

Notes on choice under uncertainty.

These notes outline the theoretical part of the uncertainty discussion in this class. You should also consult the Uncertainty chapters in Jehle and Reny, in Varian's *Microeconomic Analysis* and in Nicholson's book, as well as your own lecture notes.

1. Probability:

X - **state space** (what may occur)

In this class we, primarily, deal with money lotteries (i.e. $X \subset \mathbb{R}$). In what follows, unless otherwise defined, every state is a real number $x \in X$, which denotes how much money a person has.

$E \subset X$ - **event** (e.g., the person has between \$2 and \$3)

Probability is a function P that assigns to each possible event a number such that

i) $0 \leq P(E) \leq 1$

ii) $P(X) = 1$

iii) For any two disjoint events ($A \cap B = \emptyset$) $P(A \cup B) = P(A) + P(B)$

(Note: strictly speaking, property (iii) has to be strengthened slightly to hold for countably infinite unions, but for the purposes of this class we will never need this).

An immediate consequence:

$$P(X \setminus A) = 1 - P(A) \text{ (in particular, } P(\emptyset) = 0)$$

Note: we say that event E occurs with **positive** probability if $P(E)$ is **strictly bigger** than zero.

Comment: in economics we often think of probabilities as **subjective**: probability of an event is *what an individual believes it to be*.

2. Lotteries

Let $X = \{x_1, x_2, \dots, x_N\}$ be the **state space** and $P = (p_1, \dots, p_N)$ a **probability distribution**. We call the pair $L = \{X; P\} = \{x_1, x_2, \dots, x_N; p_1, p_2, \dots, p_N\}$ a **lottery** (some books use the term **gamble** instead).

Examples:

A coin toss: $L = \{Head, Tail; \frac{1}{2}, \frac{1}{2}\}$

Throwing dice: $L = \{1, 2, 3, 4, 5, 6; \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}\}$

Lottery from the St. Petersburg paradox (due to Bernoulli): $L = \{2, 4, 8, 16, \dots, 2^n, \dots; \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \dots, \frac{1}{2^n}, \dots\}$

If states are real numbers ($X \subset \mathbb{R}$) we can define *expected value of a lottery* (for X consisting N elements) as

$$E(X|L) = \sum_{i=1}^N p_i x_i$$

Note: expected value may not always exist. *Example:* try computing the expected value for the St. Petersburg paradox lottery.

From now on we shall assume that X is fixed and known. Given a fixed state space X we shall be interested in the set $\mathcal{L}(X)$ of all possible lotteries over X . When X is known, we may suppress it in writing and just identify the lottery with a probability distribution $L = \{p_1, p_2, \dots, p_N\}$ (notice, this is just an n -dimensional vector)

Besides the **simple** lotteries defined above, we may want to consider **compound** lotteries, *i.e.* lotteries over lotteries. For instance, what if individual is not sure, which of the lotteries $L = \{p_1, p_2, \dots, p_N\}$ or $L' = \{p'_1, p'_2, \dots, p'_N\}$ he/she is facing, but assigns a probability α to it being L (and, correspondingly, $(1 - \alpha)$ to it being L'). Consider the resulting *compound* lottery

$$L'' = \{L, L'; \alpha, 1 - \alpha\}$$

It is not hard to see (why?), that L'' can actually be viewed as a *simple* lottery

$$L'' = \{\alpha p_1 + (1 - \alpha) p'_1, \alpha p_2 + (1 - \alpha) p'_2, \dots, \alpha p_N + (1 - \alpha) p'_N\} = \alpha L + (1 - \alpha) L'$$

A compound lottery is, thus, simply an average of simple lotteries (note, that an average of the lotteries is still a lottery, while the sum of lotteries is not - think why). Nonetheless, it is sometimes convenient to use the compound lottery interpretation, as we shall see shortly.

3. Preferences over lotteries.

Consider space X and the corresponding lottery space $\mathcal{L}(X)$. We are interested in the way individuals choose actions, when these choices result in uncertain

consequences - i.e., in the way individuals choose among lotteries.. The standard hypothesis in economics is to assume that individuals have preferences \succsim over the lottery space $\mathcal{L}(X)$ (we say $L \succsim L'$ - lottery L is at least as good as lottery L' ; think of these in the same way as you thought about preferences over states).

As in the case of preferences over certainty, define the properties of preferences reflexivity, transitivity and completeness (rationality). Recall, that completeness means that the individual can compare any two lotteries, while transitivity means that if, in his view, one lottery is at least as good as another, which, in turn, is at least as good as the third, then this individual will rank the first lottery as being at least as good as the third. Since the formal definitions are, essentially, identical to the certainty case we shall skip them.

4. Utility function representation and Expected Utility.

As in the certainty case, given rationality of preferences \succsim , together with an additional regularity assumption (“continuity”), we can conclude that there exists a utility function

$$U : \mathcal{L}(X) \rightarrow \mathbb{R}$$

which represents these preferences, i.e.

$$L \succsim L' \Leftrightarrow U(L) \geq U(L')$$

However, if the theory is to be workable, we want to be able to take expectations of utilities. However, consider the “no uncertainty” lotteries $L_1 = \{1, 0, 0, \dots, 0\}$, $L_2 = \{0, 1, 0, 0 \dots 0\}$, ..., $L_N = \{0, 0, \dots, 1\}$. It seems reasonable to view utility of such a lottery as simply utility of having the corresponding state for sure. Let us denote

$$u(x_k) = U(L_k)$$

It turns out, however, that IT IS NOT generally true that utility of a lottery is equal to the expected utility of the states, given the probability distribution of the lottery.

The expected utility theory which is sketched below was developed by John von Neumann and Oscar Morgenstern in their book *Game Theory and Economic Behaviour*.

Assumption (independence axiom): for any lotteries L, L', L'' and for any α between zero and one

$$L \succsim L' \Leftrightarrow \alpha L + (1 - \alpha) L'' \succsim \alpha L' + (1 - \alpha) L''$$

Interpretation of independence: suppose, when comparing L and L' I choose L ; now suppose that, whether I choose L or L' , with some probability $(1 - \alpha)$ I will be facing a different lottery L'' . Independence says, that this additional uncertainty *which is the same whatever my choice is*, should not matter for my choice.

Von Neumann and Morgenstern showed that under assumptions of rationality, continuity AND independence it is possible to choose a particular - utility function U , which represents the preferences AND has an expected utility form:

$$U(L) = \sum_{i=1}^N p_i u(x_i) = E(u(X)|L)$$

The "small" u (defined over the STATES) is known as **Bernoulli utility function**, while the "big" U (defined over the lotteries) is known as **von Neumann - Morgenstern (vNM) utility function**. (Note, that the nomenclature is not yet settled - many books use a different terminology, actually calling u **vNM utility function**, while U - **expected utility function**).

Knowledge of an individual's Bernoulli utility function u over the states, uniquely defines his/her preferences over the larger domain of the lotteries. Since it carries a lot more information than an arbitrary utility function over X , a monotonic transformation may, actually, distort some of this information. It is only when the two utility functions are *linear* monotonic transformations of one another:

$$w(x) = au(x) + b$$

where $a > 0$, b - any real number, that they give rise to representations of the same preferences over uncertainty.

Comment: since we are primarily dealing with lotteries over amounts of money, and since it is not unreasonable to assume that without uncertainty everyone would prefer more money to less money, we shall always assume in what follows, that all Bernoulli utility functions are strictly increasing.

5. Attitude towards risk. Risk Aversion.

In fact, the concavity or convexity of a Bernoulli utility function (properties that are not generally preserved by arbitrary monotonic transformations but are preserved by the linear transformations) actually carries a lot of information about individual attitude to risk. In fact, the mathematical property known as **Jensen's inequality** tells us that for any concave u

$$E(u(X)|L) = \sum_{i=1}^N p_i u(x_i) \leq u\left(\sum_{i=1}^N p_i x_i\right) = u(E_L(X)|L)$$

In words: an individual, whose preferences are represented by a concave Bernoulli utility function prefers to have the amount of money equal to the expected value of the lottery to having the lottery itself. An individual like this is known as **risk-averse**. The inequality is reversed for convex u - such individuals are called **risk-loving**. Finally, the inequality becomes *equality* when u is linear (e.g., $u(x) = x$). Such individual cares only about the expected value of the lottery, and is indifferent to risk; he/she is known as **risk-neutral**. We shall primarily consider risk-averse individuals in what follows.

A risk-averse individual would, generally, exchange a lottery for a fixed sum of money somewhat smaller than the lottery's expected value. This amount of money, denoted $c(L; u)$, given by the equality

$$u(c(L; u)) = E(u(X)|L)$$

the individual's **certainty equivalent** of the lottery. The smaller is the certainty equivalent, the more the individual is willing to pay to get rid of uncertainty.

Consider two individuals, with preferences represented using Bernoulli utility functions u and w , respectively. If, for every lottery L it turns out that $c(L; u) \leq c(L; w)$ it is reasonable to say that the individual with Bernoulli utility function u is more risk-averse than the one with Bernoulli utility function w .

Alternatively, we may want to say that an individual with a "more concave" utility function is the more risk-averse one. One way to formally define this is to say that u is more concave than w if it is a concave transformation of w , *i.e.*

$$u(x) = \varphi(w(x))$$

for some increasing and concave function φ .

A third way to define "more concave" (and, therefore, more risk-averse) is to define a curvature coefficient of the Bernoulli utility function, which is known as

Arrow-Pratt coefficient of absolute risk aversion:

$$r_A(x; u) = -\frac{u''(x)}{u'(x)}$$

Fortunately, it turns out that the three definitions of “more risk-averse” coincide. In other words, the following three statements are equivalent:

- (i) $c(L; u) \leq c(L; w)$ for every lottery L
- (ii) $u(x) = \varphi(w(x))$ for some increasing concave φ
- (iii) $r_A(x; u) \geq r_A(x; w)$ for every state x in X

If you find that any of the above holds you may conclude that individual with B.u.f. u is more risk averse than the one with B.u.f w .

You should keep in mind, that this is only an incomplete ranking: it is possible that you can't compare two risk-averse individuals in this way. However, within an important subclass of utility functions, known as **constant absolute risk-aversion (CARA) utility functions**, the ranking is complete. These are the utility functions of the form

$$u(x) = -ae^{-rx} + b$$

For these, the coefficient of absolute risk-aversion

$$r_A(x; u) = r$$

The problem with CARA utility functions is, that such individuals are concerned only with the *absolute* size of the risk they are exposed to. Whether a possible loss of \$1000 represents .5%, 5%. or 50% of his net worth, an individual with CARA preferences will view it in the same manner. If you believe, that the possibility of a loss of \$1000 for Bill Gates now that he has his billions, is less unpleasant, than when he was a poor Harvard drop-out, you believe that his preferences exhibit decreasing, not constant absolute risk-aversion.

Therefore, it may make sense to consider the Arrow-Pratt **coefficient of relative risk-aversion**

$$r_R(x; u) = -x \frac{u''(x)}{u'(x)}$$

which focuses on the size of the gamble as a proportion of one's wealth. Constant relative risk-aversion implies decreasing absolute risk-aversion (think why?)

You will commonly encounter the *constant relative risk-aversion (CRRA)* utility functions

$$u(x) = \frac{x^{1-\sigma} - 1}{1-\sigma}, 0 < \sigma < 1$$

$$u(x) = \ln(x), \sigma = 1$$

where

$$r_R(x; u) = \sigma$$

6. Ranking of lotteries by payoff (FOSD) and by risk (SOSD).

When can we unambiguously say that one lottery gives a higher payoff than another? Clearly, simply considering expected value of the lottery is insufficient (why?), and expected utilities reflect necessarily subjective individual preferences.

We say that lottery L has higher payoff than lottery L' (L **First Order Stochastically Dominates (FOSD)** L') if every individual who prefers more money to less money agrees that he/she prefers L to L' . In other words

$$E(u(X)|L) \geq E(u(X)|L')$$

for every increasing u .

It turns out, that this is equivalent to saying that lottery L can be obtained from lottery L' by an upward shift of the probability distribution (recall discussion in class). Recalling a definition from your probability theory class, an alternative way of saying the same thing is to say that the distribution function F associated with L is always smaller than or equal to the distribution function G associated with L'

If lotteries have the same expected value we, obviously, can't compare them with respect to payoff. However, we can compare them with respect to risk. We say that lottery L is *less risky* than lottery L' (L **Second Order Stochastically Dominates (SOSD)** L') if their expected values are equal

$$E(X|L) = E(X|L')$$

and every *risk-averse* individual who prefers more money to less money agrees that he/she prefers L to L' . In other words

$$E(u(X) | L) \geq E(u(X) | L')$$

for every increasing and *concave* u .

It turns out, that this is equivalent to saying that lottery L' can be obtained from lottery L by a *mean-preserving spread* of the probability distribution (once again, recall discussion in class).