GAME THEORY AND THE MINIMAX THEOREM

MIRA SCARVALONE

ABSTRACT. Game theory is a very important branch of applied mathematics with many uses in the social sciences, biological sciences, and philosophy. Game theory attempts to mathematically explain behavior in situations in which an individual's outcome depends on the actions of others. Arguably the most important result in game theory, the Minimax Theorem was stated in 1928 by mathematician John von Neumann in his paper *Zur Theorie Der Gesellschaftsspiele*, and forms the basis for all subsequent findings in the subject. This paper will provide a brief introduction to zero-sum games and the notion of equilibrium, as well as an elementary proof of the theorem.

CONTENTS

1. Introduction to Games 1
2. Zero-Sum Games 2
3. The Minimax Theorem 3
References 5

1. INTRODUCTION TO GAMES

The notion of a game in this context is similar to certain familiar games like chess or bridge. Starting from a beginning point, each player performs a sequence of moves, often involving choices at each step. Each outcome, or terminal point, at the end of the game is alloted a particular payoff, usually in the form of money, given to each of its players.

Definition 1.1. An n-person game is one in which there are n players, and a *payoff function*, which assigns an n-vector to each terminal vertex of the game, indicating each player's earnings.

Definition 1.2. A *strategy* refers to a player's plan specifying which choices it will make in every possible situation, leading to an eventual outcome. Let Σ_i denote the set of all strategies for player i. In order to decide which strategy is best, player i will have to choose the strategy which maximizes its payoff (i.e., the i-th component of the payoff function).

Letting π denote the probability of a certain combination of strategies occuring, we can derive a mathematical expression for the payoff function, given player i uses strategy $\sigma_i \; \epsilon \; \Sigma_i$:

$\pi \; (\sigma_1, \sigma_2, ..., \sigma_n) = (\pi_1(\sigma_1 ... \sigma_n), \pi_2(\sigma_1 ... \sigma_n), ..., \pi_n(\sigma_1 ... \sigma_n))$

where σ_1 represents player *1*'s strategy, σ_2 represents player *2*'s strategy, and so on, while π_1 represents the probability of player *1* choosing strategy σ_1, π_2 represents the probability of player *2* choosing strategy σ_2, and so on.

It is possible to express this function through an n-dimensional array of n-vectors, called the *normal form* of the game.

Definition 1.3. A strategy n-tuple $(\sigma_1, \sigma_2, ..., \sigma_n)$ is said to be in *equilibrium* if and only if no player has any reason to change its strategy, assuming the other players do not change theirs. For example, if the strategy n-tuple $(\sigma_1, \sigma_2, ..., \sigma_n)$ is in equilibrium, for any i = 1,...n, and any $\hat{\sigma i} \; \epsilon \; \Sigma_i$,

$\pi_i \; (\sigma_1, ..., \sigma_{i-1}, \hat{\sigma_i}, \sigma_{i+1}, ..., \sigma_n) \leq \pi_i \; (\sigma_1, \sigma_2, ..., \sigma_n)$

2. Zero-Sum Games

Definition 2.1. A *zero sum game* is one in which, at each terminal vertex, the payoff function's components add up to zero. In other words, every amount that one player wins must be lost by another player.

The normal form of a two-person zero-sum game is just a matrix with each row representing one of player 1's strategies, and each column representing one of player 2's strategies. The expected payoff (for player 1), if player 1 chooses its i-th strategy and player 2 chooses its j-th strategy, will be element a_{ij} of the matrix. Since this matrix represents the different payoffs for player 1, it makes sense that player 1 will try to maximize the element a_{ij} which is chosen, while player 2 will try to minimize it, since that value is equivalent to player 2's loss. For the rest of this paper, we will be dealing with two-person zero-sum games.

Definition 2.2. A *saddle point* is an element in the game matrix that is both the largest in its column and the smallest in its row. Not all game matrices have saddle points, but if they do, they will clearly be the equilibrium strategies, since they both maximize player 1's payoff, and minimize player 2's loss.

Definition 2.3. A *gain-floor* is the minimum payoff player 1 will receive given player 2's attempt to minimize their payoff. Mathematically, the gain-floor can be represented as such: υ_1'=max$_i$\{min$_j$ a_{ij}\} Conversely, the *loss-ceiling* is the maximimum loss player 2 can experience given player 1's attempt to maximize their payoff. Mathematically, the loss-ceiling can be represented as such: υ_2'=min$_j$\{max$_i$ a_{ij}\}

Lemma 2.4. *Player 1's gain-floor cannot exceed player 2's loss-ceiling, i.e.* $\upsilon_1' \leq \upsilon_2'$

Proof. Let υ_1' and υ_2' be defined as gain-floors and loss-ceilings, respectively. Since υ_1'=max$_i$\{min$_j$ a_{ij}\}, υ_1' is less than all other payoffs in column j, while υ_2' is greater than all other payoffs in row i. Let x be the payoff in row i and column j, i.e., $x = a_{ij}$. Since $\upsilon_1' \leq x \leq \upsilon_2'$, $\upsilon_1' \leq \upsilon_2'$

□

It is notable that if $\upsilon_1' = \upsilon_2'$, a saddle point exists, but if the inequality is strict, the game does not have a saddle point.

Definition 2.5. A *mixed strategy* is a probability distribution on the set of a player's pure strategies. When a player has a finite number of m strategies, its mixed strategy can be expressed as an m-vector, $x = (x_1, ..., x_m)$ such that $x_i \geq 0$ and $\sum_{i=1}^{m} x_i = 1$

Suppose players 1 and 2 are playing the matrix game A, where $A = (a_{ij})$ is an mxn matrix. Let X denote the set of all mixed strategies for player 1, and Y represent the set of all mixed strategies for player 2. If player 1 chooses mixed strategy x while player 2 chooses mixed strategy y, then the expected payoff can be written as $A(x, y) = \sum_{i=1}^{m} \sum_{j=1}^{n} x_i a_{ij} y_j$.
In matrix notation, this can be expressed as $A(x, y) = xAy^T$
Player 1's expected gain floor, if using strategy x, can be expressed as
$v(x) = \min_y xAy^T$, which can be thought of as a weighted average of the expected payoffs for player 1 using x against player 2's pure strategies.
Thus the minimum will occur with pure strategy j: $v(x) = \text{min} x A_j$ where A_j is the j-th column of the matrix A. Thus player 1 should choose x in order to maximize $v(x)$, so as to obtain $v_I = \max_x \min_j xA_j$. This strategy x is known as player 1's *maximin* strategy.
Conversely, if player 2 chooses strategy y, it will obtain the expected loss-ceiling $v(y) = \max_i A_i y^T$ where A_i is the i-th row of A, and will choose y in order to obtain $v_{II} = \min_y \max_i A_i y^T$. This strategy y is known as player 2's *minimax* strategy.
Values v_I and v_{II} are known as the *values* of the game to player 1 and player 2, respectively.

3. The Minimax Theorem

Lemma 3.1. *Theorem of the Supporting Hyperplane: Let B be a closed convex set of points in an n-dimensional euclidean space, and let $x = (x_1, ..., x_n) \notin B$. Then $\exists$ numbers $p_1, ..., p_n, p_{n+1}$ such that the following conditions hold:*

$\sum_{i=1}^{n} p_i x_i = p_{n+1}$ (3.2)

and

$\sum_{i=1}^{n} p_i y_i > p_{n+1}$ *for all* $y \epsilon B$ (3.3)

Proof. Let z be the point in B whose distance to x is a minimum. This point exists and is unique because B is closed and convex. Let $p_i = z_i - x_i$ for i=1,...,n
$p_{n+1} = \sum_{i=1}^{n} x_i z_i$ - $\sum_{i=1}^{n} x_i^2$

Since $\sum_{i=1}^{n} p_i x_i = \sum_{i=1}^{n} (z_i - x_i) x_i = \sum_{i=1}^{n} x_i z_i$ - $\sum_{i=1}^{n} x_i^2 = p_{n+1}$, condition (3.2) holds.
To prove that condition (3.3) holds as well, we will see that
$\sum_{i=1}^{n} p_i z_i = \sum_{i=1}^{n} z_i^2$ - $\sum_{i=1}^{n} z_i x_i$, so $\sum_{i=1}^{n} p_i z_i$ - $p_{n+1} = \sum_{i=1}^{n} z_i^2$ - $2 \sum_{i=1}^{n} z_i x_i + \sum_{i=1}^{n} x_i^2 = \sum_{i=1}^{n} (z_i - x_i)^2 > 0$ Thus, $\sum_{i=1}^{n} p_i z_i > p_{n+1}$

Now suppose condition (3.3) does not hold, i.e. $\exists$ y ϵ B such that $\sum_{i=1}^{n} p_i y_i \leq p_{n+1}$.

Since B is a convex set, the line between y and z must be entirely contained in B, so $\forall\ 0 \leq$ r ≤ 1, $w_r = ry + (1-r)z\ \epsilon$ B
The square of the distance from x to w_r can be written as $\rho^2(x, w_r) = \sum_{i=1}^{n}(x_i - ry_i - (1-r)z_i)^2$, so $\frac{\partial\rho^2}{\partial r} = 2\sum_{i=1}^{n}(z_i - y_i)(x_i - ry_i - (1-r)z_i)$
$= 2\sum_{i=1}^{n}(z_i - x_i)y_i - 2\sum_{i=1}^{n}(z_i - x_i)z_i + 2\sum_{i=1}^{n} r(z_i - y_i)^2$
$= 2\sum_{i=1}^{n} p_i y_i - 2\sum_{i=1}^{n} p_i z_i + 2r\sum_{i=1}^{n}(z_i - y_i)^2$
If this value is evaluated at r = 0, $\frac{\partial\rho^2}{\partial r} = 2\sum_{i=1}^{n} p_i y_i - 2\sum_{i=1}^{n} p_i z_i$
Since the first term on the right side is less than or equal to $2p_{n+1}$, and the second term is greater than $2p_{n+1}$, $\frac{\partial\rho^2}{\partial r} < 0$.
Thus, for r close enough to zero, $\rho(x, w_r) < \rho(x, z)$.
However, this contradicts the assumption that z is the point in B whose distance to x is a minimum, so it follows that $\forall$ y ϵ B, condition (3.3) holds.

□

Lemma 3.4. *Theorem of The Alternative for Matrices: Let $A = (a_{ij})$ be an mxn matrix. Then either of the following two conditions must hold:*

The point 0 is contained in the convex hull of the m+n points $a_1 = (a_{11}, ..., a_{m1}), a_2 = (a_{12}, ..., a_{m2}), ..., a_n = (a_{1n}, ..., a_{mn})$ *and* $e_1 = (1, 0, ...0), e_2 = (0, 1, 0, ...0), ..., e_m = (0, 0, ..., 1)$(3.5)

$\exists\ x_1, ..., x_m$ *such that* $x_i > 0$, $\sum_{i=1}^{m} x_i = 1$, *and* $\sum_{i=1}^{n} a_{ij}x_i > 0$ *for* $j = 1, ..., n$ (3.6)

Proof. Let A be an m × n matrix, and suppose condition (3.5) holds. Then we are finished with this circumstance.
Now suppose (3.5) does not hold. According to Lemma (3.1), $\exists$ numbers $p_1, ..., p_{m+1}$ such that $\sum_{j=1}^{m} 0 \cdot p_j = p_{m+1}$ since 0 is not contained in the convex hull of these points. Therefore $p_{m+1} = 0$, and since $\sum_{j=1}^{m} p_j y_j > p_{m+1}$, $\sum_{j=1}^{m} p_j y_j > 0\ \forall$ y in the convex set. More specifically, this statement holds if y is one of the m+n vectors listed above, since they are included within the larger convex hull.
Therefore $\sum a_{ij} p_i > 0\ \forall\ j$, so $p_i > 0\ \forall$ i.It follows that $\Sigma p_i > 0$.
Let $x_i = p_i / \Sigma p_i$. Therefore $\sum a_{ij} x_i > 0$, and $x_i > 0$.
By definition,$\Sigma x_i > 0$.
Therefore, condition (3.6) holds if (3.5) does not. □

Theorem 3.7. *The Minimax Theorem:* $\max_x \min_j xA_j = \min_y \max_i A_i y^T$, *i.e.* $v_I = v_{II}$

Proof. Let A be a matrix game. According to Lemma 3.4, either condition (3.4) or (3.5) must be true.
Suppose condition (3.4) holds. Thus, 0 is contained in the convex hull of the m+n points, and is therefore a convex linear combination of the m+n vectors. Consequently, $\exists\ s_1, ..., s_{m+n}$ such that $\sum_{j=1}^{n} s_j a_{ij} + sn + i = 0$ for $i = 1, ..., m$, $s_j \geq 0$ for $j = 1, ..., m+n$, and $\sum_{j=1}^{m+n} s_j = 1$
If all the numbers $s_1, ..., s_n$ were 0, then 0 would be a convex linear combination of the m unit vectors $e_1, ..., e_m$, which would be a contradiction since these are independent vectors. Therefore, at least one of the numbers $s_1, ..., s_n$ is positive,

meaning $\sum_{j=1}^{n} s_j > 0$.
Let $y_j = s_j \diagup \sum_{j=1}^{n} s_j$. Therefore $y_j \geq 0$, $\sum_{j=1}^{n} y_j = 1$, and
$\sum_{j=1}^{n} a_{ij} y_i = -s_{n+1} \diagup \sum_{j=1}^{n} s_j \leq 0 \; \forall$ i.
Thus $v(y) \leq 0$ and $v_{II} \leq 0$.
Now suppose that condition (3.5) holds instead. Then $v(x) > 0$, and $v_I > 0$.
Since it is impossible for both conditions not to hold, it is not posible to have a situation in which $v_I \leq 0 < v_{II}$.
Now let's look at a different matrix game B, where $B = (b_{ij})$, and $b_{ij} = a_{ij} + k$, where k is a random variable. $\forall$ x, y, $xBy^T = xAy^T + k$.
Therefore, $v_I(B) = v_I(A) + k$, and $v_{II}(B) = v_{II}(A) + k$. Since it is impossible for $v_I(B) < 0 < v_{II}(B)$, we cannot obtain a situation in which $v_I(A) + k < 0 < v_{II}(A) + k$, where by subtracting k from both sides we get $v_I(A) < -k < v_{II}(A)$.
But since k was arbitrary, it is impossible for $v_I < v_{II}$.
Using the same method as lemma 2.4 we see that $v_I \leq v_{II}$. Therefore $v_I = v_{II}$ $\square$

References

[1] John Von Neumann and Oskar Morgenstern. Theory of Games and Economic Behavior. Princeton University Press. 1947.