin{aligned}
\text{risk adjustment} &= \beta_i \sigma_m p \\
&= \beta_i \sigma_m \frac{r_m - r_f}{\sigma_m} \\
&= \beta_i (r_m - r_f).
\end{aligned}$$

[2] The Greek letter β, beta, is pronounced "bait-uh." For those of you who know some statistics, the beta of a stock is defined to be $\beta_i = \text{cov}(\tilde{r}_i, \tilde{r}_m)/\text{var}(\tilde{r}_m)$. That is, β_i is the covariance of the return on the stock with the market return divided by the variance of the market return.

Now we can state the equilibrium condition in markets for risky assets: in equilibrium all assets should have the same risk-adjusted rate of return. The logic is just like the logic used in Chapter 12: if one asset had a higher risk-adjusted rate of return than another, everyone would want to hold the asset with the higher risk-adjusted rate. Thus in equilibrium the risk-adjusted rates of return must be equalized.

If there are two assets i and j that have expected returns r_i and r_j and betas of β_i and β_j, we must have the following equation satisfied in equilibrium:

$$r_i - \beta_i(r_m - r_f) = r_j - \beta_j(r_m - r_f).$$

This equation says that in equilibrium the risk-adjusted returns on the two assets must be the same—where the risk adjustment comes from multiplying the total risk of the asset by the price of risk.

Another way to express this condition is to note the following. The risk-free asset, by definition, must have $\beta_f = 0$. This is because it has zero risk, and β measures the amount of risk in an asset. Thus for any asset i we must have

$$r_i - \beta_i(r_m - r_f) = r_f - \beta_f(r_m - r_f) = r_f.$$

Rearranging, this equation says

$$r_i = r_f + \beta_i(r_m - r_f)$$

or that the expected return on any asset must be the risk-free return plus the risk adjustment. This latter term reflects the extra return that people demand in order to bear the risk that the asset embodies. This equation is the main result of the **Capital Asset Pricing Model (CAPM)**, which has many uses in the study of financial markets.

13.4 How Returns Adjust

In studying asset markets under certainty, we showed how prices of assets adjust to equalize returns. Let's look at the same adjustment process here.

According to the model sketched out above, the expected return on any asset should be the risk-free return plus the risk premium:

$$r_i = r_f + \beta_i(r_m - r_f).$$

In Figure 13.4 we have illustrated this line in a graph with the different values of beta plotted along the horizontal axis and different expected returns on the vertical axis. According to our model, all assets that are held

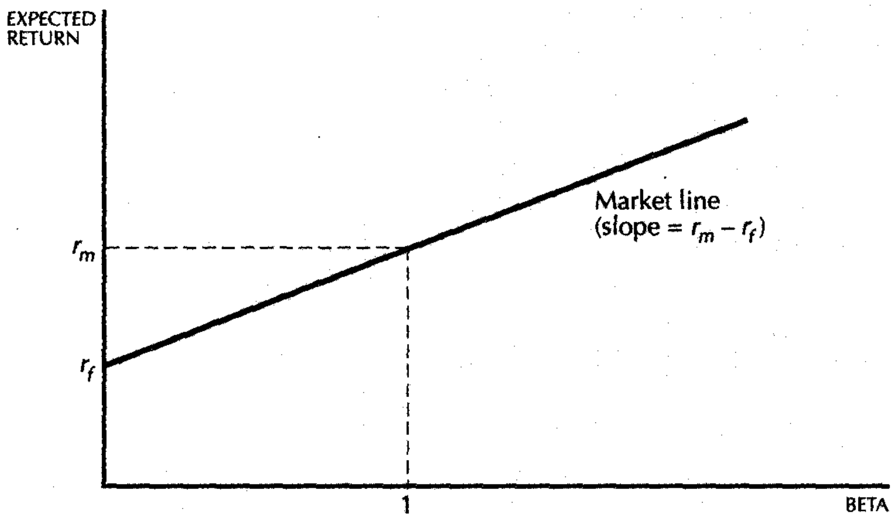

The market line. The market line depicts the combinations of expected return and beta for assets held in equilibrium.

in equilibrium have to lie along this line. This line is called the **market line**.

What if some asset's expected return and beta didn't lie on the market line? What would happen?

The expected return on the asset is the expected change in its price divided by its current price:

$$r_i = \text{ expected value of } \frac{p_1 - p_0}{p_0}.$$

This is just like the definition we had before, with the addition of the word "expected." We have to include "expected" now since the price of the asset tomorrow is uncertain.

Suppose that you found an asset whose expected return, adjusted for risk, was higher than the risk-free rate:

$$r_i - \beta_i(r_m - r_f) > r_f.$$

Then this asset is a very good deal. It is giving a higher risk-adjusted return than the risk-free rate.

When people discover that this asset exists, they will want to buy it. They might want to keep it for themselves, or they might want to buy it and sell it to others, but since it is offering a better tradeoff between risk and return than existing assets, there is certainly a market for it.

But as people attempt to buy this asset they will bid up today's price: p_0 will rise. This means that the expected return $r_i = (p_1 - p_0)/p_0$ will fall. How far will it fall? Just enough to lower the expected rate of return back down to the market line.

Thus it is a good deal to buy an asset that lies above the market line For when people discover that it has a higher return given its risk than assets they currently hold, they will bid up the price of that asset.

This is all dependent on the hypothesis that people agree about the amount of risk in various assets. If they disagree about the expected returns or the betas of different assets, the model becomes much more complicated

EXAMPLE: Ranking Mutual Funds

The Capital Asset Pricing Model can be used to compare different investments with respect to their risk and their return. One popular kind of investment is a mutual fund. These are large organizations that accept money from individual investors and use this money to buy and sell stocks of companies. The profits made by such investments are then paid out to the individual investors.

The advantage of a mutual fund is that you have professionals managing your money. The disadvantage is they charge you for managing it. These fees are usually not terribly large, however, and most small investors are probably well advised to use a mutual fund.

But how do you choose a mutual fund in which to invest? You want one with a high expected return of course, but you also probably want one with a minimum amount of risk. The question is, how much risk are you willing to tolerate to get that high expected return?

One thing that you might do is to look at the historical performance of various mutual funds and calculate the average yearly return and the beta—the amount of risk—of each mutual fund you are considering. Since we haven't discussed the precise definition of beta, you might find it hard to calculate. But there are books where you can look up the historical betas of mutual funds.

If you plotted the expected returns versus the betas, you would get a diagram similar to that depicted in Figure 13.5.[3] Note that the mutual funds with high expected returns will generally have high risk. The high expected returns are there to compensate people for bearing risk.

One interesting thing you can do with the mutual fund diagram is to compare investing with professional managers to a very simple strategy like investing part of your money in an **index fund**. There are several

[3] See Michael Jensen, "The Performance of Mutual Funds in the Period 1945–1964," *Journal of Finance*, 23 (May 1968), 389–416, for a more detailed discussion of how to examine mutual fund performance using the tools we have sketched out in this chapter. Mark Grinblatt and Sheridan Titman have examined more recent data in "Mutual Fund Performance: An Analysis of Quarterly Portfolio Holdings," *The Journal of Business*, 62 (July 1989), 393–416.

indices of stock market activity like the Dow-Jones Industrial Average, or the Standard and Poor's Index, and so on. The indices are typically the average returns on a given day of a certain group of stocks. The Standard and Poor's Index, for example, is based on the average performance of 500 large stocks in the United States.

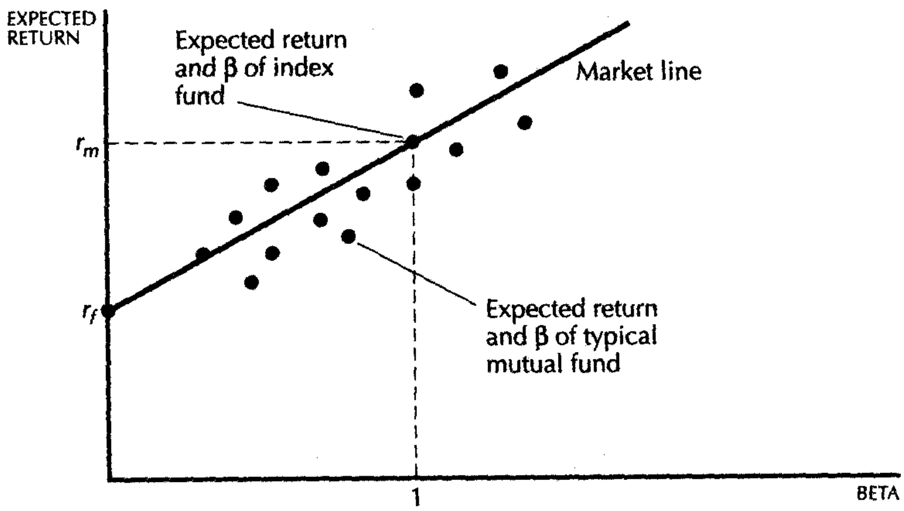

Mutual funds. Comparing the returns on mutual fund investment to the market line.

Figure 13.5

An **index fund** is a mutual fund that holds the stocks that make up such an index. This means that you are guaranteed to get the average performance of the stocks in the index, virtually by definition. Since holding the average is not a very difficult thing to do—at least compared to trying to beat the average—index funds typically have low management fees. Since an index fund holds a very broad base of risky assets, it will have a beta that is very close to 1—it will be just as risky as the market as a whole, because the index fund holds nearly all the stocks in the market as a whole.

How does an index fund do as compared to the typical mutual fund? Remember the comparison has to be made with respect to both risk and return of the investment. One way to do this is to plot the expected return and beta of a Standard and Poor's Index fund, and draw the line connecting it to the risk-free rate, as in Figure 13.5. You can get any combination of risk and return on this line that you want just by deciding how much money you want to invest in the risk-free asset and how much you want to invest in the index fund.

Now let's count the number of mutual funds that plot below this line. These are mutual funds that offer risk and return combinations that are dominated by those available by the index fund/risk-free asset combinations. When this is done, it turns out that the vast majority of the risk-return combinations offered by mutual funds are below the line. The number of funds that plot above the line is no more than could be expected by chance alone.

But seen another way, this finding might not be too surprising. The stock market is an incredibly competitive environment. People are always trying to find undervalued stocks in order to purchase them. This means that on average, stocks are usually trading for what they're really worth. If that is the case, then betting the averages is a pretty reasonable strategy—since beating the averages is almost impossible.

Summary

1. We can use the budget set and indifference curve apparatus developed earlier to examine the choice of how much money to invest in risky and riskless assets.

2. The marginal rate of substitution between risk and return will have to equal the slope of the budget line. This slope is known as the price of risk.

3. The amount of risk present in an asset depends to a large extent on its correlation with other assets. An asset that moves opposite the direction of other assets helps to reduce the overall risk of your portfolio.

4. The amount of risk in an asset relative to that of the market as a whole is called the **beta** of the asset.

5. The fundamental equilibrium condition in asset markets is that risk-adjusted returns have to be the same.

REVIEW QUESTIONS

1. If the risk-free rate of return is 6%, and if a risky asset is available with a return of 9% and a standard deviation of 3%, what is the maximum rate of return you can achieve if you are willing to accept a standard deviation of 2%? What percentage of your wealth would have to be invested in the risky asset?

2. What is the price of risk in the above exercise?

3. If a stock has a β of 1.5, the return on the market is 10%, and the risk-free rate of return is 5%, what expected rate of return should this stock offer according to the Capital Asset Pricing Model? If the expected value of the stock is \$100, what price should the stock be selling for today?