

Pol. Sci. 217
Lecture #3
Expected Utility

This far we have constrained ourselves to the study of individual, static, decision problems under certainty. In this lecture we will take the first (and, from the point of view of the present course, perhaps the most important) departure: uncertainty.

Why Uncertainty?

Why is uncertainty such an important departure? Remember that our goal is to build up the analysis of multi-person decision problems. According to the methodological claims of RC, this means that we want a theory that tells us, for any given problem in which many people interact, how each individual chooses her actions, in a given environment. But, with many people acting at the same time, their decisions also become part of the environment. But, since any given individual cannot tell *a priori* what the decisions of the others are, she must treat her environment as uncertain. In that sense, learning to deal with uncertainty is the key link between decision theory and game theory. In some sense, most games, even games where no obvious source of uncertainty exists, require the agents to reason of uncertain outcomes. Of course, this statement is very vague and even inaccurate, but it will become clearer as we move on.

Lotteries

As seen in the last lecture, decision problems are fully described once we know the possible outcomes and the preferences of the decision-maker. This is also true for uncertainty. Only that now we need to modify our definitions, both of outcomes and preferences.

Lotteries are simply uncertain outcomes. When we choose a lottery, we choose a set of possible outcomes, to each of whom we attach a probability of occurring. For convenience, in what follows we will focus on lotteries whose prizes are in money. This is only for convenience and you are encouraged to abandon that framework as soon as it becomes an impediment to your intuition. In particular, later on we will deal with non-monetary decision problems.

In order to introduce the basic concepts, let's use one example. Three different agents are offered a lottery described as follows: A coin will be tossed with probability p for heads (H) and $1 - p$ for tails (T). If H , the gambler will receive \$100, if T \$25. Now we ask, how much should anyone be willing to pay for the opportunity of playing such a lottery? To find out, we need to know what are the *preferences* of the agents about lotteries. But just as lotteries are built out of simple outcomes, preferences over lotteries are also built of preferences over simple outcomes. So, we need to start with the simple preferences. To these we turn next.

Bernoulli Utility Functions

First, a remark: the name Bernoulli utility functions is not commonly used in the literature. In fact, it has been recently coined by a leading Microeconomics textbook. Whether or not it'll stick is anybody's guess. I use it here because I agree with that textbook's authors as to the need to have a very solid understanding of preferences under certainty before moving on to the problem of preferences over lotteries. The name is as good as any other and has the advantage that pays tribute to a prominent (family of) mathematician(s), which is all the more appropriate given that the utility functions over lotteries will also pay tribute to other scientists.

Let the three agents of our example be labeled A, B, C . Furthermore, let their preferences over money be represented by (Bernoulli) utility functions $u_A(x), u_B(x), u_C(x)$, such that:

$$\begin{aligned} u_A(x) &= \sqrt{x} \\ u_B(x) &= x \\ u_C(x) &= x^2 \end{aligned}$$

There is a subtle point involved here. In a departure from last lecture, these functions are not just *ordinal* utility functions. We are not just ordering outcomes from worst to best (in particular, more money is better than less money). We are doing something else: we are giving some meaning to *changes* in preferences as we change the amount of money. For example, with this information, we could answer the following question: suppose we have \$100 and want to give them to the agent that will benefit the most from it. To whom should we give it? (Interpersonal comparisons of utility.)

Von Neumann-Morgenstern Utility Functions

This is the key step. Remember that we are trying to find how much an agent should pay to be allowed to play the lottery. Let's call q_A, q_B, q_C the amounts each agent would have to pay. Then, for each individual, the decision whether or not gamble is the same as comparing two lotteries: one offering, as said before, 100 with probability p and 25 with probability $1 - p$, and another offering q with probability 1. (I drop the dollar sign for esthetic purposes.) Is there a utility function representing the preferences over such lotteries, just as there is a utility function representing the preferences over simple outcomes? The Expected Utility Theorem answers this question in the affirmative and, furthermore, specifies exactly how such utility function looks like.

The Expected Utility Theorem claims that, under some assumptions (discussed in the textbook), the preferences over lotteries are represented by a utility function of the following form:

Let o_1, o_2, \dots, o_n be the outcomes of a lottery \mathcal{L} , occurring with probabilities p_1, p_2, \dots, p_n respectively. If $u(\cdot)$ is the Bernoulli utility function over outcomes, the *von Neumann-Morgenstern utility function* is:

$$EU(\mathcal{L}) = p_1 u(o_1) + p_2 u(o_2) + \dots + p_n u(o_n)$$

With this in hand, let's consider first agent B . If q_B is the maximal amount of money she should pay for the opportunity of playing the gamble, this means that q_B is such that she is indifferent between the two lotteries, that is:

$$\begin{aligned} EU(q_B) &= EU(L) \\ u_B(q_B) &= p u_B(100) + (1 - p) u_B(25) \\ q_B &= 100p + 25(1 - p) \\ q_B &= 75p + 25 \end{aligned}$$

Notice one small point that should give us some assurance that we're on the right track. If $p = 0, q_B = 25$ and, likewise, if $p = 1, q_B = 100$. This is as it ought to be. If the gamble gives no probability whatsoever to winning 100, why should anybody pay more than 25. On the other hand, if the big price of 100 is absolutely sure, then paying anything up to 100 should make sense.

For agent A :

$$\begin{aligned}
EU(q_A) &= EU(L) \\
u_A(q_A) &= pu_A(100) + (1-p)u_A(25) \\
\sqrt{q_A} &= 10p + 5(1-p) \\
\sqrt{q_A} &= 5p + 5 \\
q_A &= 25p^2 + 50p + 25
\end{aligned}$$

Finally, for agent C :

$$\begin{aligned}
EU(q_C) &= EU(L) \\
u_C(q_C) &= pu_C(100) + (1-p)u_C(25) \\
q_C^2 &= 10000p + 625(1-p) \\
q_C^2 &= 9375p + 625 \\
q_C &= \sqrt{9375p + 625}
\end{aligned}$$

Attitudes toward risk

We know that when $p = 1$ or $p = 0$ the three agents should pay the same simply because the lottery is no longer uncertain (becomes degenerate). But, what about intermediate values of p ? I claim that agent C should be the most willing to pay, while A should be the least willing to pay. (We can verify this by plotting q_A , q_B and q_C against p .) For example, for $p = 0.5$, $q_A = 56.25$, $q_B = 62.5$, $q_C \approx 72.89$.

Intuitively, A is less willing to gamble than B , which in turn is less willing to gamble than C . We can make this more precise. Let's first focus on agent B . The expected *value* of the lottery is 62.5, which happens to be exactly the same as B 's expected *utility* from the lottery. That is, B is indifferent between a lottery with expected value of 62.5 and having 62.5 for sure. Technically, B is *risk-neutral*. Intuitively, agent B is made neither better nor worse off by the fact of a reward becoming risky. The situation is different for the other agents. Agent A prefers having 56.25 for sure than playing a lottery with an expected value of 62.5. That is, A is *risk-averse*. Finally, C is willing to pay up to 72.89 for the chance of playing a lottery with an expected value of 62.5. Thus, C is a *risk-loving* agent.

In general, the attitudes toward risk of the agents are defined by the concavity of their Bernoulli utility functions. Concave functions imply risk-aversion, linear functions imply risk neutrality and convexity means that

the agent is risk-loving. Notice, furthermore, that concave utility functions imply, at the same time that the agent is subject to diminishing marginal utility. Incidentally, this explains why risk-aversion is so embedded in economic analysis, while risk-loving is considered more of a rarity.