EXAMPLE: The Effect of Taxation on Investment in Risky Assets

How does the level of investment in a risky asset behave when you tax its return? If the individual pays taxes at rate t, then the after-tax returns will be $(1-t)r_g$ and $(1-t)r_b$. Thus the first-order condition determining his optimal investment, x, will be

$$EU'(x) = \pi u'(w + x(1-t)r_g)(1-t)r_g + (1-\pi)u'(w + x(1-t)r_b)(1-t)r_b = 0.$$

Canceling the $(1-t)$ terms, we have

$$EU'(x) = \pi u'(w + x(1-t)r_g)r_g + (1-\pi)u'(w + x(1-t)r_b)r_b = 0. \qquad (12.6)$$

Let us denote the solution to the maximization problem without taxes—when $t = 0$—by x^* and denote the solution to the maximization problem with taxes by $\hat{x}$. What is the relationship between x^* and $\hat{x}$?

Your first impulse is probably to think that $x^* > \hat{x}$—that taxation of a risky asset will tend to discourage investment in it. But that turns out to be exactly wrong! Taxing a risky asset in the way we described will actually *encourage* investment in it!

In fact, there is an exact relation between x^* and $\hat{x}$. It must be the case that

$$\hat{x} = \frac{x^*}{1-t}.$$

The proof is simply to note that this value of $\hat{x}$ satisfies the first-order condition for the optimal choice in the presence of the tax. Substituting this choice into equation (12.6) we have

$$\begin{aligned} EU'(\hat{x}) &= \pi u'(w + \frac{x^*}{1-t}(1-t)r_g)r_g \\ &\quad + (1-\pi)u'(w + \frac{x^*}{1-t}(1-t)r_b)r_b \\ &= \pi u'(w + x^* r_g)r_g + (1-\pi)u'(w + x^* r_b)r_b = 0, \end{aligned}$$

where the last equality follows from the fact that x^* is the optimal solution when there is no tax.

What is going on here? How can imposing a tax increase the amount of investment in the risky asset? Here is what is happening. When the tax is imposed, the individual will have less of a gain in the good state, but he will also have *less of a loss in the bad state.* By scaling his original investment up by $1/(1-t)$ the consumer can reproduce the same *after-tax* returns that he had before the tax was put in place. The tax reduces his expected return, but it also reduces his risk: by increasing his investment the consumer can get exactly the same pattern of returns he had before and thus completely offset the effect of the tax. A tax on a risky investment represents a tax on the gain when the return is positive—but it represents a subsidy on the loss when the return is negative.

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."

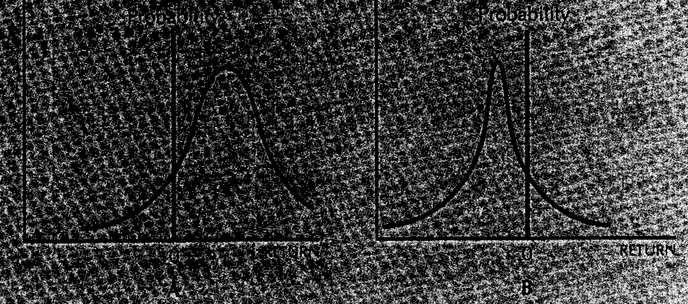

Figure 13.1

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

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

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

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

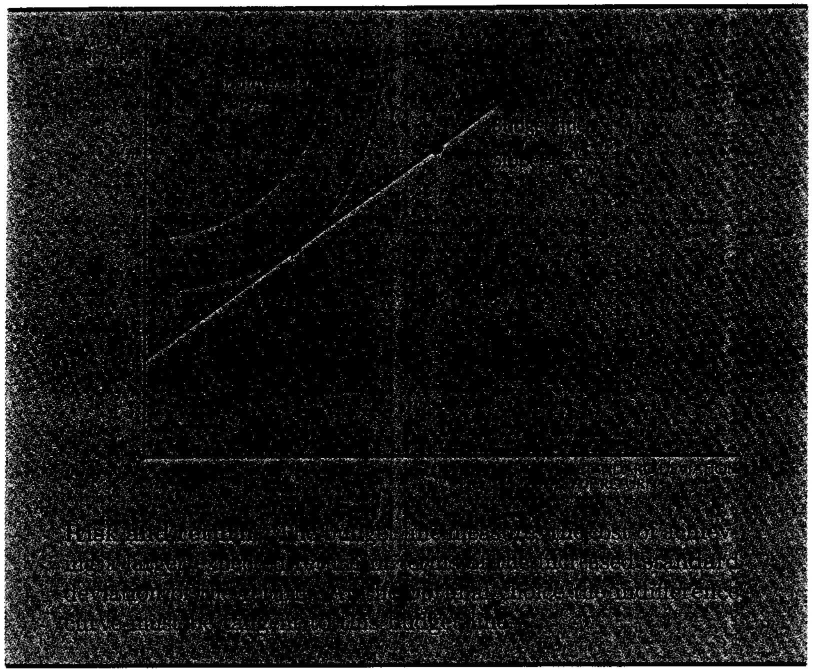

Figure 13.2

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

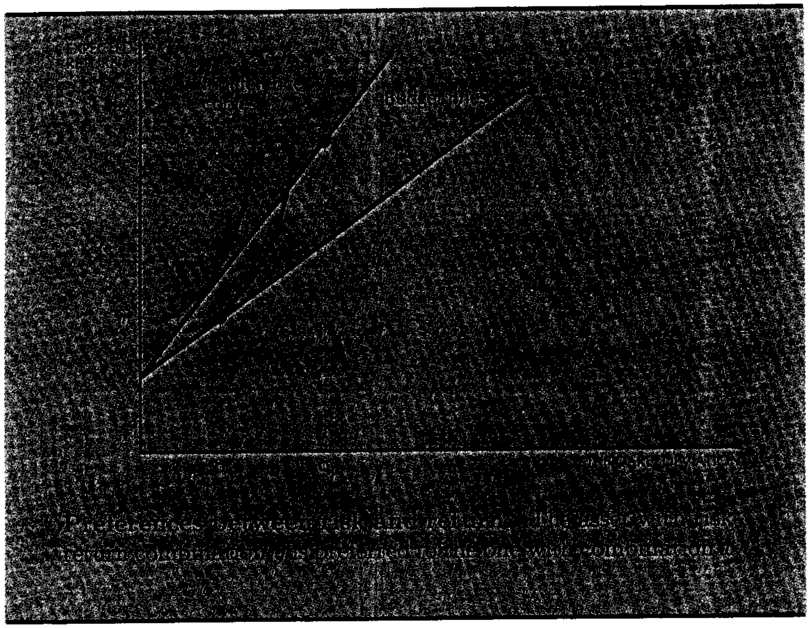

Figure 13.3

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;