CS 405/6001: Game Theory and Algorithmic Mechanism Design **Jul–Nov 2026**

Lecture 5: Matrix Games, Saddle Points, and Security Values

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

5.1 Recap

Last time we made precise what a risk-averse player can guarantee without trusting the opponent to cooperate: the **maxmin value**. We showed that a dominant strategy is always a maxmin strategy, and that every PSNE payoff is at least the maxmin value. Today we specialize to **two-player zero-sum (matrix) games**, where these ideas sharpen considerably: security and stability turn out to coincide exactly when a pure equilibrium exists, and when it does not, we are forced to consider **mixed strategies**.

5.2 The Ghost and the Sentinel: A Running Example

Definition 5.1 (Two-player zero-sum games) *A normal-form game $\langle N, (S_i)_{i\in N}, (u_i)_{i\in N}\rangle$ with $N = \{1,2\}$ is **zero-sum** if $u_1 + u_2 \equiv 0$. Writing $u := u_1 = -u_2$, the entire game is captured by a single matrix of player 1's payoffs, since player 2's payoffs are just its negative – hence the name **matrix game**.*

Example 5.2 (The Ghost and the Sentinel) *A red-team exercise: **Ghost** tries to slip a probe past **Sentinel**, a monitoring system. Ghost picks an ingress channel, L or R; Sentinel picks a channel to watch. If the two match, Sentinel catches Ghost; if they differ, Ghost slips through. Writing $+1/-1$ for a win/loss to Ghost gives the zero-sum matrix*

$$u = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix},$$

rows indexed by Ghost's channel, columns by Sentinel's. This is a matrix game in the sense above: any head-to-head contest where one side's gain is exactly the other's loss – an attacker racing a detector, or a penalty shooter against a goalkeeper – has this structure.

Example 5.3 (Scaling up: three attack vectors) *Now let Ghost choose among three moves, T (Trojan payload), M (man-in-the-middle), B (botnet flood), and Sentinel among three defenses, L (log monitoring), C (content filtering), R (rate limiting), with Ghost's threat-bounty payoffs*

$$u = \begin{pmatrix} 3 & -5 & -2 \\ 1 & 4 & 1 \\ 6 & -3 & -5 \end{pmatrix}.$$

Checking every cell for a profitable unilateral deviation shows that Example 5.2 has no *PSNE, while this 3×3 exercise has exactly one, at (M, R).*

5.3 Saddle Points

The PSNE of a matrix game has a convenient equivalent description purely in terms of the single matrix u.

Definition 5.4 (Saddle point) *(s_1^*, s_2^*) is a **saddle point** of a matrix u if*

$$u(s_1^*, s_2^*) \geqslant u(s_1, s_2^*) \quad \forall s_1 \in S_1 \qquad \text{and} \qquad u(s_1^*, s_2^*) \leqslant u(s_1^*, s_2) \quad \forall s_2 \in S_2,$$

i.e., $u(s_1^*, s_2^*)$ is simultaneously the maximum in its column and the minimum in its row: the largest value player 1 (the row player) can secure there, and the smallest value player 2 (the column player) can hold player 1 to.

In Example 5.2, no cell is simultaneously a row-minimum and a column-maximum, so there is no saddle point – matching the absence of a PSNE we found directly. In Example 5.3, the cell (M, R), value 1, is exactly this: it is the minimum of row M ($\min(1, 4, 1)$, minimum 1 at both T and R column entries) and it is also the column-maximum at column R ($\max(-2, 1, -5)$). In short, the same cell (M, R) that we identified as the PSNE is also the saddle point.

Theorem 5.5 *In a matrix game with utility matrix u, (s_1^*, s_2^*) is a saddle point if and only if it is a PSNE.*

Proof: By Definition 5.4, (s_1^*, s_2^*) is a saddle point iff $u(s_1^*, s_2^*) \geqslant u(s_1, s_2^*)$ $\forall s_1 \in S_1$ and $u(s_1^*, s_2^*) \leqslant u(s_1^*, s_2)$ $\forall s_2 \in S_2$. Since $u \equiv u_1 \equiv -u_2$, the second inequality (flipping sign, which reverses its direction) becomes $u_2(s_1^*, s_2^*) \geqslant u_2(s_1^*, s_2)$ $\forall s_2 \in S_2$. So the saddle-point conditions are exactly

$$u_1(s_1^*, s_2^*) \geqslant u_1(s_1, s_2^*) \ \forall s_1 \in S_1 \qquad \text{and} \qquad u_2(s_1^*, s_2^*) \geqslant u_2(s_1^*, s_2) \ \forall s_2 \in S_2,$$

which is precisely the statement that neither player has a profitable unilateral deviation, i.e., (s_1^*, s_2^*) is a PSNE. ■

5.4 Security Values: Maxmin and Minmax

Recall from last lecture the maxmin value of a player. In a matrix game the two players' security values acquire special notation and a clean relationship.

Definition 5.6 (Maxmin and minmax values of a matrix game)

$$\underline{v} = \max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2) \qquad (\textbf{\textit{maxmin}}, \textit{ player 1's security value}),$$

$$\overline{v} = \min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2) \qquad (\textbf{\textit{minmax}}, \textit{ the value player 2 can hold player 1 to}).$$

Since $u_2 = -u$, player 2's own maxmin value (in the sense of Lecture 4) is exactly $-\overline{v}$: Sentinel's guaranteed catch rate is the negative of Ghost's minmax value.

Lemma 5.7 *For any matrix game, $\overline{v} \geqslant \underline{v}$.*

Proof: By definition of min, for every s_1, s_2,

$$u(s_1, s_2) \geqslant \min_{t_2 \in S_2} u(s_1, t_2).$$

The right-hand side does not depend on s_2, and the inequality holds for every s_1, so taking max over s_1 on both sides (the right-hand side being a constant with respect to s_1, this dominates for every particular s_1):

$$\max_{t_1 \in S_1} u(t_1, s_2) \geqslant \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \qquad \forall s_2 \in S_2.$$

The right-hand side is now a constant (independent of s_2), so taking min over s_2 of the left-hand side preserves the inequality:

$$\min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2) \geqslant \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2), \qquad \text{i.e.,} \qquad \overline{v} \geqslant \underline{v}.$$

■

Revisiting our two running examples with actual numbers: in Example 5.2, row minimums are $-1, -1$ so $\underline{v} = -1$, and column maximums are $1, 1$ so $\overline{v} = 1$; since $\overline{v} = 1 > -1 = \underline{v}$, Ghost and Sentinel must keep guessing – no PSNE, consistent with Theorem 5.5. In Example 5.3, one can check $\overline{v} = \underline{v} = 1$, attained at (M, R): with three vectors on the table, Sentinel's rate-limiting defense settles the exercise.

5.5 When Does a Matrix Game Have a PSNE?

The two examples above suggest that $\overline{v} = \underline{v}$ is exactly the condition for a PSNE to exist. We now make this precise, and identify *which* strategies form the equilibrium.

Define

$$s_1^* \in \arg\max_{s_1 \in S_1} \min_{s_2 \in S_2} u(s_1, s_2) \qquad \textbf{(maxmin strategy of player 1)},$$

$$s_2^* \in \arg\min_{s_2 \in S_2} \max_{s_1 \in S_1} u(s_1, s_2) \qquad \textbf{(minmax strategy of player 2)}.$$

Theorem 5.8 (PSNE existence theorem for matrix games) *A matrix game has a PSNE (equivalently, a saddle point) if and only if $\overline{v} = \underline{v} = u(s_1^*, s_2^*)$, where s_1^* and s_2^* are, respectively, a maxmin strategy of player 1 and a minmax strategy of player 2.*

Corollary 5.9 *Whenever $\overline{v} = \underline{v}$, the profile (s_1^*, s_2^*) built from the maxmin and minmax strategies above is itself a PSNE.*

Proof: ($\Rightarrow$) Suppose (s_1^*, s_2^*) is a PSNE. Since (s_1^*, s_2^*) is a PSNE, $u(s_1^*, s_2^*) \geqslant u(s_1, s_2^*)$ for every $s_1 \in S_1$, so

$$u(s_1^*, s_2^*) \geqslant \max_{t_1 \in S_1} u(t_1, s_2^*) \geqslant \min_{t_2 \in S_2} \max_{t_1 \in S_1} u(t_1, t_2) = \overline{v},$$

the second inequality because s_2^* is one particular strategy in the minimization defining $\overline{v}$. By the symmetric argument applied to player 2 (using that (s_1^*, s_2^*) is also a best response for player 2, and $u_2 = -u$), we get $\underline{v} \geqslant u(s_1^*, s_2^*)$. But Lemma 5.7 gives $\overline{v} \geqslant \underline{v}$, so chaining all three:

$$u(s_1^*, s_2^*) \geqslant \overline{v} \geqslant \underline{v} \geqslant u(s_1^*, s_2^*) \implies u(s_1^*, s_2^*) = \overline{v} = \underline{v}.$$

This chain of equalities also forces s_1^* and s_2^* themselves to be a maxmin strategy for player 1 and a minmax strategy for player 2, respectively (each inequality above is in fact tight at exactly these strategies).

($\Leftarrow$) Suppose $\overline{v} = \underline{v} = u(s_1^*, s_2^*)$ where s_1^*, s_2^* are maxmin/minmax strategies as defined above. For any $s_2 \in S_2$,

$$u(s_1^*, s_2) \geqslant \min_{t_2 \in S_2} u(s_1^*, t_2) = \max_{t_1 \in S_1} \min_{t_2 \in S_2} u(t_1, t_2) = \underline{v},$$

the middle equality because s_1^* is a maxmin strategy. By the symmetric argument, $u(s_1, s_2^*) \leqslant \overline{v}$ for every $s_1 \in S_1$. Since $\underline{v} = \overline{v} = u(s_1^*, s_2^*)$ by hypothesis, substituting gives $u(s_1^*, s_2) \geqslant u(s_1^*, s_2^*)$... more precisely, reading the two displayed inequalities the other way: $u(s_1^*, s_2^*) \geqslant u(s_1, s_2^*)\ \forall s_1$ (from the second display, since $\overline{v} = u(s_1^*, s_2^*)$) and $u(s_1^*, s_2^*) \leqslant u(s_1^*, s_2)\ \forall s_2$ (from the first display, since $\underline{v} = u(s_1^*, s_2^*)$). This is exactly the saddle-point condition of Definition 5.4, so by Theorem 5.5, (s_1^*, s_2^*) is a PSNE. ■

5.6 Mixed Strategies

Example 5.2 has no PSNE: whatever channel Ghost commits to, Sentinel profits by matching it, and vice versa. The only way out is for both to be *unpredictable*. This motivates randomizing over pure strategies.

Definition 5.10 (Mixed strategy) *For a finite set A, let $\Delta A = \{p \in [0,1]^{|A|} : \sum_{a \in A} p(a) = 1\}$ be the set of probability distributions over A. A **mixed strategy** for player i is an element $\sigma_i \in \Delta S_i$, i.e., a function $\sigma_i : S_i \to [0,1]$ with $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.*

Since players choose independently (we remain in the non-cooperative setting), the joint probability of player 1 playing s_1 and player 2 playing s_2 is $\sigma_1(s_1)\sigma_2(s_2)$. Utility at a mixed profile is defined as the corresponding expectation over utilities at pure profiles:

$$u_i(\sigma_1, \dots, \sigma_n) = \sum_{s_1 \in S_1} \cdots \sum_{s_n \in S_n} \sigma_1(s_1) \cdots \sigma_n(s_n)\, u_i(s_1, \dots, s_n).$$

We overload u_i to mean the utility at pure *or* mixed strategies; since utility at a mixed profile is an expectation, all the usual rules of expectation – linearity, conditioning, etc. – apply.

Example 5.11 (Computing an expected utility) *Suppose in Example 5.2, Ghost probes L with probability $2/3$ and R with probability $1/3$, while Sentinel watches L with probability $4/5$ and R with probability $1/5$. Ghost's expected utility is*

$$u_1(\sigma_1, \sigma_2) = \frac{2}{3} \cdot \frac{4}{5} \cdot (-1) + \frac{2}{3} \cdot \frac{1}{5} \cdot (1) + \frac{1}{3} \cdot \frac{4}{5} \cdot (1) + \frac{1}{3} \cdot \frac{1}{5} \cdot (-1) = -\frac{8}{15} + \frac{2}{15} + \frac{4}{15} - \frac{1}{15} = -\frac{3}{15} = -\frac{1}{5}.$$

5.7 Summary and Preview

We introduced matrix (two-player zero-sum) games and showed that a PSNE is exactly a saddle point of the payoff matrix (Theorem 5.5); proved that minmax never falls below maxmin (Lemma 5.7); and characterized exactly when a PSNE exists, namely when maxmin equals minmax, with the equilibrium given by the maxmin/minmax strategy pair (Theorem 5.8). When maxmin and minmax disagree, as in Example 5.2, no pure equilibrium exists and unpredictability is essential, which led us to mixed strategies and the computation of expected utility (Example 5.11). Next lecture: mixed strategy Nash equilibrium, a characterization theorem that lets us actually compute one, and Nash's 1951 theorem guaranteeing one always exists.

References

[MSZ13] M. Maschler, E. Solan and S. Zamir, *Game Theory*, Chapter 4, Cambridge University Press, 2013.

[Nas51] J. Nash, "Non-cooperative games," *Annals of Mathematics*, 54(2), 1951, pp. 286–295.

[vN28] J. von Neumann, "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, 100(1), 1928, pp. 295–320.