Decision Making: Utilities

Daphne Koller
Stanford University

December 10, 2002

1 Introduction

In our oil wildcatter example, we assumed that his goal was to maximize profit. In many real-world situations, many other factors enter into preference for one consequence versus another. As a simple example, consider the problem of deciding whether to have a party indoors or outdoors. The consequence of these two actions depends on the whether, e.g., whether it's raining or dry. Given our two actions and the two states of the world, there are four possible consequences:

indoors	dry	regret
indoors	wet	relief
outdoors	dry	perfect
outdoors	wet	disaster

Our preferences between the consequences is clear. However, that is not enough to make a decision. As I'm uncertain about the state of the world, I need to be able to decide between *lotteries* over consequences. E.g., I need to be able to figure out whether I prefer regret with probability 0.7 and relief with probability 0.3 to perfect with probability 0.7 and disaster with probability 0.3. In more complicated problems (e.g., where my actions can actually affect the probabilities with which things happen), these lotteries will be even more complicated. For example, consider a (real-world) decision making situation: giving advice to pregnant women on which prenatal tests to take. Here, you have to ascribe preferences over much more complicated lotteries, e.g., "Downs syndrome baby with probability 0.03, perfect baby with probability 0.9, miscarriage with probability 0.07".

In order to make the decision making process feasible, and certainly in order to automate it, we need to simplify things a little. Specifically, we would like to reduce the problem of reasoning about lotteries to one of reasoning about individual consequences. We would then put together our preferences about individual consequences to get preferences about lotteries. In the earliest days of decision theory, Von Neumann and Morgenstern provided us with axiomatic foundations which justify such a simplification.

2 Decision-Theoretic Foundations

At a high level, Savage proved that if our preferences over lotteries satisfy some basic "rationality" assumptions, then we can duplicate our complicated preferences over lotteries using a much simpler mechanism. Let's look at some of these assumptions. First, some notation. For two lotteries ℓ_1, ℓ_2, if the agent prefers ℓ_1, we say that $\ell_1 \succ \ell_2$. We also allow the agent to be indifferent between the two lotteries, in which case $\ell_1 \sim \ell_2$.

Orderability: An agent must know what he wants. I.e., for all lotteries ℓ_1, ℓ_2, ℓ_3, either

$$(\ell_1 \prec \ell_2) \vee (\ell_1 \succ \ell_2) \vee (\ell_1 \sim \ell_2)$$

This is not a trivial assumption. It's hard for people to come up with preferences over lotteries. For example: Do you prefer a lottery where you get \$4 million with 20% chance or one where you get \$3 million with 25% chance? This is a hard question. In general, it's asking a bit much of the agent to ascribe preferences between complex lotteries.

Transitivity: Preferences are transitive. This one seems fairly unobjectionable, although there are still circumstances where one could argue the other way.

$$(\ell_1 \prec \ell_2) \wedge (\ell_2 \prec \ell_3) \rightarrow (\ell_1 \prec \ell_3)$$

Continuity: This one essentially says that if ℓ_2 is somewhere in the middle for you between ℓ_1 and ℓ_3, then there should be some lottery between ℓ_1 and ℓ_3 which is equivalent to ℓ_2.

$$(\ell_1 \prec \ell_2 \prec \ell_3) \Rightarrow \exists p(\ell_2 \sim [\ell_1 : p; \ell_3 : (1-p)]$$

For the simple example of consequences, we might have that Regret $\sim$ [Perfect : 0.7; Disaster : 0.3].

Monotonicity: You prefer that better things happen with higher probability.

$$(\ell_1 \succ \ell_2), (p \geq q) \Rightarrow ([\ell_1 : p; \ell_2 : (1-p)] \succ [\ell_1 : q; \ell_2 : (1-q)]$$

Substitutability:

$$(\ell_1 \sim \ell_2) \Rightarrow ([\ell_1 : p; \ell_3 : (1-p)] \sim [\ell_2 : p; \ell_3 : (1-p)]$$

Decomposability: You can't prefer lotteries to each other just because one has more "gambling points" (no fun in gambling).

$$[\ell_1 : p, [\ell_2 : q, \ell_3 : (1-q)] : (1-p)] \sim [\ell_1 : p, \ell_2 : (1-p)q, \ell_3 : (1-p)(1-q)]$$

We can now get the following result.

Theorem 2.1: *Assuming that we have an agent whose preferences over lotteries satisfy these axioms, then there exists a function U from consequences to real numbers such that for any $\ell_1 = [y_1 : p_1; \ldots; y_n : p_n]$ and $\ell_2 = [z_1 : q_1; \ldots; z_m : q_m]$, $\ell_1 \prec \ell_2$ if and only if*

$$\sum_i p_i U(y_i) < \sum_j q_j U(z_j)$$

That is, the utility of any lottery is completely determined by the utility and probability of the consequences.

Proof: Our goal is to take a preference relation $\prec$ that satisfies these axioms, and to construct a utility function U over consequences such that $\prec$ is equivalent to implementing the MEU principle over the utility function U. The proof is essentially as follows: We take the least and most preferred consequences c_{min} and c_{max}. By orderability and transitivity, these must exist. We assign $U(c_{min}) := 0$ and $U(c_{max}) := 1$. By orderability, we have that for any other consequence c:

$$c_{min} \prec c \prec c_{max}$$

By continuity, there must exist a probability p such that

$$[c:1] \sim [c_{min}:(1-p); c_{max}:p]$$

We assign $U(c) := p$. The axioms can then be used to show that the assignment of utilities to lotteries resulting from applying the expected utility principle results in an ordering which is consistent with our preferences. ∎

From an operational perspective, this discussion gives us a formal justification for the principle of maximum expected utility, even in general situations where our consequences are complex things. When we have a set of consequences, we ascribe a numerical *utility* to each one. If we have a set of actions that induce different lotteries over our consequences, we should choose the action whose *expected utility* is largest; this is equivalent to choosing the action that induces the lottery which we prefer the most.

Returning to our example of the party, we might have

regret	7
relief	8
perfect	10
disaster	0

Then

$$\mathrm{EU}(indoors) = 0.7 \cdot 7 + 0.3 \cdot 8 = 7.3$$

while

$$\mathrm{EU}(outdoors) = 0.7 \cdot 10 + 0.3 \cdot 0 = 7$$

Therefore, we will choose to have our party indoors.

3 Real-life utilities

What are people's utility functions like in real-life settings?

3.1 Utility of money

In situations involving a monetary element, it is tempting to think of utilities as money. Is this correct? Assume that someone offers you the option of getting \$3 million for sure, or getting \$4 million with probability 0.8. Most people pick the first.

This does not mean that you are irrational. It just means that the utility function is not exactly money. While it is usually true that our utility is monotonic in money (more money is better), the exact shape of the graph is not quite a straight line. Let c_0 represent your current financial status. Assume (purely for simplicity) that you assign a utility of 0 to c_0. If you assign a utility of 10 to the consequence $c_0 + 3million$ and 12 to the consequence $c_0 + 4million$. Then the expected utility of the gamble is $0.2 \cdot 0 + 0.8 \cdot 12 = 9.6 < 10$. Therefore, *with this utility function* your decision is completely rational.

This observation, that utility is not linear in money, was originally made by Bernoulli. He designed an example called the St. Petersburg paradox.

Example 3.1: Suppose you are offered a chance to play a game where a fair coin is tossed repeatedly until it comes up heads. If the first head appears on the nth toss, you get $\$2^n$. How much would you be willing to pay in order to play this game?

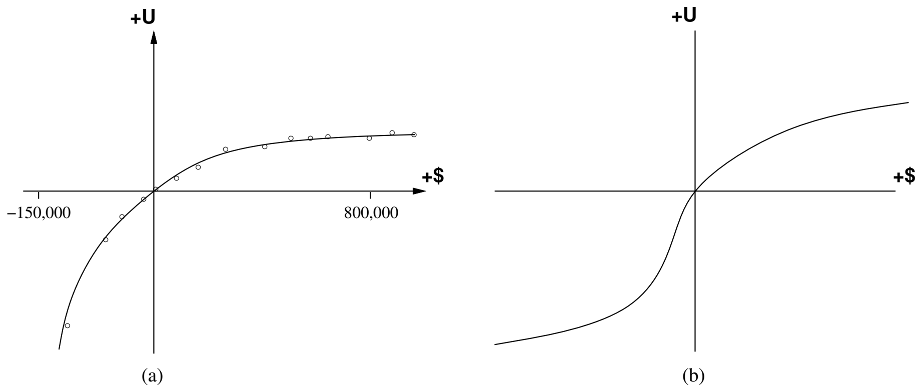

Figure 1: The utility of money (a) Empirical data for a certain Mr. Beard over a limited range. (b) A typical curve for the full range.

The probability of the event H_n—the first head showing up on the nth toss—is $1/2^n$. Therefore, your expected winnings are:

$$\sum_n P(H_n) \mathit{Payoff}(H_n) = \sum_n \frac{1}{2^n} 2^n = 1 + 1 + 1 + \ldots = \infty.$$

Therefore, you should be willing to pay any amount to play this game. This is clearly false for most people. ∎

Bernoulli posited that the utility of money was logarithmic. That is, that the utility of the outcome c_k, corresponding to your current financial status plus $\$k$, is proportional to $\log k$. In this case, we get

$$\sum_n P(H_n) U(\mathit{Payoff}(H_n)) = \sum_n \frac{1}{2^n} U(c_{2^n}) = \sum_n \frac{n}{2^n}$$

which turns out to be 2, the number that most people are willing to pay.

Empirical psychological studies show that people's utility functions in a certain range are often are logarithmic with money. One person's utility function, from a study done in 1960, is shown in Figure 1(a). The empirical data are almost a perfect fit to a utility function:

$$U(c_n) = -263.31 + 22.09 \log(n + 150000)$$

3.2 Attitudes towards risk

The shape of the curve describes the agent's attitude towards risk. A convex curve, like this one, says that the agent is *risk averse*: he prefers a sure thing to a gamble with the same payoff. Consider the risk-averse curve of Figure 2. We see that the utility of a lottery such as $[\$1000 : 0.5, 0 : 0.5]$ is lower than the utility of getting \$500 for sure. Indeed, empirical research shows that the utility of this lottery is pretty much the same as the utility of getting \$400 for sure. This amount is called the *certainty equivalent* of the lottery. It is the amount of "sure-thing" money that people are willing to trade for a lottery. The difference between the expected monetary reward of \$500 and the certainty equivalent of \$00 is called the *insurance premium*, and for a good reason. The premium people pay to the insurance company is precisely to guarantee a sure thing (a sure small loss) as opposed to a lottery where one of the consequences involves a large negative utility (the price of rebuilding the house if it burns down).

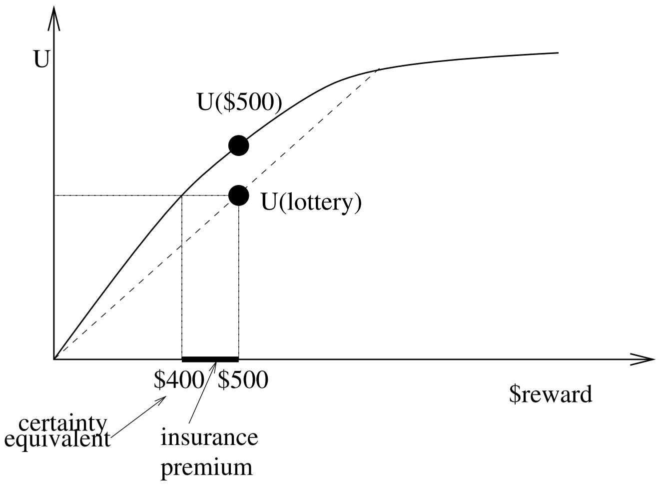

Figure 2: Risk averse utility curve.

People are typically risk averse. But they are often risk seeking when the certain loss is small (relative to their financial situation). Indeed, state lotteries exploit precisely this phenomenon. When the agent prefers the lottery to the certainty equivalent, he is said to be *risk seeking*; this will be the case when the curve is concave. Finally, if the agent's utility curve is linear in money, he is said to be *risk neutral.* Most utility curves are locally linear, which means we can assume risk neutrality for small risks and rewards. We also note that people are rarely consistent about risk throughout the entire monetary range. In fact, many people's utility curve changes to risk-seeking at large negative amounts. Going into debt is usually considered disastrous, so there is a steep slope where our monetary status turns negative. However, preferences between different levels of debt can be the reverse of preferences associated with levels of wealth. For example, someone who is already $10 million in debt might choose to accept a gamble on a fair coin with $10 million payoff on "heads" and a $20 million loss on tails. While this might be desperate, it is not irrational. It simple reflects the person's utilities for being $10 million in debt, $30 million in debt, and having no debt at all.

3.3 Rationality

Decision theory allows people to have a completely arbitrary utility function, one which is risk averse, risk seeking, or risk neutral, and even one that changes over the range. Does that mean that people are rational with respect to some utility function? The answer is no. People's preferences are rarely consistent with a utility function. For example, we asked two questions during class:

- Which do you prefer: a lottery (A) where you get $4 million with 20% chance or one (B) where you get $3 million with 25% chance?
- Which would you prefer: (C) $4 million with 80% probability or (D) $3 million for sure?

In experiments by Tversky & Kahneman [1982], most people prefer (A) to (B), but (D) to (C). These preferences are not consistent with any utility function. To see that, assume (for simplidity) that $U(\$0) = 0$. In this case, the preference of (D) to (C) implies that

$$0.8 \cdot U(c_{4,000,000}) < U(c_{3,000,000}).$$

On the other hand, the preference of (A) to (B) implies that

$$\begin{aligned} 0.2 \cdot U(c_{4,000,000}) &> 0.25 \cdot U(c_{3,000,000}) \\ 0.8 \cdot U(c_{4,000,000}) &> 1 \cdot U(c_{3,000,000}) \end{aligned}$$

Thus, people's preferences are rarely rational.

The interesting thing is that, after participating in an experiment like this, people's preferences often change. For example, Keeney & Raiffa (Harvard Business School) report that "Most of the subjects, however, can reconcile their inconsistencies and feel that they have learned an important lesson about how they want to behave. As a consequence, some subjects cancel their automobile collision insurance and take out more term insurance on their lives."

3.4 Multi-attribute utilities

In real life, money is not the only thing. For example, in a medical decision making problem, there is risk of death, quality of life issues, cost of various treatments, etc. We have to factor in all of these aspects into a single utility function. Thus, if our utility depends on several variables $Z_1, \ldots, Z_k$, we have to define a utility function $U : Val(\mathbf{Z}) \mapsto I\!R$. As usual, the size of this representation is exponential in k.

It is possible to consider structured models of utility functions, which allow us to decompose this exponential representation. The simplest model is that of *additive independence*. It assumes that we can represent the utility function U as a sum:

$$U(Z_1, \ldots, Z_k) = U_1(Z_1) + \cdots + U_k(Z_k).$$

For example, if we are planning a vacation, our utility might depend on the cost C, the length of the flight L, the weather at the destination W, and the food quality F. It might be possible to decompose the utility function as:

$$U(C, L, W, F) = U_1(C) + U_2(L) + U_3(W) + U_4(F)$$

Unfortunately, the situation is rarely that nice. For example, our utility for the vacation would depend not only on the weather, but also on the type of activity available A. Clearly, if the type of activity is hiking, we really prefer that it not rain; on the other hand, if the primary activity is going to museums, we might not care as much. Thus, we cannot decompose the utility as a simple sum over the attributes.

We can generalize this decomposition in several ways. For example, we can decompose additively over *clusters* of variables:

$$U(C, L, W, F) = U_1(C) + U_2(L) + U_3(W, A) + U_4(F)$$

There are also more general decompositions, called *conditional additive independence* and *generalized additive independence*. For example, consider the three attributes corresponding to the hotel room quality Q, the availability of a good beach B, and the weather W. It might be the case that our utility for different beach qualities depends on the weather, as does our utility for different room qualities, but given that the weather is fixed, there is no interaction between B and Q. This example is an instance of conditional additive independence, and corresponds to a decomposition of

$$U(Q, W, B) = U(Q, W) + U(W, B)$$

Deciding on the actual numerical values of utilities is quite hard in practice. Consider the problem of medical decision making. There, different procedures carry different risks to the patient. In order to compare them to each other, and factor in other things like quality of life and costs of the procedures, so you have to ascribe a utility to the loss of human life. People hate to do that, but in real life, tradeoffs are made all the time. Airplanes are not overhauled after each trip, even though that would probably improve safety. Not all cars are made with antilock brakes and airbags, even though they are known to save lives. In fact, the refusal the "put a monetary value on life" often means that life is undervalued. Ross Shachter tells a story about a government agency that commissioned a study on removing asbestos from schools. The study assumed a particular dollar value for the life of a child, and argued that the rational choice under that assumption was to remove the asbestos. The government agency was morally outraged at the idea that a value can be placed on human life, and rejected the report out of hand. They then decided against removing the asbestos.

There has been some work done on placing a value on human life. The two main metrics used are *micromorts*—a 1 in a million chance of death, and $QALY$—quality-adjusted life year (a year in good health with no infirmities). A number of studies across a wide range of individuals have shown that a micromort is worth about $20 in 1980 dollars.