

CHAPTER 12

UNCERTAINTY

Uncertainty is a fact of life. People face risks every time they take a shower, walk across the street, or make an investment. But there are financial institutions such as insurance markets and the stock market that can mitigate at least some of these risks. We will study the functioning of these markets in the next chapter, but first we must study individual behavior with respect to choices involving uncertainty.

12.1 Contingent Consumption

Since we now know all about the standard theory of consumer choice, let's try to use what we know to understand choice under uncertainty. The first question to ask is what is the basic “thing” that is being chosen?

The consumer is presumably concerned with the **probability distribution** of getting different consumption bundles of goods. A probability distribution consists of a list of different outcomes—in this case, consumption bundles—and the probability associated with each outcome. When a consumer decides how much automobile insurance to buy or how much to

invest in the stock market, he is in effect deciding on a pattern of probability distribution across different amounts of consumption.

For example, suppose that you have \$100 now and that you are contemplating buying lottery ticket number 13. If number 13 is drawn in the lottery, the holder will be paid \$200. This ticket costs, say, \$5. The two outcomes that are of interest are the event that the ticket is drawn and the event that it isn't.

Your original endowment of wealth—the amount that you would have if you did not purchase the lottery ticket—is \$100 if 13 is drawn, and \$100 if it isn't drawn. But if you buy the lottery ticket for \$5, you will have a wealth distribution consisting of \$295 if the ticket is a winner, and \$95 if it is not a winner. The original endowment of probabilities of wealth in different circumstances has been changed by the purchase of the lottery ticket. Let us examine this point in more detail.

In this discussion we'll restrict ourselves to examining monetary gambles for convenience of exposition. Of course, it is not money alone that matters; it is the consumption that money can buy that is the ultimate "good" being chosen. The same principles apply to gambles over goods, but restricting ourselves to monetary outcomes makes things simpler. Second, we will restrict ourselves to very simple situations where there are only a few possible outcomes. Again, this is only for reasons of simplicity.

Above we described the case of gambling in a lottery; here we'll consider the case of insurance. Suppose that an individual initially has \$35,000 worth of assets, but there is a possibility that he may lose \$10,000. For example, his car may be stolen, or a storm may damage his house. Suppose that the probability of this event happening is $p = .01$. Then the probability distribution the person is facing is a 1 percent probability of having \$25,000 of assets, and a 99 percent probability of having \$35,000.

Insurance offers a way to change this probability distribution. Suppose that there is an insurance contract that will pay the person \$100 if the loss occurs in exchange for a \$1 premium. Of course the premium must be paid whether or not the loss occurs. If the person decides to purchase \$10,000 dollars of insurance, it will cost him \$100. In this case he will have a 1 percent chance of having \$34,900 (\$35,000 of other assets – \$10,000 loss + \$10,000 payment from the insurance payment – \$100 insurance premium) and a 99 percent chance of having \$34,900 (\$35,000 of assets – \$100 insurance premium). Thus the consumer ends up with the same wealth no matter what happens. He is now fully insured against loss.

In general, if this person purchases K dollars of insurance and has to pay a premium γK , then he will face the gamble:¹

probability .01 of getting $\$25,000 + K - \gamma K$

¹ The Greek letter γ , gamma, is pronounced "gam-ma."

and

probability .99 of getting $\$35,000 - \gamma K$.

What kind of insurance will this person choose? Well, that depends on his preferences. He might be very conservative and choose to purchase a lot of insurance, or he might like to take risks and not purchase any insurance at all. People have different preferences over probability distributions in the same way that they have different preferences over the consumption of ordinary goods.


In fact, one very fruitful way to look at decision making under uncertainty is just to think of the money available under different circumstances as different goods. A thousand dollars after a large loss has occurred may mean a very different thing from a thousand dollars when it hasn't. Of course, we don't have to apply this idea just to money: an ice cream cone if it happens to be hot and sunny tomorrow is a very different good from an ice cream cone if it is rainy and cold. In general, consumption goods will be of different value to a person depending upon the circumstances under which they become available.

Let us think of the different outcomes of some random event as being different **states of nature**. In the insurance example given above there were two states of nature: the loss occurs or it doesn't. But in general there could be many different states of nature. We can then think of a **contingent consumption plan** as being a specification of what will be consumed in each different state of nature—each different outcome of the random process. *Contingent* means depending on something not yet certain, so a contingent consumption plan means a plan that depends on the outcome of some event. In the case of insurance purchases, the contingent consumption was described by the terms of the insurance contract: how much money you would have if a loss occurred and how much you would have if it didn't. In the case of the rainy and sunny days, the contingent consumption would just be the *plan* of what would be consumed given the various outcomes of the weather.

People have preferences over different plans of consumption, just like they have preferences over actual consumption. It certainly might make you feel better now to know that you are fully insured. People make choices that reflect their preferences over consumption in different circumstances, and we can use the theory of choice that we have developed to analyze those choices.

If we think about a contingent consumption plan as being just an ordinary consumption bundle, we are right back in the framework described in the previous chapters. We can think of preferences as being defined over different consumption plans, with the "terms of trade" being given by the budget constraint. We can then model the consumer as choosing the best consumption plan he or she can afford, just as we have done all along.

Let's describe the insurance purchase in terms of the indifference-curve analysis we've been using. The two states of nature are the event that the loss occurs and the event that it doesn't. The contingent consumptions are the values of how much money you would have in each circumstance. We can plot this on a graph as in Figure 12.1.



Your endowment of contingent consumption is \$25,000 in the “bad” state—if the loss occurs—and \$35,000 in the “good” state—if it doesn’t occur. Insurance offers you a way to move away from this endowment point. If you purchase K dollars’ worth of insurance, you give up γK dollars of consumption possibilities in the good state in exchange for $K - \gamma K$ dollars of consumption possibilities in the bad state. Thus the consumption you lose in the good state, divided by the extra consumption you gain in the bad state, is

$$\frac{\Delta C_g}{\Delta C_b} = -\frac{\gamma K}{K - \gamma K} = -\frac{\gamma}{1 - \gamma}.$$

This is the slope of the budget line through your endowment. It is just as if the price of consumption in the good state is $1 - \gamma$ and the price in the bad state is γ .

We can draw in the indifference curves that a person might have for contingent consumption. Here again it is very natural for indifference curves to have a convex shape: this means that the person would rather have a constant amount of consumption in each state than a large amount in one state and a low amount in the other.

Given the indifference curves for consumption in each state of nature, we can look at the choice of how much insurance to purchase. As usual, this will be characterized by a tangency condition: the marginal rate of substitution between consumption in each state of nature should be equal to the price at which you can trade off consumption in those states.

Of course, once we have a model of optimal choice, we can apply all of the machinery developed in early chapters to its analysis. We can examine how the demand for insurance changes as the price of insurance changes, as the wealth of the consumer changes, and so on. The theory of consumer behavior is perfectly adequate to model behavior under uncertainty as well as certainty.

EXAMPLE: Catastrophe Bonds

We have seen that insurance is a way to transfer wealth from good states of nature to bad states of nature. Of course there are two sides to these transactions: those who buy insurance and those who sell it. Here we focus on the sell side of insurance.

The sell side of the insurance market is divided into a retail component, which deals directly with end buyers, and a wholesale component, in which insurers sell risks to other parties. The wholesale part of the market is known as the **reinsurance market**.

Typically, the reinsurance market has relied on large investors such as pension funds to provide financial backing for risks. However, some reinsurers rely on large individual investors. Lloyd's of London, one of the most famous reinsurance consortia, generally uses private investors.

Recently, the reinsurance industry has been experimenting with **catastrophe bonds**, which, according to some, are a more flexible way to provide reinsurance. These bonds, generally sold to large institutions, have typically been tied to natural disasters, like earthquakes or hurricanes.

A financial intermediary, such as a reinsurance company or an investment bank, issues a bond tied to a particular insurable event, such as an earthquake involving, say, at least \$500 million in insurance claims. If there is no earthquake, investors are paid a generous interest rate. But if the earthquake occurs and the claims exceed the amount specified in the bond, investors sacrifice their principal and interest.

Catastrophe bonds have some attractive features. They can spread risks widely and can be subdivided indefinitely, allowing each investor to bear

only a small part of the risk. The money backing up the insurance is paid in advance, so there is no default risk to the insured.

From the economist's point of view, "cat bonds" are a form of **state contingent security**, that is, a security that pays off if and only if some particular event occurs. This concept was first introduced by Nobel laureate Kenneth J. Arrow in a paper published in 1952 and was long thought to be of only theoretical interest. But it turned out that all sorts of options and other derivatives could be best understood using contingent securities. Now Wall Street rocket scientists draw on this 50-year-old work when creating exotic new derivatives such as catastrophe bonds.

12.2 Utility Functions and Probabilities

If the consumer has reasonable preferences about consumption in different circumstances, then we will be able to use a utility function to describe these preferences, just as we have done in other contexts. However, the fact that we are considering choice under uncertainty does add a special structure to the choice problem. In general, how a person values consumption in one state as compared to another will depend on the *probability* that the state in question will actually occur. In other words, the rate at which I am willing to substitute consumption if it rains for consumption if it doesn't should have something to do with how likely I think it is to rain. The preferences for consumption in different states of nature will depend on the beliefs of the individual about how likely those states are.

For this reason, we will write the utility function as depending on the probabilities as well as on the consumption levels. Suppose that we are considering two mutually exclusive states such as rain and shine, loss or no loss, or whatever. Let c_1 and c_2 represent consumption in states 1 and 2, and let π_1 and π_2 be the probabilities that state 1 or state 2 actually occurs.

If the two states are mutually exclusive, so that only one of them can happen, then $\pi_2 = 1 - \pi_1$. But we'll generally write out both probabilities just to keep things looking symmetric.

Given this notation, we can write the utility function for consumption in states 1 and 2 as $u(c_1, c_2, \pi_1, \pi_2)$. This is the function that represents the individual's preference over consumption in each state.

EXAMPLE: Some Examples of Utility Functions

We can use nearly any of the examples of utility functions that we've seen up until now in the context of choice under uncertainty. One nice example is the case of perfect substitutes. Here it is natural to weight each

consumption by the probability that it will occur. This gives us a utility function of the form

$$u(c_1, c_2, \pi_1, \pi_2) = \pi_1 c_1 + \pi_2 c_2.$$

In the context of uncertainty, this kind of expression is known as the **expected value**. It is just the average level of consumption that you would get.

Another example of a utility function that might be used to examine choice under uncertainty is the Cobb–Douglas utility function:

$$u(c_1, c_2, \pi, 1 - \pi) = c_1^\pi c_2^{1-\pi}.$$

Here the utility attached to any combination of consumption bundles depends on the pattern of consumption in a nonlinear way.

As usual, we can take a monotonic transformation of utility and still represent the same preferences. It turns out that the logarithm of the Cobb–Douglas utility will be very convenient in what follows. This will give us a utility function of the form

$$\ln u(c_1, c_2, \pi_1, \pi_2) = \pi_1 \ln c_1 + \pi_2 \ln c_2.$$

12.3 Expected Utility

One particularly convenient form that the utility function might take is the following:

$$u(c_1, c_2, \pi_1, \pi_2) = \pi_1 v(c_1) + \pi_2 v(c_2).$$

This says that utility can be written as a weighted sum of some function of consumption in each state, $v(c_1)$ and $v(c_2)$, where the weights are given by the probabilities π_1 and π_2 .

Two examples of this were given above. The perfect substitutes, or expected value utility function, had this form where $v(c) = c$. The Cobb–Douglas didn't have this form originally, but when we expressed it in terms of logs, it had the linear form with $v(c) = \ln c$.

If one of the states is certain, so that $\pi_1 = 1$, say, then $v(c_1)$ is the utility of certain consumption in state 1. Similarly, if $\pi_2 = 1$, $v(c_2)$ is the utility of consumption in state 2. Thus the expression

$$\pi_1 v(c_1) + \pi_2 v(c_2)$$

represents the average utility, or the expected utility, of the pattern of consumption (c_1, c_2) .

For this reason, we refer to a utility function with the particular form described here as an **expected utility function**, or, sometimes, a **von Neumann-Morgenstern utility function**.²

When we say that a consumer's preferences can be represented by an expected utility function, or that the consumer's preferences have the expected utility property, we mean that we can choose a utility function that has the additive form described above. Of course we could also choose a different form; any monotonic transformation of an expected utility function is a utility function that describes the same preferences. But the additive form representation turns out to be especially convenient. If the consumer's preferences are described by $\pi_1 \ln c_1 + \pi_2 \ln c_2$ they will also be described by $c_1^{\pi_1} c_2^{\pi_2}$. But the latter representation does not have the expected utility property, while the former does.

On the other hand, the expected utility function can be subjected to some kinds of monotonic transformation and still have the expected utility property. We say that a function $v(u)$ is a **positive affine transformation** if it can be written in the form: $v(u) = au + b$ where $a > 0$. A positive affine transformation simply means multiplying by a positive number and adding a constant. It turns out that if you subject an expected utility function to a positive affine transformation, it not only represents the same preferences (this is obvious since an affine transformation is just a special kind of monotonic transformation) but it also still has the expected utility property.

Economists say that an expected utility function is "unique up to an affine transformation." This just means that you can apply an affine transformation to it and get another expected utility function that represents the same preferences. But any other kind of transformation will destroy the expected utility property.

12.4 Why Expected Utility Is Reasonable

The expected utility representation is a convenient one, but is it a reasonable one? Why would we think that preferences over uncertain choices would have the particular structure implied by the expected utility function? As it turns out there are compelling reasons why expected utility is a reasonable objective for choice problems in the face of uncertainty.

The fact that outcomes of the random choice are consumption goods that will be consumed in different circumstances means that ultimately *only one* of those outcomes is actually going to occur. Either your house

² John von Neumann was one of the major figures in mathematics in the twentieth century. He also contributed several important insights to physics, computer science, and economic theory. Oscar Morgenstern was an economist at Princeton who, along with von Neumann, helped to develop mathematical game theory.

will burn down or it won't; either it will be a rainy day or a sunny day. The way we have set up the choice problem means that only one of the many possible outcomes is going to occur, and hence only one of the contingent consumption plans will actually be realized.

This turns out to have a very interesting implication. Suppose you are considering purchasing fire insurance on your house for the coming year. In making this choice you will be concerned about wealth in three situations: your wealth now (c_0), your wealth if your house burns down (c_1), and your wealth if it doesn't (c_2). (Of course, what you really care about are your consumption possibilities in each outcome, but we are simply using wealth as a proxy for consumption here.) If π_1 is the probability that your house burns down and π_2 is the probability that it doesn't, then your preferences over these three different consumptions can generally be represented by a utility function $u(\pi_1, \pi_2, c_0, c_1, c_2)$.

Suppose that we are considering the tradeoff between wealth now and one of the possible outcomes—say, how much money we would be willing to sacrifice now to get a little more money if the house burns down. *Then this decision should be independent of how much consumption you will have in the other state of nature—how much wealth you will have if the house is not destroyed.* For the house will either burn down or it won't. If it happens to burn down, then the value of extra wealth shouldn't depend on how much wealth you would have if it *didn't* burn down. Bygones are bygones—so what *doesn't* happen shouldn't affect the value of consumption in the outcome that *does* happen.

Note that this is an *assumption* about an individual's preferences. It may be violated. When people are considering a choice between two things, the amount of a third thing they have typically matters. The choice between coffee and tea may well depend on how much cream you have. But this is because you consume coffee together with cream. If you considered a choice where you rolled a die and got either coffee, *or* tea, *or* cream, then the amount of cream that you might get shouldn't affect your preferences between coffee and tea. Why? Because you are either getting one thing or the other: if you end up with cream, the fact that you might have gotten either coffee or tea is irrelevant.

Thus in choice under uncertainty there is a natural kind of "independence" between the different outcomes because they must be consumed separately—in different states of nature. The choices that people plan to make in one state of nature should be independent from the choices that they plan to make in other states of nature. This assumption is known as the **independence assumption**. It turns out that this implies that the utility function for contingent consumption will take a very special structure: it has to be additive across the different contingent consumption bundles.

That is, if c_1 , c_2 , and c_3 are the consumptions in different states of nature, and π_1 , π_2 , and π_3 are the probabilities that these three different states of

nature materialize, then if the independence assumption alluded to above is satisfied, the utility function must take the form

$$U(c_1, c_2, c_3) = \pi_1 u(c_1) + \pi_2 u(c_2) + \pi_3 u(c_3).$$

This is what we have called an expected utility function. Note that the expected utility function does indeed satisfy the property that the marginal rate of substitution between two goods is independent of how much there is of the third good. The marginal rate of substitution between goods 1 and 2, say, takes the form

$$\begin{aligned} \text{MRS}_{12} &= - \frac{\Delta U(c_1, c_2, c_3) / \Delta c_1}{\Delta U(c_1, c_2, c_3) / \Delta c_2} \\ &= - \frac{\pi_1 \Delta u(c_1) / \Delta c_1}{\pi_2 \Delta u(c_2) / \Delta c_2}. \end{aligned}$$

This MRS depends only on how much you have of goods 1 and 2, not how much you have of good 3.

12.5 Risk Aversion


We claimed above that the expected utility function had some very convenient properties for analyzing choice under uncertainty. In this section we'll give an example of this.

Let's apply the expected utility framework to a simple choice problem. Suppose that a consumer currently has \$10 of wealth and is contemplating a gamble that gives him a 50 percent probability of winning \$5 and a 50 percent probability of losing \$5. His wealth will therefore be random: he has a 50 percent probability of ending up with \$5 and a 50 percent probability of ending up with \$15. The *expected value* of his wealth is \$10, and the expected utility is

$$\frac{1}{2}u(\$15) + \frac{1}{2}u(\$5).$$

This is depicted in Figure 12.2. The expected utility of wealth is the average of the two numbers $u(\$15)$ and $u(\$5)$, labeled $.5u(5) + .5u(15)$ in the graph. We have also depicted the utility of the expected value of wealth, which is labeled $u(\$10)$. Note that in this diagram the expected utility of wealth is less than the utility of the expected wealth. That is,

$$u\left(\frac{1}{2}15 + \frac{1}{2}5\right) = u(10) > \frac{1}{2}u(15) + \frac{1}{2}u(5).$$



Risk aversion. For a risk-averse consumer the utility of the expected value of wealth, $u(10)$, is greater than the expected utility of wealth, $.5u(5) + .5u(15)$.


In this case we say that the consumer is **risk averse** since he prefers to have the expected value of his wealth rather than face the gamble. Of course, it could happen that the preferences of the consumer were such that he prefers a random distribution of wealth to its expected value, in which case we say that the consumer is a **risk lover**. An example is given in Figure 12.3.

Note the difference between Figures 12.2 and 12.3. The risk-averse consumer has a *concave* utility function—its slope gets flatter as wealth is increased. The risk-loving consumer has a *convex* utility function—its slope gets steeper as wealth increases. Thus the curvature of the utility function measures the consumer's attitude toward risk. In general, the more concave the utility function, the more risk averse the consumer will be, and the more convex the utility function, the more risk loving the consumer will be.

The intermediate case is that of a linear utility function. Here the consumer is **risk neutral**: the expected utility of wealth is the utility of its expected value. In this case the consumer doesn't care about the riskiness of his wealth at all—only about its expected value.

EXAMPLE: The Demand for Insurance

Let's apply the expected utility structure to the demand for insurance that we considered earlier. Recall that in that example the person had a wealth



Risk loving. For a risk-loving consumer the expected utility of wealth, $.5u(5) + .5u(15)$, is greater than the utility of the expected value of wealth, $u(10)$.

of \$35,000 and that he might incur a loss of \$10,000. The probability of the loss was 1 percent, and it cost him γK to purchase K dollars of insurance. By examining this choice problem using indifference curves we saw that the optimal choice of insurance was determined by the condition that the MRS between consumption in the two outcomes—loss or no loss—must be equal to $-\gamma/(1-\gamma)$. Let π be the probability that the loss will occur, and $1-\pi$ be the probability that it won't occur.

Let state 1 be the situation involving no loss, so that the person's wealth in that state is

$$c_1 = \$35,000 - \gamma K,$$

and let state 2 be the loss situation with wealth

$$c_2 = \$35,000 - \$10,000 + K - \gamma K.$$

Then the consumer's optimal choice of insurance is determined by the condition that his MRS between consumption in the two outcomes be equal to the price ratio:

$$\text{MRS} = -\frac{\pi \Delta u(c_2) / \Delta c_2}{(1-\pi) \Delta u(c_1) / \Delta c_1} = -\frac{\gamma}{1-\gamma}. \quad (12.1)$$

Now let us look at the insurance contract from the viewpoint of the insurance company. With probability π they must pay out K , and with

probability $(1 - \pi)$ they pay out nothing. No matter what happens, they collect the premium γK . Then the expected profit, P , of the insurance company is

$$P = \gamma K - \pi K - (1 - \pi) \cdot 0 = \gamma K - \pi K.$$

Let us suppose that on the average the insurance company just breaks even on the contract. That is, they offer insurance at a “fair” rate, where “fair” means that the expected value of the insurance is just equal to its cost. Then we have

$$P = \gamma K - \pi K = 0,$$

which implies that $\gamma = \pi$.

Inserting this into equation (12.1) we have

$$\frac{\pi \Delta u(c_2) / \Delta c_2}{(1 - \pi) \Delta u(c_1) / \Delta c_1} = \frac{\pi}{1 - \pi}.$$

Canceling the π 's leaves us with the condition that the optimal amount of insurance must satisfy

$$\frac{\Delta u(c_1)}{\Delta c_1} = \frac{\Delta u(c_2)}{\Delta c_2}. \quad (12.2)$$

This equation says that the *marginal utility of an extra dollar of income if the loss occurs should be equal to the marginal utility of an extra dollar of income if the loss doesn't occur.*

Let us suppose that the consumer is risk averse, so that his marginal utility of money is declining as the amount of money he has increases. Then if $c_1 > c_2$, the marginal utility at c_1 would be less than the marginal utility at c_2 , and vice versa. Furthermore, if the marginal utilities of income are equal at c_1 and c_2 , as they are in equation (12.2), then we must have $c_1 = c_2$. Applying the formulas for c_1 and c_2 , we find

$$35,000 - \gamma K = 25,000 + K - \gamma K,$$

which implies that $K = \$10,000$. This means that when given a chance to buy insurance at a “fair” premium, a risk-averse consumer will always choose to fully insure.

This happens because the utility of wealth in each state depends only on the total amount of wealth the consumer has in that state—and not what he *might* have in some other state—so that if the total amounts of wealth the consumer has in each state are equal, the marginal utilities of wealth must be equal as well.

To sum up: if the consumer is a risk-averse, expected utility maximizer and if he is offered fair insurance against a loss, then he will optimally choose to fully insure.

12.6 Diversification

Let us turn now to a different topic involving uncertainty—the benefits of diversification. Suppose that you are considering investing \$100 in two different companies, one that makes sunglasses and one that makes raincoats. The long-range weather forecasters have told you that next summer is equally likely to be rainy or sunny. How should you invest your money?

Wouldn't it make sense to hedge your bets and put some money in each? By diversifying your holdings of the two investments, you can get a return on your investment that is more certain, and therefore more desirable if you are a risk-averse person.

Suppose, for example, that shares of the raincoat company and the sunglasses company currently sell for \$10 apiece. If it is a rainy summer, the raincoat company will be worth \$20 and the sunglasses company will be worth \$5. If it is a sunny summer, the payoffs are reversed: the sunglasses company will be worth \$20 and the raincoat company will be worth \$5. If you invest your entire \$100 in the sunglasses company, you are taking a gamble that has a 50 percent chance of giving you \$200 and a 50 percent chance of giving you \$50. The same magnitude of payoffs results if you invest all your money in the sunglasses company: in either case you have an expected payoff of \$125.

But look what happens if you put half of your money in each. Then, if it is sunny you get \$100 from the sunglasses investment and \$25 from the raincoat investment. But if it is rainy, you get \$100 from the raincoat investment and \$25 from the sunglasses investment. Either way, you end up with \$125 for sure. By diversifying your investment in the two companies, you have managed to reduce the overall risk of your investment, while keeping the expected payoff the same.

Diversification was quite easy in this example: the two assets were perfectly negatively correlated—when one went up, the other went down. Pairs of assets like this can be extremely valuable because they can reduce risk so dramatically. But, alas, they are also very hard to find. Most asset values move together: when GM stock is high, so is Ford stock, and so is Goodrich stock. But as long as asset price movements are not *perfectly* positively correlated, there will be some gains from diversification.

12.7 Risk Spreading

Let us return now to the example of insurance. There we considered the situation of an individual who had \$35,000 and faced a .01 probability of a \$10,000 loss. Suppose that there were 1000 such individuals. Then, on average, there would be 10 losses incurred, and thus \$100,000 lost each year. Each of the 1000 people would face an *expected loss* of .01 times \$10,000, or

\$100 a year. Let us suppose that the probability that any person incurs a loss doesn't affect the probability that any of the others incur losses. That is, let us suppose that the risks are *independent*.

Then each individual will have an expected wealth of $.99 \times \$35,000 + .01 \times \$25,000 = \$34,900$. But each individual also bears a large amount of risk: each person has a 1 percent probability of losing \$10,000.

Suppose that each consumer decides to *diversify* the risk that he or she faces. How can they do this? Answer: by selling some of their risk to other individuals. Suppose that the 1000 consumers decide to insure one another. If anybody incurs the \$10,000 loss, each of the 1000 consumers will contribute \$10 to that person. This way, the poor person whose house burns down is compensated for his loss, and the other consumers have the peace of mind that they will be compensated if that poor soul happens to be themselves! This is an example of **risk spreading**: each consumer spreads his risk over all of the other consumers and thereby reduces the amount of risk he bears.

Now on the average, 10 houses will burn down a year, so on the average, each of the 1000 individuals will be paying out \$100 a year. But this is just on the average. Some years there might be 12 losses, and other years there might be 8 losses. The probability is very small that an individual would actually have to pay out more than \$200, say, in any one year, but even so, the risk is there.

But there is even a way to diversify this risk. Suppose that the homeowners agree to pay \$100 a year for certain, whether or not there are any losses. Then they can build up a cash reserve fund that can be used in those years when there are multiple fires. They are paying \$100 a year for certain, and on average that money will be sufficient to compensate homeowners for fires.

As you can see, we now have something very much like a cooperative insurance company. We could add a few more features: the insurance company gets to invest its cash reserve fund and earn interest on its assets, and so on, but the essence of the insurance company is clearly present.

12.8 Role of the Stock Market

The stock market plays a role similar to that of the insurance market in that it allows for risk spreading. Recall from Chapter 11 that we argued that the stock market allowed the original owners of firms to convert their stream of returns over time to a lump sum. Well, the stock market also allows them to convert their risky position of having all their wealth tied up in one enterprise to a situation where they have a lump sum that they can invest in a variety of assets. The original owners of the firm have an incentive to issue shares in their company so that they can spread the risk of that single company over a large number of shareholders.

Similarly, the later shareholders of a company can use the stock market to reallocate their risks. If a company you hold shares in is adopting a policy that is too risky for your taste—or too conservative—you can sell those shares and purchase others.

In the case of insurance, an individual was able to reduce his risk to zero by purchasing insurance. For a flat fee of \$100, the individual could purchase full insurance against the \$10,000 loss. This was true because there was basically no risk in the aggregate: if the probability of the loss occurring was 1 percent, then on average 10 of the 1000 people would face a loss—we just didn't know which ones.

In the case of the stock market, there is risk in the aggregate. One year the stock market as a whole might do well, and another year it might do poorly. Somebody has to bear that kind of risk. The stock market offers a way to transfer risky investments from people who don't want to bear risk to people who are willing to bear risk.

Of course, few people outside of Las Vegas *like* to bear risk: most people are risk averse. Thus the stock market allows people to transfer risk from people who don't want to bear it to people who are willing to bear it if they are sufficiently compensated for it. We'll explore this idea further in the next chapter.

Summary

1. Consumption in different states of nature can be viewed as consumption goods, and all the analysis of previous chapters can be applied to choice under uncertainty.
2. However, the utility function that summarizes choice behavior under uncertainty may have a special structure. In particular, if the utility function is linear in the probabilities, then the utility assigned to a gamble will just be the expected utility of the various outcomes.
3. The curvature of the expected utility function describes the consumer's attitudes toward risk. If it is concave, the consumer is a risk averter; and if it is convex, the consumer is a risk lover.
4. Financial institutions such as insurance markets and the stock market provide ways for consumers to diversify and spread risks.

REVIEW QUESTIONS

1. How can one reach the consumption points to the left of the endowment in Figure 12.1?

2. Which of the following utility functions have the expected utility property? (a) $u(c_1, c_2, \pi_1, \pi_2) = a(\pi_1 c_1 + \pi_2 c_2)$, (b) $u(c_1, c_2, \pi_1, \pi_2) = \pi_1 c_1 + \pi_2 c_2^2$, (c) $u(c_1, c_2, \pi_1, \pi_2) = \pi_1 \ln c_1 + \pi_2 \ln c_2 + 17$.
3. A risk-averse individual is offered a choice between a gamble that pays \$1000 with a probability of 25% and \$100 with a probability of 75%, or a payment of \$325. Which would he choose?
4. What if the payment was \$320?
5. Draw a utility function that exhibits risk-loving behavior for small gambles and risk-averse behavior for larger gambles.
6. Why might a neighborhood group have a harder time self insuring for flood damage versus fire damage?

APPENDIX

Let us examine a simple problem to demonstrate the principles of expected utility maximization. Suppose that the consumer has some wealth w and is considering investing some amount x in a risky asset. This asset could earn a return of r_g in the “good” outcome, or it could earn a return of r_b in the “bad” outcome. You should think of r_g as being a positive return—the asset increases in value, and r_b being a negative return—a decrease in asset value.

Thus the consumer’s wealth in the good and bad outcomes will be

$$\begin{aligned} W_g &= (w - x) + x(1 + r_g) = w + xr_g \\ W_b &= (w - x) + x(1 + r_b) = w + xr_b. \end{aligned}$$

Suppose that the good outcome occurs with probability π and the bad outcome with probability $(1 - \pi)$. Then the expected utility if the consumer decides to invest x dollars is

$$EU(x) = \pi u(w + xr_g) + (1 - \pi)u(w + xr_b).$$

The consumer wants to choose x so as to maximize this expression.

Differentiating with respect to x , we find the way in which utility changes as x changes:

$$EU'(x) = \pi u'(w + xr_g)r_g + (1 - \pi)u'(w + xr_b)r_b. \quad (12.3)$$

The second derivative of utility with respect to x is

$$EU''(x) = \pi u''(w + xr_g)r_g^2 + (1 - \pi)u''(w + xr_b)r_b^2. \quad (12.4)$$

If the consumer is risk averse his utility function will be concave, which implies that $u''(w) < 0$ for every level of wealth. Thus the second derivative of expected utility is unambiguously negative. Expected utility will be a concave function of x .

Consider the change in expected utility for the first dollar invested in the risky asset. This is just equation (12.3) with the derivative evaluated at $x = 0$:

$$\begin{aligned} EU'(0) &= \pi u'(w)r_g + (1 - \pi)u'(w)r_b \\ &= u'(w)[\pi r_g + (1 - \pi)r_b]. \end{aligned}$$

The expression inside the brackets is the **expected return** on the asset. If the expected return on the asset is negative, then expected utility must decrease when the first dollar is invested in the asset. But since the second derivative of expected utility is negative due to concavity, then utility must continue to decrease as additional dollars are invested.

Hence we have found that if the *expected value* of a gamble is negative, a risk averter will have the highest *expected utility* at $x^* = 0$: he will want no part of a losing proposition.

On the other hand, if the expected return on the asset is positive, then increasing x from zero will increase expected utility. Thus he will always want to invest a little bit in the risky asset, no matter how risk averse he is.

Expected utility as a function of x is illustrated in Figure 12.4. In Figure 12.4A the expected return is negative, and the optimal choice is $x^* = 0$. In Figure 12.4B the expected return is positive over some range, so the consumer wants to invest some positive amount x^* in the risky asset.




Figure
12.4

How much to invest in the risky asset. In panel A, the optimal investment is zero, but in panel B the consumer wants to invest a positive amount.

The optimal amount for the consumer to invest will be determined by the condition that the derivative of expected utility with respect to x be equal to zero. Since the second derivative of utility is automatically negative due to concavity, this will be a global maximum.

Setting (12.3) equal to zero we have

$$EU'(x) = \pi u'(w + xr_g)r_g + (1 - \pi)u'(w + xr_b)r_b = 0. \quad (12.5)$$

This equation determines the optimal choice of x for the consumer in question.

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

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

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.