Game Theory

Lecture Notes By

Y. Narahari

Department of Computer Science and Automation

Indian Institute of Science

Bangalore, India

August 2012

Chapter 7: Von Neumann - Morgenstern Utilities

Note: *This is a only a draft version, so there could be flaws. If you find any errors, please do send email to* `hari@csa.iisc.ernet.in`. *A more thorough version would be available soon in this space.*

Utilities play a central role in game theory. They capture the preferences the players have for different outcomes in terms of real numbers thus enabling real-valued functions to be used in game theoretic analysis. The utility theory developed by Von Neumann and Oskar Morgenstern provides the foundation for using utilities to represent preferences. This chapter introduces their utility theory.

1 Introduction

The outcomes in a strategic form game are typically n-dimensional vectors of strategies, where n is the number of players. Suppose X is the set outcomes in a given game. Each player has preferences on the different outcomes which can be expressed formally in terms of a binary relation called *preference relation* defined on X. The utility function of the player maps the outcomes to real numbers, so as to reflect the preference the player has for these outcomes. For example, in the case of BOS game, the set of outcomes is

$$X = \{(A, A),\ (A, B),\ (B, A),\ (B, B)\}$$

The utility function of player 1 is

$$u_1(A, A) = 2;\ u_1(A, B) = 0;\ u_1(B, A) = 0;\ u_1(B, B) = 1$$

The utility function of player 2 is

$$u_1(A, A) = 1;\ u_1(A, B) = 0;\ u_1(B, A) = 0;\ u_1(B, B) = 2$$

The real numbers $2,\ 0,\ 0,\ 1$ above capture the preference level the players have for the four outcomes of the game. Note that the utility function is a single dimensional function that maps (possibly complex) multi-dimensional information into real numbers to capture preferences. The question arises whether it is possible at all to capture all the preferences without losing any information. Utility theory deals with this problem in a systematic and scientific way. There are many different utility theories which have been developed over the last century. The theory developed by von Neumann and Morgenstern [1] is one of the most influential among these and certainly the most relevant for game theory. In this chapter, we undertake a study of various issues involved in coming up with a satisfactory way of defining utilities in a game setting. The discussion is based on the development of these ideas in the books by Straffin [2], Shoham and Leyton-Brown [3], and Myerson [4].

2 Ordinal Utilities

Consider a game with two players 1 and 2 and four outcomes, $X = \{u, v, w, x\}$. Suppose player 1 prefers u the most, followed by v, w, and x in that order. Let us denote this by, $u \succ v \succ w \succ x$. Assume that player $2's$ preferences are exactly the reverse, that is $x \succ w \succ v \succ u$. If it is required to assign real numbers to these outcomes to reflect the precedence ordering, then there are innumerable ways. One possible immediate assignment would be:

$$\begin{array}{ll} \text{Player 1}: & u:4;\ v:3;\ w:2;\ x:1 \\ \text{Player 2}: & u:-4;\ v:-3;\ w:-2;\ x:-1 \end{array}$$

Clearly, there will exist an uncountably infinite number of utility functions $u_1 : X \rightarrow \Re$ and $u_2 : X \rightarrow \Re$ that represent the preferences of players 1 and 2, respectively. A scale on which larger numbers represent more preferred outcomes in a way that only the order of the numbers matters and not their absolute or relative magnitude is called *ordinal scale*. Utility numbers determined from preferences in this way are called *ordinal utilities*.

3 Cardinal Utilities

A utility scale such that not only the orders of numbers but also the ratios of differences of numbers is meaningful is called an *interval scale*. Numbers which are chosen according to an interval scale reflecting the underlying preferences are called *cardinal utilities*.

Cardinal utilities are necessitated because we need to meaningfully capture the notion of mixed strategies. For example, consider the zero-sum game shown in Figure 1. Assume that $a > b$, $d > c$, $d > b$, and $a > c$. It can be verified that the above game has a Nash equilibrium with the equilibrium strategy of player 1 given by

$$\left(A : \frac{(d-c)}{(d-c)+(a-b)}\ ; \qquad B : \frac{(a-b)}{(d-c)+(a-b)} \right)$$

while the equilibrium strategy of player 2 is given by

$$\left(A : \frac{(d-b)}{(d-b)+(a-c)}\ ; \qquad B : \frac{(a-c)}{(d-b)+(a-c)} \right)$$

For these mixed strategies to make sense, the numbers a, b, c, d must be assigned in such a way that the ratios of the differences

$$\frac{d-c}{a-b}\ ;\ \frac{d-b}{a-c}$$

	2	
1	A	B
A	$a,\ -a$	$b,\ -b$
B	$c,\ -c$	$d,\ -d$

Figure 1: A zero sum game

are meaningful. von Neumann and Morgenstern came up with an extremely elegant theory for determining cardinal utilities.

4 Von Neumann - Morgenstern Utilities

Let X be a set of outcomes. Consider a player i and suppose we focus on the preferences that the player has over the outcomes in X. These preferences can be expressed in the form of a binary relation $\succeq$ on X. Given $x_1, x_2 \in X$, let us define the following:

- $x_1 \succeq x_2$: outcome x_1 is weakly preferred to outcome x_2
- $x_1 \succ x_2$: outcome x_1 is strictly preferred to outcome x_2
- $x_1 \sim x_2$: outcomes x_1 and x_2 are equally preferred (player i is indifferent between x_1 and x_2)

Note immediately that

- $x_1 \succ x_2 \Longleftrightarrow x_1 \succeq x_2$ and $\sim (x_2 \succeq x_1)$
- $x_1 \sim x_2 \Longleftrightarrow x_1 \succeq x_2$ and $x_2 \succeq x_1$

It is clear that the relation $\succeq$ is reflexive. To describe the interaction of preferences with uncertainty about which outcome will be selected, the notion of a lottery (or probability distribution) is a natural tool that can be used. Suppose $X = \{x_1, x_2, \ \ldots, x_m\}$. Then a lottery on X is a probability distribution

$$[p_1 : x_1;\ p_2 : x_2;\ \ldots;\ \ p_m : x_m]$$

Note that

$$p_j \geq 0 \text{ for } j = 1, 2, \ldots, m \text{ and } \sum_{j=1}^{m} p_j = 1.$$

We now present, in the form of six axioms, several natural and desirable properties that we would like preferences to satisfy. These axioms are: completeness, transitivity, substitutability, decomposability, monotonicity, and continuity. These axioms were enunciated by Von Neumann and Morgenstern.

Axiom 1 (Completeness)

This can be formally expressed as

$$\forall x_1, x_2 \in X,\ x_1 \succ x_2 \text{];or } \ x_2 \succ x_1 \ \text{ or } \ x_1 \sim x_2.$$

The completeness property means that the preference relation $\succeq$ induces an ordering on X which allows for ties among outcomes.

Axiom 2 (Transitivity)

This states that

$$x_1 \succeq x_2 \text{ and } x_2 \succeq x_3 \Longrightarrow x_1 \succeq x_3 \;\; \forall x_1, x_2, x_3 \in X$$

To see why transitivity is natural requirement, we have to just visualize what would happen if transitivity is not satisfied. Suppose $x_1 \succeq x_2$ and $x_2 \succeq x_3$ but $x_3 \succ x_1$. Assume that the player in question is willing to pay a certain amount of money if she is allowed to exchange a current outcome with a more preferable outcome. Then the above three relationships will lead to the conclusion that the player is willing to pay a non-zero sum of money to exchange outcome x_3 with the same outcome! Such a situation is popularly known as a *money pump* situation and clearly corresponds to an inconsistent situation.

We now extend the relation $\succeq$ to lotteries over outcomes using the following axioms.

Axiom 3 (Substitutability)

If $x_1 \sim x_2$, then for all sequences of one or more outcomes $x_3, \ldots, x_m$, and sets of probabilities $p, p_3, \ldots, p_m$ such that

$$p + \sum_{j=3}^{m} p_j = 1,$$

the lotteries $[p : x_1;\; p_3 : x_3\;;\; \ldots\;;\; p_m : x_m]$ and $[p : x_2;\; p_3 : x_3\;;\; \ldots\;;\; p_m : x_m]$ are indifferent to the player. We write this as

$$[p : x_1;\; p_3 : x_3\;;\; \ldots\;;\; p_m : x_m] \sim [p : x_2;\; p_3 : x_3\;;\; \ldots\;;\; p_m : x_m]$$

Substitutability implies that the outcome x_1 can always be substituted with outcome x_2 as long as the above technical condition is satisfied.

Axiom 4 (Decomposability)

Suppose σ is a lottery over X and let $P_\sigma(x_i)$ denote the probability that x_i is selected by σ. An example of σ for $X = \{x_1, x_2, x_3\}$ would be

$$\sigma = [0.6 : x_1;\;\; 0.4 : [0.4 : x_1;\;\; 0.6 : x_2]]$$

This would mean that $P_\sigma(x_1) = 0.76; P_\sigma(x_2) = 0.24$; and $P_\sigma(x_3) = 0$. The decomposability axiom states that

$$P_{\sigma_1}(x_i) = P_{\sigma_2}(x_i) \; \forall x_i \in X \Longrightarrow \sigma_1 \sim \sigma_2$$

As a consequence of this axiom, the following lotteries will all be indifferent to a player:

$$\sigma_1 = [0.76 : x_1;\;\; 0.24 : x_2;\;\; 0 : x_3]]$$

$$\sigma_2 = [0.6 : x_1;\;\; 0.4 : [0.4 : x_1;\;\; 0.6 : x_2]]$$

$$\sigma_3 = [0.4 : x_1;\;\; 0.6 : [0.6 : x_1;\;\; 0.4 : x_2]]$$

$$\sigma_4 = [0.5 : [x_1;\;\; 0.8 : 0.2 : x_2];\;\; 0.5 : [0.72 : x_1;\;\; 0.28 : x_2]]$$

Axiom 5 (Monotonicity)

Consider a player who prefers outcome x_1 to outcome x_2. Suppose σ_1 and σ_2 are two lotteries over $\{x_1, x_2\}$. Monotonicity implies that the player would prefer the lottery that assigns higher probability to x_1. More formerly, $\forall x_1, x_2 \in X$,

$$x_1 \succ x_2 \text{ and } 1 \geq p > q \geq 0 \Longrightarrow [p : x_1;\;\; 1-p : x_2] \succ [q : x_1;\;\; 1-q : x_2]$$

Intuitively, monotonicity means that players prefer more of a good thing.

Axiom 6 (Continuity)

This axiom states that $\forall x_1, x_2, x_3 \in X$,

$$x_1 \succ x_2 \text{ and } x_2 \succ x_3 \Longrightarrow \;\exists\, p \in [0,1] \ni x_2 \sim [p : x_1;\;\; 1-p : x_3]$$

A Useful Lemma

We now state (without proof) a lemma and then state and prove an important theorem.

Lemma: Suppose a relation $\succeq$ satisfies completeness, transitivity, decomposability, and monotonicity. Then if $x_1 \succ x_2$ and $x_2 \succ x_3$, there would exist a probability p such that

$$x_2 \succ [q : x_1; 1-q : x_3] \;\forall\; 0 \leq q < p$$

$$[r : x_1; 1-r : x_3] \succ x_2 \;\forall\; 1 \geq r > p$$

The proof is left as an exercise (see problems at the end of the chapter). Using axioms (1) to (6) and the above lemma, we are now in a position to state and prove the key result due to von Neumann and Morgenstern [1].

5 Von Neumann - Morgenstern Theorem

Theorem: Given a set of outcomes X and a preference relation $\succeq$ on X that satisfies completeness, transitivity, substitutability, decomposability, monotonicity and continuity, there exists a utility function $u : X \to [0,1]$ with the following properties:

1. $u(x_1) \geq u(x_2)$ iff $x_1 \succeq x_2$

2. $u([p_1 : x_1;\; p_2 : x_2\;;\; \ldots\;;\; p_m : x_m]) = \sum_{j=1}^{m} p_j u(x_j)$

Proof : First we look at the degenerate case when $x_i \sim x_j \;\forall\; i,j \in \{1, 2, \ldots, m\}$. That is, the player is indifferent among all $x_i \in X$. Consider the function $u(x_i) = 0 \;\;\forall\; x_i \in X$. Part 1 of the theorem follows immediately. Part 2 follows from decomposability.

If this degenerate case is not satisfied, then there must exist at least one most preferred outcome and at least one least preferred outcome with the former different from the latter. Suppose $\bar{x} \in X$ is a most preferred outcome and $\underline{x} \in X$ is a least preferred outcome. Clearly, $\bar{x} \succ \underline{x}$. Now, given any $x_i \in X$, by continuity, there exists a probability p_i uniquely such that

$$x_i \sim [p_i : \bar{x};\;\; 1 - p_i : \underline{x}]$$

Define $u : X \to [0, 1]$ as

$$u(x_i) = p_i \; \forall x_i \in X$$

For this choice of u, we will now prove Part 1 and Part 2.

Proof of Part 1: Suppose $x_1, x_2 \in X$. Let us define two lotteries σ_1 and σ_2 in the following way, corresponding to x_1 and x_2, respectively.

$$x_1 \sim \sigma_1 = [\; u(x_1) : \bar{x};\;\; 1 - u(x_1) : \underline{x} \;]$$

$$x_2 \sim \sigma_2 = [\; u(x_2) : \bar{x};\;\; 1 - u(x_2) : \underline{x} \;]$$

We will show that $u(x_1) \geq u(x_2) \iff x_1 \succeq x_2$. First we prove that $u(x_1) \geq u(x_2) \implies x_1 \succeq x_2$. Suppose $u(x_1) > u(x_2)$. Since $\bar{x} \succ \underline{x}$, then by monotonicity we can conclude that

$$x_1 \sim \sigma_1 \succ \sigma_2 \sim x_2$$

Using transitivity, substitutability, and decomposability, we get $x_1 \succ x_2$.

Suppose $u(x_1) = u(x_2)$. Then σ_1 and σ_2 are identical lotteries which means

$$x_1 \sim \sigma_1 \equiv \sigma_2 \sim x_2$$

Transitivity now yields $x_1 \sim x_2$. We have thus shown that

$$u(x_1) \geq u(x_2) \implies x_1 \succeq x_2$$

It remains to show that

$$x_1 \succeq x_2 \implies u(x_1) \geq u(x_2)$$

We show the above by proving the contrapositive:

$$u(x_1) < u(x_2) \implies x_2 \succ x_1$$

Note that the contrapositive above can be written down by virtue of completeness. The above statement has already been proved when we showed above that

$$u(x_1) > u(x_2) \implies x_1 \succ x_2$$

All that we have to do is to swap x_1 and x_2 to get the implication for the current case.

Proof of Part 2: First we define

$$u^* = u([p_1 : x_1;\;\; p_2 : x_2;\;\; \ldots;\;\; p_m : x_m])$$

By the definition of u, for each $x_j \in X$, we have

$$x_j \sim [u(x_j) : \bar{x}; \quad 1 - u(x_j) : \underline{x}]$$

Using substitutability, we can replace each x_j (in the definition of u^*) by the corresponding lottery. This yields

$$u^* = u([p_1 : [u(x_1) : \bar{x}; 1 - u(x_1) : \underline{x}] \ \ ; \ldots ; \ \ p_m : [u(x_m)\bar{x}; 1 - u(x_m) : \underline{x}]])$$

Note that the above "nested" lottery only selects between the two outcomes $\bar{x}$ and $\underline{x}$. Using decomposability, we get

$$u^* = u\left(\left[\left(\sum_{j=1}^{m} p_j u(x_j)\right) : \bar{x}; \quad \left(1 - \sum_{j=1}^{m} p_j u(x_j)\right) : \underline{x}\right]\right)$$

We can now use the definition of u to immediately obtain

$$u^* = \sum_{j=1}^{m} p_j u(x_j)$$

This proves Part 2 of the theorem.

Note: In the above theorem, the range of the utility function is [0,1]. It would be useful to have a utility function which is not confined to the range [0,1]. The following result extends utility functions to a wide range of possibilities.

Result: Every positive linear transformation (affine transformation) of a utility function, that is a transformation of the form,

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

where a and b are constants and $a > 0$, yields another utility function (in this case (U)) that satisfies properties (1) and (2) of the above theorem.

The proof of the result is left as an exercise. An interesting consequence of the above result is that some two player games which do not appear to be zero-sum are in fact zero-sum games, as seen by the following examples.

Example: A Constant Sum Game

Consider the constant sum game shown in Figure 2. The constant sum here is equal to 1. By subtracting this constant sum from the utilities of one of the players (say player 2), we end up with the zero-sum game in Figure 3.

Example: A Non-Zero Sum Game

Consider the two player non-zero, non-constant sum game shown in Figure 4. Using affine transformation $g(x) = \frac{1}{2}(x - 17)$ on the utilities of player 1, we get a zero-sum game shown in Figure 5.

	2	
1	A	B
A	2, −1	5, −4
B	−6, 7	−1, 2

Figure 2: A constant sum game

	2	
1	A	B
A	2, −2	5, −5
B	−6, 6	−1, 1

Figure 3: An equivalent zero-sum game

	2	
1	A	B
A	27, −5	17, 0
B	19, −1	23, −3

Figure 4: A non-zero sum game

	2	
1	A	B
A	5, −5	0, 0
B	1, −1	3, −3

Figure 5: An equivalent zero-sum game

6 A Procedure for Computing Von Neumann-Morgenstern Utilities

Given a set of outcomes X, the theory of Von Neumann-Morgenstern utilities provides a way of constructing cardinal utilities on those outcomes. The key observation is that cardinal utilities can be constructed by asking the player concerned appropriate questions about lotteries. We explain this with an example (a simplified version of the one appearing in chapter 9 of [2]).

Suppose $X = \{x_1, x_2, x_3\}$ and assume without loss of generality that the player in question has the following preference ordering: $x_1 \succ x_2 \succ x_3$. We start by assigning numbers to the most preferred outcome x_1 and least preferred outcome x_3 in an arbitrary way, respecting only the fact that x_1 gets a larger number than x_3. Suppose we choose the numbers 200 and 100 respectively $(u(x_1 = 200; u(x_3) = 100)$. We now try to fix a number for x_2. For this, we ask questions such as the following: would you prefer x_2 with probability 1 or a lottery that gives you x_1 with probability $\frac{1}{2}$ and x_3 with probability $\frac{1}{2}$. If the player prefers the certain event x_2 to lottery, the implication is that x_2 ranks higher than the midpoint between x_1 and x_3, which means x_2 must be assigned a number greater than 150. This situation is pictorially depicted in Figure 6.

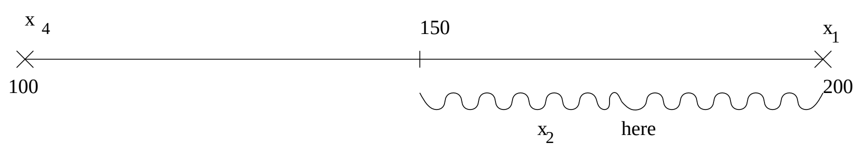

Figure 6: Scenario 1

A possible next question to the player would be: Do you prefer x_2 for certain or the outcome x_1 with probability 0.75 and the outcome x_3 with probability 0.25 ? If the player prefers the lottery, then the situation will be depicted in Figure 7.

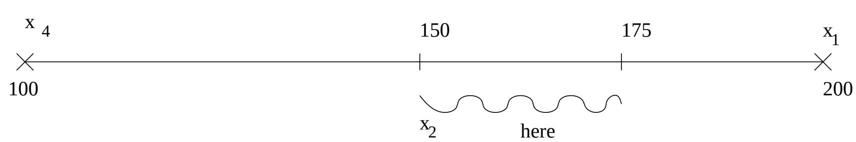

Figure 7: Scenario 2

After a logical sequence of such questions, we would eventually find a lottery such that player 1 is indifferent between x_2 and perhaps the lottery $[0.7 : x_1;\ \ 0.3 : x_3]$. This means we assign the number 170 to x_2 as shown in Figure 8.

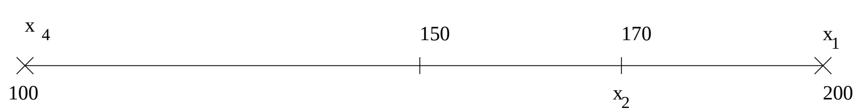

Figure 8: Final assignment

The existence of a unique such solution is guaranteed by Von Neumann - Morgenstern utility theory as long as our exploration is within the axiomatic framework.

7 Utilities and Money

It is tempting to interpret utilities in monetary terms. However, it is not always appropriate to represent utilities by money. There are many reasons for this. First, utility of an individual is not necessarily related to quantity of money. A simple example would be the utility derived by a desperately needy person through a certain amount of money (say Rs. 100) derived through the same amount of money by a rich person. Secondly, money may not always be involved in every transaction that we undertake. An example would be kidney exchange or barter transaction.

8 To Probe Further

As already stated, the material of this chapter has been culled out of the treatment that appears in [3], [2], and [4]. The reader must consult these references for more insights. The treatment [4] is rigorous and comprehensive. An exhaustive account appears in the original classic work of von Neumann and Morgenstern [1].

9 Problems

1. Complete the proof of Lemma 1. (Proof is available in [3])
2. Complete the proof of the result that affine transformations of a utility function do not affect properties 1 and 2 of the von Neumann and Morgenstern utilities.
3. Straffin [2] describes a simple graphical way of investigating whether or not a given two player non-zerosum game is equivalent to a zero-sum game. This involves plotting of the utilities of player 1 and player 2 on the X-Y plane. Try to work this out.

References

[1] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[2] Philip D. Straffin Jr. *Game Theory and Strategy*. The Mathematical Association of America, 1993.

[3] Yoam Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, USA, 2009, 2009.

[4] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, USA, 1997.