Two-person zero-sum games

Moty Katzman

October 28, 2016

"Maxminization"

You play a game and want to ensure you don't do too badly.
You find the worst possible outcome of each of your actions and choose the action with highest possible *worst* outcome.
E.g., in a 2-person game (S, T, u_1, u_2), if we restrict to pure strategies Alice plays the $s \in S$ which maximizes $\min_{t \in T} u_1(s, t)$ and Bob plays the $t \in T$ which maximizes $\min_{s \in S} u_2(s, t)$.
This guarantees Alice at least $\max_{s \in S} \min_{t \in T} u_1(s, t)$ and Bob $\max_{t \in T} \min_{s \in S} u_2(s, t)$.

“Maxminimizing” with mixed strategies

If P and Q denote the mixed-strategies for Alice and Bob, Alice could play the $p \in P$ which maximizes $\min_{q \in Q} u_1(p, q)$ and Bob plays the $q \in Q$ which maximizes $\min_{p \in P} u_2(p, q)$.

Zero-sum games

Definition. A two-person zero-sum game a game in strategic form (S_1, S_2, u_1, u_2) where $u_1(s_1, s_2) + u_2(s_1, s_2) = 0$ for all $s_1 \in S_1$ and $s_2 \in S_2$.
("My gain is your loss, your gain is my loss.")
In a zero-sum game a player's best response is always the one that does most damage to the other!
Games in which one side wins and the other loses, or both sides draw, can be modeled as zero-sum games, e.g., chess and Rock-Scissors-Paper.
Wars are seldom zero-sum games.
Unlike cooperative games, zero-sum games can be solved, in some sense.

Pure strategies: looking for saddle points

Example. Consider the zero-sum game

	a	b	c
I	1,-1	1,-1	8, -8
II	5, -5	2, -2	4, -4
III	7, -7	0, 0	0, 0

We don't need to keep track of both sides' payoffs, and by convention we keep track of the row-player payoffs only. The abbreviated form of the game is

	a	b	c
I	1	1	8
II	5	2	4
III	7	0	0

and we set $u = u_1$ to be the row-players's utility function.
How should both sides maxminimize this? Consider the case where only pure strategies are played.

Bob's responses to each of Alice's strategies are the ones resulting in minimal score for Alice:

	a	b	c
I	1	1	8
II	5	2	4
III	7	0	0

So Alice would choose the strategy $s = s^*$ which maximizes

$$\max_{s \in \{I, II, III\}} \min_{t \in \{a, b, c\}} u(s, t),$$

i.e., $s^* = II$, and choosing this guarantees her a payoff of 2.

Alice's responses to each of Bob's strategies are the ones resulting in maximal score for Alice:

	a	b	c
I	1	1	$\boxed{8}$
II	5	$\boxed{2}$	4
III	$\boxed{7}$	0	0

So Bob would choose the strategy $t = t^*$ which minimizes

$$\min_{t \in \{a,b,c\}} \max_{s \in \{I,II,III\}} u(s,t),$$

i.e., $t^* = b$, and this strategy guarantees a payoff of -2 for Bob. Note that (II, b) is a Nash equilibrium; this is because (II, b) is a *saddle-point* for u, i.e., $u(II, b) \geq u(s, b)$ for all $s \in \{I, II, III\}$ and $u(II, b) \leq u(II, t)$ for all $t \in \{a, b, c\}$.

In this example

$$\max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t) = \min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t) = 2$$

This implies that for s^* and t^* such that

$$\min_{t\in\{a,b,c\}} u(s^*,t) = \max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t)$$

$$\max_{s\in\{I,II,III\}} u(s,t^*) = \min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t),$$

(s^*, t^*) is a Nash equilibrium.
As long as Bob plays t^* Alice cannot get more than 2,
as long as Alice plays s^* Bob cannot get more than -2.

Another example

Consider the zero-sum game

	a	b	c
I	2, -2	5, -5	0, 0
II	3, -3	1, -1	2, -2
III	4, -4	3, -3	6, -6

Knowing that this game is zero-sum, we don't need to keep track of both sides' payoffs, and by convention we keep track of the row-player payoffs only. The abbreviated form of the game is

	a	b	c
I	2	5	0
II	3	1	2
III	4	3	6

and we set $u = u_1$ to be the row-players's utility function.

Bob's responses to each of Alice's strategies are the ones resulting in minimal score for Alice:

	a	b	c
I	2	5	$\boxed{0}$
II	3	$\boxed{1}$	2
III	4	$\boxed{3}$	6

So Alice would choose the strategy $s = s^*$ which maximizes $\min_{t \in \{a,b,c\}} u(s,t)$,. i.e. the value $s = s^*$ for which

$$\max_{s \in \{I,II,III\}} \min_{t \in \{a,b,c\}} u(s,t) = \min_{t \in \{a,b,c\}} u(s^*,t),$$

i.e., $s^* = III$, and choosing this guarantees her a payoff of 3.

Alice's responses to each of Bob's actions are the ones resulting in maximal score for Alice:

	a	b	c
I	2	$\boxed{5}$	0
II	3	1	2
III	$\boxed{4}$	3	$\boxed{6}$

So Bob would choose the strategy $t = t^*$ which minimizes $\max_{s \in \{I,II,III\}} u(s,t)$, i.e., the value $t = t^*$ for which

$$\min_{t \in \{a,b,c\}} \max_{s \in \{I,II,III\}} u(s,t) = \max_{s \in \{I,II,III\}} u(s,t^*),$$

i.e., $t^* = a$, and this strategy guarantees a payoff of -4 for Bob. Note that

$$\min_{t \in \{a,b,c\}} \max_{s \in \{I,II,III\}} u(s,t) = 4 > 3 = \max_{s \in \{I,II,III\}} \min_{t \in \{a,b,c\}} u(s,t).$$

Note also that (III, a) is *not* a Nash equilibrium. Is this the best that can be achieved?

minmax $\geq$ maxmin

Lemma. Consider a finite zero-sum game (S, T, u). We have

$$\min_{t \in T} \max_{s \in S} u(s,t) \geq \max_{s \in S} \min_{t \in T} u(s,t).$$

Proof. Write $m = \max_{s \in S} \min_{t \in T} u(s,t)$. For all $s \in S$ and $t' \in T$ we have $u(s,t') \geq \min_{t \in T} u(s,t)$ so $\max_{s \in S} u(s,t') \geq \max_{s \in S} \min_{t \in T} u(s,t) = m$ so in particular $\min_{t \in T} \max_{s \in S} u(s,t) \geq m$. $\square$

Theorem. Consider a finite zero-sum game (S, T, u). Let $s^* \in S$ and $t^* \in T$ be such that

$$\min_{t \in T} \max_{s \in S} u(s, t) = \max_{s \in S} u(s, t^*)$$

and

$$\max_{s \in S} \min_{t \in T} u(s, t) = \min_{t \in T} u(s^*, t).$$

The strategy profile (s^*, t^*) is a saddle-point for u (and hence a Nash equilibrium) if and only if

$$\min_{t \in T} \max_{s \in S} u(s, t) = \max_{s \in S} \min_{t \in T} u(s, t).$$

Proof. Write $m = \max_{s \in S} \min_{t \in T} u(s,t)$ and $M = \min_{t \in T} \max_{s \in S} u(s,t)$.
Assume first that (s^*, t^*) is a saddle-point. Then for all $s \in S$ and $t \in T$ we have

$$u(s^*, t) \geq u(s^*, t^*) \geq u(s, t^*)$$

hence

$$\min_{t \in T} u(s^*, t) \geq u(s^*, t^*) \geq \max_{s \in S} u(s, t^*)$$

and

$$m = \max_{s \in S} \min_{t \in T} u(s,t) \geq u(s^*, t^*) \geq \min_{t \in T} \max_{s \in S} u(s,t) = M$$

and since $m \leq M$, we get equalities throughout.
Assume now that $m = M$. For all $s \in S$ and $t \in T$ we have

$$u(s^*, t) \geq \min_{\tau \in T} u(s^*, \tau) = m = M = \max_{\sigma \in S} u(\sigma, t^*) \geq u(s, t^*)$$

and in particular $u(s^*, t) \geq u(s^*, t^*)$ and $u(s, t^*) \leq u(s^*, t^*)$. $\square$.

Mixed strategies

Consider a finite zero-sum game (S, T, u) with $S = \{s_1, \ldots, s_m\}$ and $T = \{t_1, \ldots, t_n\}$.
Mixed strategies for row player:

$$\Delta^R = \{(p_1, \ldots, p_m) \mid p_1, \ldots, p_m \geq 0, p_1 + \cdots + p_m = 1\}$$

(s_i is played with probability p_i.)
Mixed strategies for column player:

$$\Delta^C = \{(q_1, \ldots, q_n) \mid q_1, \ldots, q_n \geq 0, q_1 + \cdots + q_n = 1\}$$

(t_i is played with probability q_i.)
The mixed strategy pair $(p, q) \in \Delta^R \times \Delta^C$ yields a payoff of

$$u(p, q) = \sum_{1 \leq i \leq m} \sum_{1 \leq j \leq n} p_i q_j u(s_i, t_j)$$

for the row player and $-u(p, q)$ for the column player.

Mixed strategies

Definition. Consider a zero sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. Let

$$\overline{V} = \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y)$$

and

$$\underline{V} = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

An optimal strategy for the column player is a $y^* \in \Delta^C$ for which $\overline{V} = \max_{x \in \Delta^R} u(x, y^*)$.
An optimal strategy for the row player is a $x^* \in \Delta^R$ for which $\underline{V} = \min_{y \in \Delta^C} u(x^*, y)$.

minmax $\geq$ maxmin

Previous results hold for mixed strategies as well:

Lemma. Consider a finite zero sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. We have

$$\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) \geq \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

(Same) Proof. Write $m = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y)$. For all $x \in \Delta^R$ and $y' \in \Delta^C$ we have $u(x, y') \geq \min_{y \in \Delta^C} u(x, y)$ so $\max_{x \in \Delta^R} u(x, y') \geq \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y) = m$ so in particular $\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) \geq m$. □

Theorem. Consider a finite zero-sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. Let $x^* \in \Delta^R$ and $y^* \in \Delta^C$ be optimal strategies, i.e.,

$$\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = \max_{x \in \Delta^R} u(x, y^*)$$

and

$$\max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y) = \min_{y \in \Delta^C} u(x^*, y).$$

The strategy profile (x^*, y^*) is a Nash equilibrium if and only if

$$\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

Proof. Write $m = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y)$ and $M = \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y)$.
Assume first that (x^*, y^*) is a Nash equilibrium. Then for all $x \in \Delta^R$ and $y \in \Delta^C$ we have

$$u(x^*, y) \geq u(x^*, y^*) \geq u(x, y^*)$$

hence

$$\min_{y \in \Delta^C} u(x^*, y) \geq u(x^*, y^*) \geq \max_{x \in \Delta^R} u(x, y^*)$$

and

$$m = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y) \geq u(x^*, y^*) \geq \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = M$$

and since $m \leq M$, we get equalities throughout.
Assume now that $m = M$. For all $x \in \Delta^R$ and $y \in \Delta^C$ we have

$$u(x^*, y) \geq \min_{\overline{y} \in \Delta^C} u(x^*, \overline{y}) = m = M = \max_{\overline{x} \in \Delta^R} u(\overline{x}, y^*) \geq u(x, y^*)$$

and in particular $u(x^*, y) \geq u(x^*, y^*)$ and $u(x, y^*) \leq u(x^*, y^*)$.
□

Two-person zero-sum games have mixed-strategy Nash-equilibria!

We now state one of the most important results in game theory: John von Neumann's *Minimax Theorem* (1928)

Theorem: (The Minimax Theorem) $\overline{V} = \underline{V}$.

Corollary: Two-person finite zero-sum games have at least one mixed-strategy Nash-equilibrium: any pair of optimal strategies is a Nash equilibrium.

The value of a game

Definition. Let $G = (S, T, u)$ be a zero sum game with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. The value of G is defined as the common value of

$$\overline{V} = \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y)$$

and

$$\underline{V} = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

Zero sum games in matrix form

A finite zero sum game (S, T, u) with $S = \{s_1, \dots, s_m\}$ and $T = \{t_1, \dots, t_n\}$ can be represented by a $m \times n$ matrix $M = (u(s_i, t_j))$ and mixed strategy sets

$$\Delta_m = \left\{ \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} : x_1, \dots, x_m \geq 0 \text{ and } x_1 + \cdots + x_m = 1 \right\} \subseteq \mathbb{R}^m$$

and

$$\Delta_n = \left\{ \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} : y_1, \dots, y_n \geq 0 \text{ and } y_1 + \cdots + y_n = 1 \right\} \subseteq \mathbb{R}^n$$

for the row and column players.
Now the payoff of a strategy pair $(x, y) \in \Delta_m \times \Delta_n$ is $x^t M y$ for the row player and $-x^t M y$ for the column player.

Example: symmetric games

A zero-sum game $G = (S, T, u)$ is *symmetric* if $S = T$ and for all $s_1, s_2 \in S$, $u(s_2, s_1) = -u(s_1, s_2)$. Let A be the matrix associated with a symmetric game $G = (\{1, \dots, n\}, \{1, \dots, n\}, u)$.

(a) The value V of a symmetric game is zero.

(b) An optimal strategy for one player is also an optimal strategy for the other player.

(a) For any $p \in \Delta^R$ we have
$p^T Ap = (p^T Ap)^T = p^T A^T p = -p^T Ap$, hence $p^T Ap = 0$ so
$V = \min_{y \in \Delta^C} \max_{x \in \Delta^R} x^T Ay \geq \min_{y \in \Delta^C} y^T Ay = 0$ and
$V = \max_{x \in \Delta^R} \min_{y \in \Delta^C} x^T Ay \leq \max_{x \in \Delta^R} x^T Ax = 0$.
(b) If $p \in \Delta^R$ is optimal for the row player, $p^T A$ has non-negative entries, and so $(p^T A)^T = A^T p = -Ap$ has non-negative entries and p is optimal for the column player.

Example: symmetric games

Verify that the zero-sum game

	A	B	C
I	0	2	-1
II	-2	0	3
III	1	-3	0

has Nash-equilibrium $((1/2, 1/6, 1/3), (1/2, 1/6, 1/3))$.

Proposition. Let $(\{s_1, \ldots, s_m\}, \{t_1, \ldots, t_n\}, u)$ be a zero-sum game. Let M be the $m \times n$ matrix with $M_{ij} = u(s_i, t_j)$. Suppose that $p^* \in \Delta_m$ and $q^* \in \Delta_n$ are such that the minimal coordinate in ${p^*}^t M$ and the maximal coordinate in Mq^* both equal v. Then the value of the game is v and (p^*, q^*) is an optimal strategy.
Proof. Let V be value of the game.

$$V = \max_{p \in \Delta_m} \min_{q \in \Delta_n} p^t M q \le \max_{p \in \Delta_m} p^t M q^* \le \max_{p \in \Delta_m} p^t \begin{bmatrix} v \\ \vdots \\ v \end{bmatrix} = v$$

$$V = \min_{q \in \Delta_n} \max_{p \in \Delta_m} p^t M q \ge \min_{q \in \Delta_n} {p^*}^t M q \ge \min_{q \in \Delta_n} (v, \ldots, v) q = v$$

hence $v = V$.
Now $V = \max_{p \in \Delta_m} \min_{q \in \Delta_n} u(p, q) \ge \min_{q \in \Delta_n} u(p^*, q) \ge V$, hence $\min_{q \in \Delta_n} u(p^*, q) = V$ and p^* is optimal. Also $V = \min_{q \in \Delta_n} \max_{p \in \Delta_m} u(p, q) \le \max_{p \in \Delta_m} u(p, q^*) \le V$ hence $\max_{p \in \Delta_m} u(p, q^*) = V$ and q^* is optimal.

Example: finding the value of a game

Consider the following zero-sum game

	A	B	C
I	2	-1	-1
II	-2	0	3
III	1	2	1

We verify that (p^*, q^*) is an optimal strategy profile where $p^* = (0, 0, 1)$ and $q^* = (2/5, 0, 3/5)$. Write

$$M = \max\{u(\widehat{I}, q^*), u(\widehat{II}, q^*), u(\widehat{III}, q^*)\} = \max\{1/5, 1, 1\} = 1$$

and

$$m = \min\{u(p^*, \widehat{A}), u(p^*, \widehat{B}), u(p^*, \widehat{C})\} = \min\{1, 2, 1\} = 1.$$

Example: finding the value of a game

Consider the following zero-sum game

	A	B	C	D	E
I	-1	2	-2	0	1
II	-2	-1	3	2	0
III	2	1	0	-1	-2
IV	0	0	2	1	1
V	1	-1	0	-2	1

We verify that (p^*, q^*) is an optimal strategy profile where
$p^* = (5/52, 0, 11/52, 34/52, 2/52)$ and
$q^* = (21/52, 12/52, 0, 3/52, 16/52)$.
Compute

$$\begin{bmatrix} -1 & 2 & -2 & 0 & 1 \\ -2 & -1 & 3 & 2 & 0 \\ 2 & 1 & 0 & -1 & -2 \\ 0 & 0 & 2 & 1 & 1 \\ 1 & -1 & 0 & -2 & 1 \end{bmatrix} \begin{bmatrix} 21/52 \\ 12/52 \\ 0 \\ 3/52 \\ 16/52 \end{bmatrix} = \begin{bmatrix} 19/52 \\ -\mathbf{48}/\mathbf{52} \\ 19/52 \\ 19/52 \\ 19/52 \end{bmatrix}$$

Bob's q^* guarantees not to lose more that $19/5$.

Compute

$$\begin{bmatrix} 5/52 & 0 & 11/52 & 34/52 & 2/52 \end{bmatrix} \begin{bmatrix} -1 & 2 & -2 & 0 & 1 \\ -2 & -1 & 3 & 2 & 0 \\ 2 & 1 & 0 & -1 & -2 \\ 0 & 0 & 2 & 1 & 1 \\ 1 & -1 & 0 & -2 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} 19/52 & 19/52 & \mathbf{58}/\mathbf{52} & 19/52 & 19/52 \end{bmatrix}$$

So Alices p^* guarantees her at least 19/52.
The value of this game is 19/52.

A very different example: A duel

Alice and Bob fight a duel at dawn. They start at a distance of 1 unit apart, each armed with a pistol that can fire once, and they walk toward each other. Each can shoot at will: if the person who shoots first hits the target, he survives and the other dies; otherwise if the first person to shoot misses, the second person shoots the first point blank. If both decide to shoot at the same time, a fair coin is tossed to select who shoots first.
Alice's probability of hitting Bob at a distance of d is given by $A(d)$ and Bob's probability of hitting Alice at a distance of d is given by $B(d)$, and we assume that $A(d)$ and $B(d)$ are decreasing continuous functions of d and that $A(0) = B(0) = 1$. How should this duel proceed?

Consider a strategy profile (x, y) (Alice shoots at a distance x and Bob shoots at a distance y.) The probability of Alice surviving is

$$p(x, y) = \begin{cases} A(x) & \text{if } x > y \\ 1 - B(y) & \text{if } x < y \end{cases}$$

and that of Bob is $q(x, y) = 1 - p(x, y)$.
Alice wants to shoot at a distance x^* for which

$$\sup_{0 \le x \le 1} \inf_{0 \le y \le 1} p(x, y) = \inf_{0 \le y \le 1} p(x^*, y).$$

Define $\beta(x) = \inf_{0 \le y \le 1} p(x, y)$.
If $A(x) > 1 - B(x)$, $\beta(x) = \inf_{x < y \le 1} p(x, y) = 1 - B(x)$ and if $A(x) \le 1 - B(x)$, then $\beta(x) = A(x)$.
Let $z \in [0, 1]$ be such that $1 - B(z) = A(z)$. We have

$$\beta(x) = \begin{cases} 1 - B(x) & \text{if } x < z \\ A(x) & \text{if } x \ge z \end{cases}$$

Alice will look for an x which maximizes $\beta(x)$ and this occurs at $x = z$. A similar analysis shows that Bob will also shoot at distance z.