Choice Under Uncertainty

Microeconomics I

Many important economic decisions involve an element of uncertainty. Although it is formally possible to analyze these situation using the general theory of consumer choice, there are good reasons to develop a more specialized theory: We can use the special structure of uncertain alternatives to restrict the preferences that "rational" individuals may hold. This allows us to derive stronger implications that are more suited to situations where choices are made in the presence of risk rather than where outcomes are perfectly certain.

1 Notation

A decision maker faces a choice among a number of risky alternatives. For example, an investor may need to choose between two different portfolios. Or a home owner may need to decide on whether and what insurance policy to purchase for their house. Each risky alternative may result in one of a number of possible outcomes, but which outcome will eventually occur is uncertain at the time the choice has to be made. Going back to the previous examples, it is uncertain what rate of return each of the assets in a portfolio will generate over the year. It is also uncertain as to whether the house that is being insured will suffer any damage and if so, to what the extent.

Let C be the set of all possible outcomes (or consequences). These outcomes could take many forms and so we treat C as an abstract set to allow for very general outcomes. Often times we will be working with outcomes that can be expressed in monetary terms. This will be the case especially later in this chapter, when we talk about risk-aversion.

To avoid some technicalities, we assume, for now, that the number of possible outcomes in C is finite, and we index those by $n = 1, \ldots, N$:

$$C = \{c_{1,}c_2, \ldots.., c_N\}.$$

The basic building block of the theory on choice under uncertainty is the concept of a *lottery*, which is used to represent risky alternatives. A lottery is a probability distribution over the set of possible outcomes.

Definition 1. *Given C, a* **simple lottery** *L is a list $L = (p_1, \ldots, p_N)$ with $p_n \geq 0$ for all n and $\sum_{n=1}^{N} p_n = 1$, where p_n is the probability of outcome c_n occurring.*

We assume throughout that the probabilities of the various outcomes arising are objectively known. For example, the risky alternatives might be monetary gambles on the roll of a fair dice or on the spin of an unbiased roulette wheel.

Let $\mathscr{L}(C)$ be the set of all possible simple lotteries over the set of outcomes C:

$$\mathscr{L}(C) = \Delta(C) = \{p \in \mathbb{R}_+^N : \sum_{n=1}^{N} p_n = 1\}$$

Note that a simple lottery can be presented geometrically as a point in the $(N-1)$-dimensional simplex $\Delta(C)$.

In a simple lottery, the outcomes that may arise are certain. A more general variant of a lottery, known as compound lottery, allows the outcomes of the lottery themselves to be simple lotteries. Without loss of generality, we can assume that the simple lotteries have the same outcome space.

Definition 2. *Given K simple lotteries $L_k = (p_1^k, \ldots, p_N^k)$, $k = 1, \ldots, K$, and probabilities $\alpha_k \geq 0$ with $\sum_{k=1}^{K} \alpha_k = 1$, the* **compound lottery** *$(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$ is the risky alternative that yields a simple lottery L_k with probability α_k.*

For any compound lottery $(L_1, \ldots, L_K; \alpha_1, \ldots, \alpha_K)$ we can calculate the corresponding *reduced lottery* as the simple lottery $L = (p_1, \ldots, p_N)$ that generates the same ultimate distribution over outcomes. This is done by setting $p_n = \sum_{k=1}^{K} \alpha_k p_n^k$. Therefore, the reduced lottery of any compound lottery can be obtained by vector addition: $L = \alpha_1 L_1 + \cdots + \alpha_K L_K$.

Example

Consider an outcome space with $N = 3$ and the following simple lotteries:

$$L_1 = (1, 0, 0); \;\; L_2 = (\frac{1}{4}, \frac{3}{8}, \frac{3}{8}); \;\; L_3 = (\frac{1}{4}, \frac{1}{8}, \frac{5}{8}); \;\; L_4 = (\frac{13}{24}, 0, \frac{11}{24}); \;\; L_5 = (\frac{17}{24}, \frac{7}{24}, 0).$$

Consider a compound lottery $(L_1, L_2, L_3; \frac{1}{2}, \frac{1}{3}, \frac{1}{6})$. The reduced lottery for this compound lottery, $\frac{1}{2}L_1 + \frac{1}{3}L_2 + \frac{1}{6}L_3$ can be written as

$$L' = (\frac{5}{8}, \frac{7}{48}, \frac{11}{48}).$$

Now consider a different compound lottery given by $(L_4, L_5; \frac{1}{2}, \frac{1}{2})$. The reduced lottery for this compound lottery, $\frac{1}{2}L_4 + \frac{1}{2}L_5$, can also be written as

$$L' = (\frac{5}{8}, \frac{7}{48}, \frac{11}{48}).$$

Hence, we have two different compound lotteries that yield the same reduced lottery.

2 Preferences over Lotteries

Notice that the decision maker does not choose outcomes from the set C directly, but instead chooses alternatives from the set $\mathscr{L}(C)$. Therefore, we want to study the decision maker's preferences over lotteries. The theoretical analysis to follow rests on a basic consequentialist premise: We assume that for any risky alternative, only the corresponding simple lottery over final outcomes is of relevance to the decision maker. Whether the probabilities of various outcomes arise as a result of a simple lottery or of a more complex compound lottery has no significance. Hence, by this premise, the two compound lotteries from the previous example should be equivalent in the eyes of the decision maker. In accordance with this consequentialist premise, we take the set of alternatives to be $\mathscr{L}$. The decision maker has preferences over $\mathscr{L}$, which are represented by the preference relation $\succeq$.[1]

We first assume that $\succeq$ is a rational preference relation, which requires the following two axioms.

Axiom 1. Completeness: *For all $L, L' \in \mathscr{L}$, we have that $L \succeq L'$ or $L' \succeq L$ (or both).*

Axiom 2. Transitivity: *For all $L, L', L'' \in \mathscr{L}$, if $L \succeq L'$ and $L' \succeq L''$, then $L \succeq L''$.*

Adding the next axiom is sufficient for the existence of a (continuous) utility function representation.

Axiom 3. Continuity: *For any $L, L', L'' \in \mathscr{L}$, the sets*

$$\{\alpha \in [0,1] : \alpha L + (1-\alpha)L' \succeq L''\} \subset [0,1]$$

and

$$\{\beta \in [0,1] : L'' \succeq \beta L + (1-\beta)L'\} \subset [0,1]$$

are closed.

Another way to interpret the continuity axiom is to say that for any $L, L', L'' \in \mathscr{L}$ such that $L \succeq L'' \succeq L'$ there exists $\alpha \in [0,1]$ such that $L'' \sim \alpha L + (1-\alpha)L'$.

The axioms thus far should be familiar from basic consumer choice theory where the consumer has preferences over possible commodity bundles (perfectly certain outcomes). They guarantee the existence of a utility function that represents $\succeq$, that is, a function $U : \mathscr{L} \to \mathbb{R}$

[1]A strict preference is denoted by $\succ$ and indifference by $\sim$.

such that $L \succeq L'$ if and only if $U(L) \geq U(L')$. This ensures that a mathematical representation of the preference is possible and that a rational decision maker will act as if maximizing this utility function.

The next axiom will allow us to impose considerably more structure on $U(\cdot)$.

Axiom 4. Independence: *For all $L, L', L'' \in \mathscr{L}$ and for all $\alpha \in (0,1)$,*

$$L \succeq L' \Leftrightarrow \alpha L + (1-\alpha)L'' \succeq \alpha L' + (1-\alpha)L''.$$

In words, the independence axiom means that if we were to mix each of two lotteries with a third one in exactly the same way (same proportion α), then the preference ordering of the two resulting mixtures does not depend on (is independent of) the particular third lottery used. This assumption is at the heart of the theory of choice under uncertainty, because it exploits, in a fundamental manner, the structure of uncertainty present in the model. In the theory of consumer demand (choice under certainty) there is no reason to believe that a consumer's preferences over various bundles of two goods should be independent of the quantities of a third good that he will consume; indeed, issues of complementarity and substitutability play an important role in this context. In the presence of uncertainty, however, the choice between two lotteries L and L' should be independent of the presence of a third lottery L'', since the lotteries are never consumed together but only instead of one another.

Claim 1. *If $\succeq$ satisfies independence, then the indifference curves are parallel straight lines.*

Indifference curves are straight lines if, for every pair of lotteries L, L' we have that $L \sim L'$ implies $\alpha L + (1-\alpha)L' \sim L$ for all $\alpha \in [0,1]$.

The independence axiom is intimately linked to the representability of preferences over lotteries by a utility function that has an *expected utility form.* We next define this property and study some of its features.

3 Expected Utility

Definition 3. *Given an outcome space $C = \{c_1, \ldots, c_N\}$, the utility function $U : \mathscr{L} \to \mathbb{R}$ has the* **expected utility form** *if there is an assignment of numbers $u = (u_1, \ldots, u_N)$ such that for every lottery $L = (p_1, \ldots, p_N) \in \mathscr{L}$ we have*

$$U(L) = p_1 u_1 + \cdots + p_N u_N = \sum_{n=1}^{N} p_n u_n.$$

A utility function $U : \mathscr{L} \to \mathbb{R}$ with the expected utility form is called von Neumann-Morgenstern (v.N-M) utility function.

It will be important to distinguish between the utility function $U(\cdot)$ defined on lotteries, and the utility function $u(\cdot)$ defined on outcomes. For this reason, we will call $U(\cdot)$ the *v.N-M expected utility function* and $u(\cdot)$ the *Bernoulli utility function*.

Observe that if we let L^n denote the lottery that yields outcome c_n with probability one, then $U(L^n) = u_n$. Thus, the term *expected utility* is appropriate because with the v.N-M expected utility form, the utility of a lottery can be thought of as the expected value of the utilities u_n of the N outcomes. Also notice, that the expression $U(L) = \sum_{n=1}^{N} p_n u_n$ is a linear function in the probabilities $(p_1, \ldots, p_N)$.

Proposition 4. *Suppose that $U : \mathscr{L} \to \mathbb{R}$ is a v.N-M utility function for the preference relation $\succeq$ on $\mathscr{L}$. Then $\tilde{U} : \mathscr{L} \to \mathbb{R}$ is another v.N-M utility function for $\succeq$ if and only if there are scalars $\beta > 0$ and γ such that $\tilde{U} = \beta U(L) + \gamma$ for every $L \in \mathscr{L}$.*

The expected utility property is a *cardinal* property of utility functions defined on the space of lotteries. In particular, the above result shows that the expected utility form is preserved only by strictly increasing *affine* transformations. A consequence of this is that for a utility function with the expected utility form, differences of utilities have meaning.

The expected utility theorem says that if the decision maker's rational preferences over lotteries satisfy the continuity and independence axioms, then his preferences are representable by a utility function with the expected utility form. It is the most important result in the theory of choice under uncertainty, which we next state and then prove formally.

Theorem 5. (Expected Utility Theorem) *Suppose that the rational preference relation $\succeq$ on the space of lotteries $\mathscr{L}$ satisfies the continuity and independence axioms. Then $\succeq$ admits a utility representation of the expected utility form.*

In other words, the theorem states that, if the rational preference relation $\succeq$ satisfies the continuity and independence axioms, we can assign a number u_n to each outcome c_n, $n = 1, \ldots, N$ in such a manner that for any two lotteries $L = (p_1, \ldots, p_N)$ and $L' = (p'_1, \ldots, p'_N)$, we have

$$L \succeq L' \quad \text{if and only if} \quad \sum_{n=1}^{N} p_n u_n \geq \sum_{n=1}^{N} p'_n u_n.$$

Proof: The proof is organized in a succession of five steps. Suppose $\underline{L}$ and $\overline{L}$ are, respectively the worst and best lottery in $\mathscr{L}$, so that $\overline{L} \succeq L \succeq \underline{L}$ for all $L \in \mathscr{L}$.[2] We assume that $\overline{L} \succ \underline{L}$ (since if $\overline{L} \sim \underline{L}$, then all lotteries in $\mathscr{L}$ are indifferent and the conclusion of the theorem holds trivially).

Step 1. If $L \succ L'$ and $\alpha \in (0,1)$, then $L \succ \alpha L + (1-\alpha)L' \succ L'$.
This claim follows from the independence axiom: Since $L \succ L'$, by the independence axiom we get

$$L = \alpha L + (1-\alpha)L \succ \alpha L + (1-\alpha)L' \succ \alpha L' + (1-\alpha)L' = L'.$$

Step 2. Let $\alpha, \beta \in [0,1]$. Then, $\beta\overline{L} + (1-\beta)\underline{L} \succ \alpha\overline{L} + (1-\alpha)\underline{L}$ if and only if $\beta > \alpha$.
"If" direction: Suppose that $\beta > \alpha$. Note that we can write

$$\beta\overline{L} + (1-\beta)\underline{L} = \gamma\overline{L} + (1-\gamma)(\alpha\overline{L} + (1-\alpha)\underline{L})$$

where $\gamma = \frac{\beta-\alpha}{1-\alpha} \in (0,1]$. By Step 1, we have that $\overline{L} \succ \alpha\overline{L} + (1-\alpha)\underline{L}$. Applying Step 1 again, we obtain $\gamma\overline{L} + (1-\gamma)(\alpha\overline{L} + (1-\alpha)\underline{L}) \succ \alpha\overline{L} + (1-\alpha)\underline{L}$, and so we conclude that

[2] The existence of worst and best lottery actually follows as a consequence of C being finite and the independence axiom.

$\beta\overline{L} + (1-\beta)\underline{L} \succ \alpha\overline{L} + (1-\alpha)\underline{L}$.
"Only if" direction: Suppose that $\beta \leq \alpha$. If $\beta = \alpha$, we must have $\beta\overline{L} + (1-\beta)\underline{L} \sim \alpha\overline{L} + (1-\alpha)\underline{L}$. So suppose that $\beta < \alpha$. By the argument used in the previous paragraph and just reversing the roles of α and β, we must have that $\alpha\overline{L} + (1-\alpha)\underline{L} \succ \beta\overline{L} + (1-\beta)\underline{L}$.

Step 3. For any $L \in \mathscr{L}$, there is a unique α_L such that $[\alpha_L\overline{L} + (1-\alpha_L)\underline{L}] \sim L$.
Existence of such an α_L is implied by the continuity axiom and the fact that $\overline{L}$ and $\underline{L}$ are, respectively, the best and worst lottery. Uniqueness follows from Step 2.

Step 4. The function $U : \mathscr{L} \to \mathbb{R}$ that assigns $U(L) = \alpha_L$ for all $L \in \mathscr{L}$ represents the preference relation $\succeq$.
By Step 3, for any two lotteries $L, L' \in \mathscr{L}$, we have

$$L \succeq L' \quad \text{if and only if} \quad \alpha_L\overline{L} + (1-\alpha_L)\underline{L} \succeq \alpha_{L'}\overline{L} + (1-\alpha_{L'})\underline{L}.$$

Thus, by Step 2, $L \succeq L'$ if and only if $\alpha_L \geq \alpha_{L'}$.

Step 5. The utility function $U(\cdot)$ that assigns $U(L) = \alpha_L$ for all $L \in \mathscr{L}$ is linear and therefore has the expected utility form.
We want to show that for any $L, L' \in \mathscr{L}$, and $\beta \in [0,1]$, we have $U(\beta L + (1-\beta)L') = \beta U(L) + (1-\beta)U(L')$. By definition we have

$$L \sim U(L)\overline{L} + (1-U(L))\underline{L}$$

and

$$L' \sim U(L')\overline{L} + (1-U(L'))\underline{L}.$$

By the independence axiom applied twice we have

$$\begin{aligned} \beta L + (1-\beta)L' &\sim \beta[U(L)\overline{L} + (1-U(L))\underline{L}] + (1-\beta)L' \\ &\sim \beta[U(L)\overline{L} + (1-U(L))\underline{L}] + (1-\beta)[U(L')\overline{L} + (1-U(L'))\underline{L}]. \end{aligned}$$

Rearranging terms in the last line we obtain

$$\beta L + (1-\beta)L' \sim [\beta U(L) + (1-\beta)U(L')]\overline{L} + [1 - \beta U(L) - (1-\beta)U(L')]\underline{L}.$$

By construction of $U(\cdot)$ in Step 4 we get

$$U(\beta L + (1-\beta)L') = \beta U(L) + (1-\beta)U(L').$$

Together, Steps 1-5 establish the existence of a utility function representing $\succeq$ that has the expected utility form.

■

4 Reasons to Doubt EU Theory

Now that we have stated and proven the expected utility theorem, we will evaluate its advantages, as well as some of the reasons it has been doubted. A first advantage of the expected utility theorem is technical: It is extremely convenient analytically. It is very easy to work with expected utility and very difficult to do without it, which is probably the reason behind its pervasive use in economics.

A second advantage is normative: Expected utility may provide a valuable plan to action. People often find it hard to think systematically about risky alternatives. But if an individual believes his choices should satisfy the axioms in which the theorem is based (notably, the independence axiom), then the theorem can be used as a guide in this decision process.

As a descriptive theory, however, the expected utility theorem (and, by implication, its central assumption, the independence axiom), is not without difficulties. The following paradoxes constitute some of the major challenges to its plausibility.

The Allais Paradox

This is the oldest and most famous challenge to the expected utility theorem. There are three possible monetary prizes (hence, the number of outcomes is $N = 3$): $c_1 = \$5,000,000$, $c_2 = \$1,000,000$, $c_3 = 0$. Consider the following two choices. First, you have to choose between the lotteries L_1 and L_2:

$$L_1 = (0, 1, 0) \qquad L_2 = (0.10, 0.89, 0.01).$$

Then you need to choose between the lotteries L_3 and L_4:

$$L_3 = (0, 0.11, 0.89) \qquad L_4 = (0.10, 0, 0.90).$$

In experiments, it is common for people to express the preferences $L_1 \succ L_2$ and $L_4 \succ L_3$. However, these choices are not consistent with expected utility theory. To see why, suppose that there was a v.N-M utility function representing these preferences and denote the utility values of the three outcomes by u_1, u_2 and u_3, respectively. Then, the choice $L_1 \succ L_2$ implies

$$u_2 > 0.10u_1 + 0.89u_2 + 0.01u_3.$$

If we add $0.89u_3 - 0.89u_2$ to both sides of the above inequality, we obtain

$$0.11u_2 + 0.89u_3 > 0.10u_1 + 0.90u_3,$$

and therefore any individual with a v.N-M utility function must have $L_3 \succ L_4$.

The Ellsberg Paradox

There are two urns, each containing 100 marbles, which are either red or blue. Urn 1 contains 49 red marbles and 51 blue marbles. Urn 2 contains an unspecified assortment of marbles. There are two choice situations and in each one you need to select one of the urns, from which a marble will be drawn at random. In choice situation one, you have a chance of winning \$1,000 if the marble drawn is blue and you win nothing if the marble drawn is red. In choice situation two, you will get nothing if the marble drawn is blue and \$1,000 if the marble drawn is red. Which urn will you choose in choice situation one? What about choice situation two?

With the information given most people would pick Urn 1 in choice situation one. This must mean that the perceived (or subjective) probability of a blue marble in Urn 2 is less than 0.51, which in turn means that the perceived probability of a red marble in Urn 2 is more than 0.49. Therefore, people who picked Urn 1 in choice situation one should choose Urn 2 in choice situation 2. However, it turns out that this does not happen overwhelmingly in actual experiments with most people choosing Urn 1 in both situations. This has something to do with the probabilities of Urn 1 being well understood and, in a way, “safe”, while the uncertainty associated with Urn 2 is shunned by the decision makers.

This example demonstrates that it is important to distinguish between *risk* and *ambiguity* according to whether the probabilities are given to us objectively or not.

5 Money Lotteries and Risk Aversion

In many economic settings, individuals seem to display aversion to risk. In this section, we formalize the notion of *risk aversion* and study some of its properties. For the rest of this chapter we will concentrate on lotteries whose outcomes are amounts of money. Notice that the theory we developed thus far relied on a finite outcome space. It is useful, however, to treat money as a continuous variable, and so we will extend the theory to an infinite set of outcomes.

We will denote different amounts of money by the continuous random variable x. Hence, a lottery can be described by means of a *cumulative distribution function* $F : \mathbb{R} \to [0, 1]$.

That is, for any x, $F(x)$ gives the probability that the realized outcome is less than or equal to x. If the distribution function of a lottery has a density function $f(\cdot)$ associated with it, then $F(x) = \int_{-\infty}^{x} f(t)dt$ for all x. The advantage of using distribution functions rather than density functions to formalize ideas is that the first is more general and does not exclude a priori the possibility of a discrete set of outcomes

It is important to note that distribution functions (and also density functions) preserve the linear structure of lotteries. For example, the final distribution of money, $F(\cdot)$, induced by a compound lottery $(L_1, \ldots L_K; \alpha_1, \ldots, \alpha_K)$ is just the weighted average of the distributions induced by each of the individual lotteries that constitute it: $F(x) = \sum_k \alpha_k F_k(x)$ where $F_k(\cdot)$ is the distribution of the payoff under lottery L_k.

From this point on, we will work with distribution functions to describe lotteries over monetary outcomes. Therefore, we define the lottery space $\mathscr{L}$ as the set of all distribution functions over nonnegative amounts of money. The application of the expected utility theorem to outcomes defined by a continuous variable tells us that under the assumptions of the theorem, there is an assignment of utility values $u(x)$ to non-negative amounts of money with the property that any $F(\cdot)$ can be evaluated by a utility function $U(\cdot)$ of the form

$$U(F) = \int u(x)dF(x).$$

The v.N-M expected utility function $U(\cdot)$ is the expectation, over the realizations of x, of the values of the Bernoulli utility function $u(x)$. Note that, as before, $U(\cdot)$ is linear in the "probabilities" $F(\cdot)$.

The strength of the expected utility representation is that it preserves the useful expectation form while allowing for the utility of monetary lotteries to depend not only on the mean but also on the higher moments of the distribution of the monetary payoffs $F(\cdot)$. To a large extent, the analytical power of the expected utility formulation hinges on specifying the Bernoulli utility function $u(\cdot)$ in such a manner that it captures interesting economic attributes of choice behavior. In the current monetary context we will assume that $u(\cdot)$

increasing, continuous, and bounded.

We next turn to the formulation of different risk attitudes in terms of the Bernoulli utility function $u(\cdot)$.

Definition 6. *A decision maker is*

- **risk averse** *(***strictly risk averse***) if for any lottery $F(\cdot)$, the degenerate lottery that yields the amount $\int x dF(x)$ with certainty is at least as good as (strictly better than) the lottery $F(\cdot)$ itself.*
- **risk neutral** *if for any lottery $F(\cdot)$ he is always indifferent between the above two lotteries.*
- **risk loving** *(***strictly risk loving***) if any lottery $F(\cdot)$ is at least as good as (strictly better than) the degenerate lottery that yields the amount $\int x dF(x)$ with certainty.*

Notice that for strictly risk averse and strictly risk loving decision makers, indifference between $F(\cdot)$ and degenerate lottery that yields the amount $\int x dF(x)$ with certainty holds only if the two lotteries are the same, i.e. if $F(\cdot)$ is degenerate.

For preferences that admit an expected utility representation with Bernoulli utility function $u(x)$, it follows directly from the definition of risk aversion that the decision maker is risk averse if and only if

$$\int u(x) dF(x) \leq u\left(\int x dF(x)\right) \quad \text{for all} \quad F(\cdot). \tag{1}$$

Inequality 1 is called *Jensen's inequality.* It constitutes the defining property of a concave function. Hence, in the context of expected utility theory, risk aversion is equivalent to the concavity of $u(\cdot)$ and strict risk aversion is equivalent to the strict concavity of $u(\cdot)$.

This all makes perfect sense. Strict concavity of $u(\cdot)$ means that the marginal utility of money is decreasing. Hence, at any level of wealth x, the utility gain from an extra dollar is smaller than the absolute value of the utility loss of having a dollar less. This in turn implies that the risk of gaining or loosing a dollar with equal probability is not worth taking.

The next definition introduces some useful concepts for the analysis of risk aversion.

Definition 7. *Given a Bernoulli utility function $u(\cdot)$ we define the* **certainty equivalent of** $F(\cdot)$*, denoted $c(F,u)$, as the amount of money for which the individual is indifferent between the gamble $F(\cdot)$ and the certain amount $c(F,u)$; that is,*

$$u(c(F,u)) = \int u(x)dF(x).$$

The **risk premium of** $F(\cdot)$*, denoted as $rp(F,u)$, is defined as the difference between the expected payoff under $F(\cdot)$ and the certainty equivalent $c(F,u)$; that is:*

$$rp(F,u) = \int xdF(x) - c(F,u).$$

Proposition 8. *Suppose a decision maker is an expected utility maximizer with a Bernoulli utility function $u(\cdot)$ on amounts of money. Then the following properties are equivalent:*

(i) The decision maker is risk averse.

(ii) $u(\cdot)$ is concave.

(iii) $c(F,u) \leq \int xdF(x)$ for all $F(\cdot)$.

(iv) $rp(F,u) \geq 0$ for all $F(\cdot)$.

Insurance

Consider a strictly risk averse decision maker who has an initial wealth of w dollars, but who runs a risk of a loss of the amount of D dollars. The probability of the loss is given by π. It is possible, however, for the decision maker to buy insurance. One unit of insurance costs q dollars and pays 1 dollar if the loss occurs. Thus, if α units of insurance are bought, the wealth of the individual will be $w - \alpha q$ if no loss occurs, which happens with probability $1 - \pi$, and $w - \alpha q - D + \alpha$ if the loss occurs, which happens with probability π. Therefore, the decision maker's expected wealth is given by

$$w - \pi D + \alpha(\pi - q).$$

We will assume that $Dq < w$ so that full insurance is affordable. The decision maker's problem is to choose the optimal level of insurance α^*. His utility maximization problem is

$$\max_{\alpha} \ (1-\pi)u(w - \alpha q) + \pi u(w - \alpha q - D + \alpha)$$

subject to the constraints $\alpha \geq 0$ and $\alpha q \leq w$. The Lagrangian can then be written as

$$\mathbb{L} = (1-\pi)u(w-\alpha q) + \pi u(w - \alpha q - D + \alpha) + \mu\alpha + \lambda(w - \alpha q)$$

where $\lambda \geq 0$ and $\mu \geq 0$ are the Lagrange multipliers. The first order conditions (FOCs) are given by

$$\frac{\partial \mathbb{L}}{\partial \alpha} = -q(1-\pi)u'(w - \alpha q) + (1-q)\pi u'(w - \alpha q - D + \alpha) + \mu - \lambda q = 0$$

$$\frac{\partial \mathbb{L}}{\partial \mu} = \alpha \geq 0; \quad \mu \geq 0; \quad \alpha\mu = 0$$

$$\frac{\partial \mathbb{L}}{\partial \lambda} = w - \alpha q \geq 0; \quad \lambda \geq 0; \quad \lambda(w - \alpha q) = 0$$

Consider $q = \pi$. This is the case of *actuarially fair* insurance, since the price of one unit of insurance is equal to the expected cost of insurance, which means that the insurance company breaks even. Then, the FOC w.r.t. α becomes

$$\frac{\partial \mathbb{L}}{\partial \alpha} = -\pi(1-\pi)\big[u'(w - \alpha\pi) - u'(w - D + (1-\pi)\alpha)\big] + \mu - \lambda q = 0$$

which needs to hold at α^*.

Suppose $\alpha^* = 0$. Then, it must be that $\lambda^* = 0$ and the FOC w.r.t. α becomes

$$\pi(1-\pi)\big[u'(w) - u'(w-D)\big] = \mu.$$

However, since $u(\cdot)$ is strictly concave, we have that $u'(w-D) > u'(w)$. Thus, we get $\mu < 0$, which violates one of the conditions w.r.t. μ. Therefore, we must have $\alpha^* > 0$, which in turn implies $\mu^* = 0$.

Suppose $\alpha^* = \frac{w}{q} = \frac{w}{\pi}$. Then, the FOC w.r.t. α becomes

$$-\pi(1-\pi)\big[u'(0) - u'(\frac{w}{\pi} - D)\big] = \lambda q.$$

Since $u(\cdot)$ is strictly concave, we have that $u'(0) > u'(\frac{w}{\pi} - D)$ (notice that $\frac{w}{\pi} - D > 0$ by the initial assumption that full insurance is affordable). Thus, we get $\lambda < 0$, which violates one of the conditions w.r.t. λ. Therefore, we must have $\alpha^* < \frac{w}{q} = \frac{w}{\pi}$, which in turn implies $\lambda^* = 0$.

We so far showed that $\alpha^* \in (0, \frac{w}{p})$, i.e. we have an interior solution. The FOC w.r.t. α becomes:

$$u'(w - \alpha\pi) = u'(w - D + (1-\pi)\alpha).$$

Since $u(\cdot)$ is strictly concave, we must have $w - \alpha\pi = w - D + (1-\pi)\alpha \Leftrightarrow \alpha^* = D$. Thus, if *insurance is actuarially fair, the decision maker insures completely.* The individual's final wealth is thus $w - \pi D$ regardless of the occurrence of the loss. Notice that by purchasing this insurance the agent has not changed his expected wealth, which is the same regardless of whether he purchases insurance or not. However, by purchasing insurance, he has reduced the variance of his wealth to zero. Given the same level of expected wealth, a risk averse individual always wants to reduce the variance of his wealth as much as possible.