

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

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

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,

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

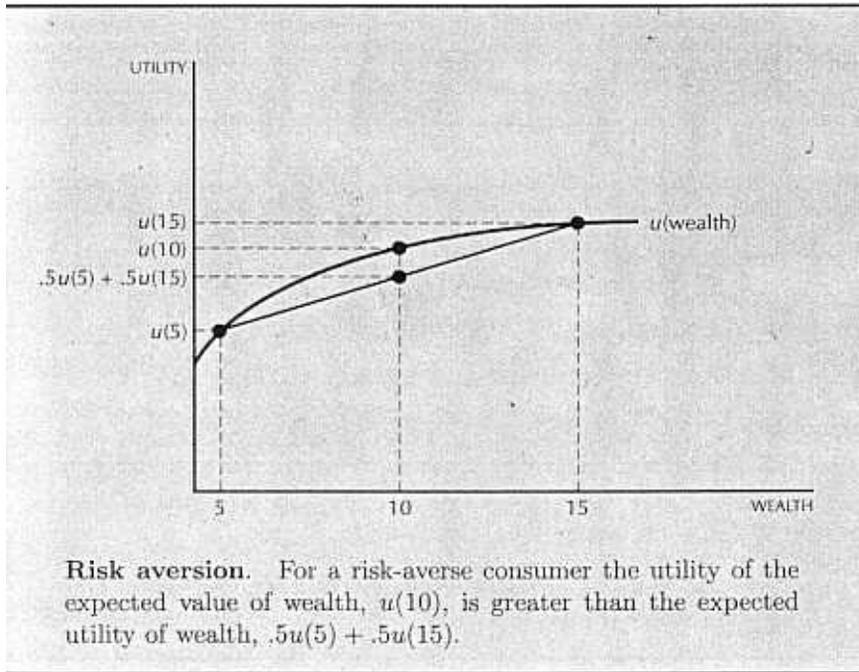


Figure 12.2

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

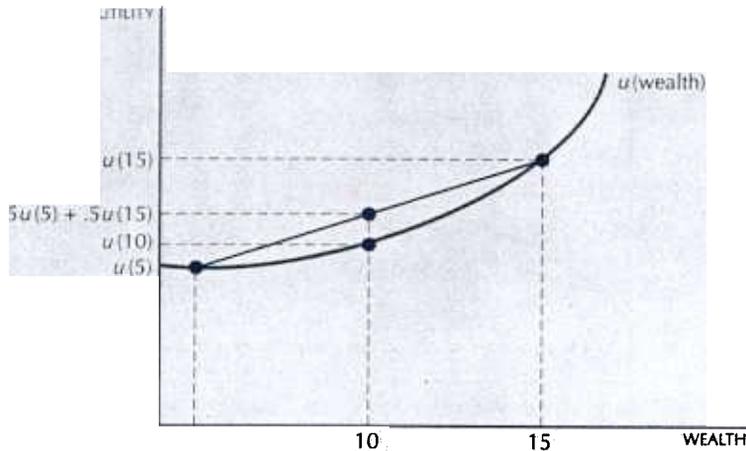


Figure 12.3

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

$$c_1 \quad \$35,000 \quad \gamma K,$$

and let state 2 be the loss situation with wealth

$$c_2 \quad \$35,000 \quad \$10,000 + K \quad \gamma K.$$

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:

$$\text{MRS} \quad \frac{\pi \Delta u(c_2) / \Delta c_2}{(1-\pi) \Delta u(c_1) / \Delta c_1} \quad \frac{\gamma}{\gamma} \quad (12.)$$

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.