Zero-Sum Games and the Minimax Theorem

Game Theory · Module 1 — Strategic Form Games

Node 1.5

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.5

Zero-sum games are the original setting of game theory. Von Neumann's 1928 minimax theorem is the first major existence result in the field and is historically the moment the subject was born. It shows that every finite two-player zero-sum game has a *value* — a unique number that the row player can guarantee from below and the column player can guarantee from above — and that this value is achieved by a pair of mixed strategies that are mutual best responses. This node is the specialized prelude to the general Nash existence theorem (2.4): zero-sum minimax is the case where the equilibrium computation reduces to a linear program. It is also the mathematical home of modern poker solving, because the two-player heads-up subgame is a zero-sum game.

Formal Statement

Let $G = (\{1,2\}, S_1, S_2, u_1, u_2)$ be a finite two-player game with $u_1 + u_2 = 0$. Writing $A_{ij} = u_1(s_i, s_j)$ for the payoff matrix, player 1 maximizes A and player 2 minimizes it.

Define the **maximin value**

$$\underline{v} = \max_{\sigma_1 \in \Delta(S_1)} \min_{\sigma_2 \in \Delta(S_2)} \sigma_1^\top A \sigma_2,$$

and the **minimax value**

$$\overline{v} = \min_{\sigma_2 \in \Delta(S_2)} \max_{\sigma_1 \in \Delta(S_1)} \sigma_1^\top A \sigma_2.$$

Minimax Theorem (von Neumann, 1928). For every finite zero-sum game,

$$\underline{v} = \overline{v}.$$

This common value is the **value of the game** $v(G)$, and the maximin–minimax strategies form a Nash equilibrium.

Two crucial facts:

1. The inequality $\underline{v} \le \overline{v}$ is trivial (swapping max and min never improves the objective).
2. The reverse inequality $\overline{v} \le \underline{v}$ is the non-trivial theorem and requires the lift from S_i to $\Delta(S_i)$. In pure strategies, $\underline{v} < \overline{v}$ is possible (a "saddle point" need not exist).

Intuition

In a zero-sum game, one player's gain is the other's loss. There is no "win–win"; strategic interaction is pure conflict. In that setting, the question "what should I play?" has a particularly clean answer: play the strategy that maximizes the *worst* outcome your opponent can impose on you.

The maximin strategy is the "paranoid" strategy: assume the opponent will play to hurt you as much as possible, and choose accordingly. In a zero-sum game, that is also the "rational" strategy, because the opponent really *is* trying to minimize your payoff — that is literally what they maximize.

The minimax theorem says this paranoid reasoning has a fixed point: there is a value that both players can simultaneously guarantee against the other's worst behavior. This value is not too low (row can force it) and not too high (column can force it against row). It is unique.

The linear-programming reformulation (below) is both the fastest proof and the fastest computational method: finding the value of a zero-sum game is equivalent to solving a linear program.

Structural Role

Zero-sum minimax is the simplest non-trivial case of Nash existence:

- **1.4 ⇒ 1.5.** Maximin strategies are best-response-optimal; minimax-optimal strategies are mutual best responses in the zero-sum context.
- **1.5 ⇒ 1.6.** Minimax provides a worst-case lower bound for the normal-form representation of any game.
- **1.5 ⇒ 2.1.** The unique minimax pair is a Nash equilibrium of the zero-sum game.
- **1.5 ⇒ 2.4.** The Nash existence theorem generalizes the zero-sum case to general-sum games, replacing linear programming duality with a fixed-point theorem.
- **1.5 ⇒ 9.2.** No-regret learning on zero-sum games converges to the value; this is a dynamic realization of the static minimax result.

Zero-sum games are also the unique setting in which "equilibrium" and "security" coincide: the minimax strategy is both a Nash equilibrium component and the solution to the paranoid decision problem. In general-sum games, these two objects diverge, and the strategic question "what should I play?" becomes ambiguous.

Mathematical Development

Proof via LP duality

Player 1 wants to maximize $\min_j(\sigma_1^\top A)_j$. Introducing v as the minimum, this becomes

$$\max_{\sigma_1, v} \; v \quad \text{s.t.} \quad (\sigma_1^\top A)_j \ge v \text{ for all } j, \; \sum_i \sigma_1(i) = 1, \; \sigma_1 \ge 0.$$

This is a linear program. Its dual is

$$\min_{\sigma_2, w} \; w \quad \text{s.t.} \quad (A\sigma_2)_i \le w \text{ for all } i, \; \sum_j \sigma_2(j) = 1, \; \sigma_2 \ge 0.$$

This is exactly the minimax LP for player 2. LP duality (strong duality holds under feasibility) gives $v^* = w^*$, which is the statement of the theorem.

The saddle-point characterization

A pair (σ_1^*, σ_2^*) is a **saddle point** of A if

$$\sigma_1^\top A\sigma_2^* \le \sigma_1^{*\top} A\sigma_2^* \le \sigma_1^{*\top} A\sigma_2 \quad \text{for all } \sigma_1, \sigma_2.$$

Minimax equilibria are exactly saddle points. The saddle-point structure is what makes zero-sum games tractable: the interests of the two players are perfectly anti-aligned, so a single function's saddle encodes both best-response conditions.

Security level and value

For player 1, the **security level** is $\max_{\sigma_1} \min_{\sigma_2} \sigma_1^\top A\sigma_2 = \underline{v}$. The minimax theorem says the security level is attained and equals $\overline{v}$. In pure strategies:

$$\underline{v}_{\text{pure}} = \max_i \min_j A_{ij}, \qquad \overline{v}_{\text{pure}} = \min_j \max_i A_{ij}.$$

These are the "pure-strategy security levels." Without mixing, $\underline{v}_{\text{pure}} < \overline{v}_{\text{pure}}$ is possible. Matching Pennies is the canonical example: $\underline{v}_{\text{pure}} = -1, \overline{v}_{\text{pure}} = +1$, but $v = 0$ in mixed strategies.

Continuity and convexity arguments

A second (more general) proof: the function $f(\sigma_1, \sigma_2) = \sigma_1^\top A\sigma_2$ is bilinear, hence continuous, concave in σ_1, and convex in σ_2 on the compact convex domains. Sion's minimax theorem then gives $\max_{\sigma_1} \min_{\sigma_2} f = \min_{\sigma_2} \max_{\sigma_1} f$. This is the "convex-analytic" form of the theorem and lifts directly to infinite action spaces under appropriate continuity assumptions.

Computing the value

For a 2×2 zero-sum game with no pure saddle, the value is found by the indifference principle. Let $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, with the interior mixed equilibrium. Row's equilibrium mix $(p, 1-p)$ satisfies

$$pa + (1-p)c = pb + (1-p)d, \quad p = \frac{d-c}{a-b-c+d}.$$

The value is $v = pa + (1-p)c = \frac{ad-bc}{a-b-c+d}$.

Worked Example

A 3×3 zero-sum game

$$A = \begin{pmatrix} 2 & -1 & 0 \\ -1 & 3 & -2 \\ 0 & -2 & 4 \end{pmatrix}.$$

Pure-strategy analysis. $\underline{v}_{\text{pure}} = \max_i \min_j A_{ij} = \max(-1, -2, -2) = -1$. $\overline{v}_{\text{pure}} = \min_j \max_i A_{ij} = \min(2, 3, 4) = 2$. Since $-1 < 2$, no pure saddle exists.

LP for Row. Maximize v subject to $2p_1 - p_2 + 0p_3 \ge v$, $-p_1 + 3p_2 - 2p_3 \ge v$, $0p_1 - 2p_2 + 4p_3 \ge v$, $p_1 + p_2 + p_3 = 1$, $p_i \ge 0$. (Solving this LP is a computer exercise; the value of the game comes out to approximately $v \approx 0.769$.)

Interpretive takeaway. Row has a "strong" row 1 (positive against col 2 is its worst: -1), and a "strong" row 3 (positive against col 3 with $+4$ but weak against col 2 with -2). The equilibrium balances these to achieve a guaranteed payoff that neither row alone can secure.

A clean analytical case

Consider

$$A = \begin{pmatrix} 0 & -1 & 3 \\ 2 & 3 & -1 \\ -3 & 1 & 2 \end{pmatrix}.$$

This can be solved via the LP or by inspection of support cardinality. In practice, for anything bigger than 2×2 by hand, the LP is the way.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Zero-sum assumption	Heads-up cash poker (minus rake) is zero-sum
Minimax value $v(G)$	The expected chip EV of GTO play against GTO opponent — often exactly 0 modulo position
Maximin strategy	The "unexploitable" strategy — guarantees at least the game value regardless of opponent
LP formulation	Linear-programming solutions of HU limit poker — CFR is an iterative alternative
Saddle point	HU GTO pair — mutual best responses where neither player can improve unilaterally

Concrete situation

Heads-up limit hold'em (HULHE) was solved to optimality in 2015 by the Cepheus program (Bowling et al., 2015) via counterfactual regret minimization, a no-regret algorithm that provably converges to the minimax value of a zero-sum game. The solution is a set of mixed strategies — one per information set — that together constitute a near-optimal pair of saddle-point strategies.

Concrete number: the HULHE value for the small-blind (dealer) is approximately $+0.088$ BB per hand, rather than exactly 0. Why? Because the ante/blind structure is not perfectly symmetric between the two players; the SB posts half a bet and acts last postflop, and these positional asymmetries combine to produce a small positive value for the button. In the strict zero-sum sense, BB's value is -0.088 BB per hand. The game is solved: if you play the Cepheus strategy, you cannot lose more than ϵ per hand, no matter who your opponent is.

For heads-up **no-limit** hold'em, the game is too big to solve exactly, but Libratus (2017) and DeepStack (2017) both used CFR variants and beat top human professionals. The game value is approximately 0 for SB (slightly positive) in no-limit as well; the exact number is not known because an exact solution is infeasible.

What this understanding buys you

Heads-up poker has a well-defined value. Whatever you believe about your edge at the table, the ceiling is the difference between your win rate and the game value. If HULHE has a value of $+0.088$ BB per hand for the button, no button player can ever sustainably win more than that per hand against equally strong competition. The value of the game is the upper bound on your edge at infinite skill.

Unexploitable = at the minimax value. When you hear "play GTO," what is being asked is that you play a strategy achieving the minimax value — a strategy against which no opponent can do better than the zero-sum value of the game. You give up exploitation potential against weak players, but you also eliminate every losing counter-strategy. This is the exact tradeoff the minimax theorem makes precise: the value is what you can guarantee, and no more.

Cross-Links

- **1.1 Players, strategies, payoffs** — the underlying objects.
- **1.2 Mixed strategies** — the lift required for the theorem.
- **1.4 Best response** — minimax strategies are mutual best responses.
- **2.1 Nash equilibrium** — minimax is the zero-sum case.
- **2.4 Nash existence theorem** — generalization to general-sum.
- **9.2 No-regret learning** — time-averaged dynamics converge to minimax value.
- **9.3 CFR** — the algorithm that solved HU limit poker.
- **Linear programming / convex optimization** — the computational backbone.

Structural Compression

Invariant: In every finite two-player zero-sum game, the saddle-point value exists, is unique, and equals both the maximin and minimax value; it is attained by a pair of mixed strategies that form mutual best responses.

Mechanism: Linear-programming duality (or Sion's minimax theorem) applied to the bilinear payoff function on the product of simplices; strong duality equates the primal optimum (maximin) to the dual optimum (minimax).

Cross-System Echo: Minimax is the mathematical skeleton of adversarial settings throughout computer science: GANs train at a saddle point of generator vs discriminator; robust optimization is a minimax over worst-case parameter realizations; adversarial examples in machine learning are the inner maximization problem. All of these are lifts of the zero-sum minimax theorem to richer function spaces.

Exercises

1. Verify that Matching Pennies has pure-strategy security levels $\underline{v}_{\text{pure}} = -1$ and $\overline{v}_{\text{pure}} = +1$, and mixed-strategy value $v = 0$.
2. Solve the 2×2 zero-sum game $A = \begin{pmatrix} 3 & -1 \\ -2 & 4 \end{pmatrix}$ for the value and optimal mixed strategies via the indifference principle.
3. Write out the primal LP and the dual LP for a general $m \times n$ zero-sum game and verify that they are duals of each other.
4. Prove the trivial inequality $\underline{v} \leq \overline{v}$ from scratch, without using LP duality.
5. Show that in a zero-sum game, any maximin strategy of Row and any minimax strategy of Column together form a Nash equilibrium.
6. Construct a 3×3 zero-sum game with a pure saddle point. Show that any iterated strict-dominance argument would find the saddle directly.
7. **Argue, structurally,** why the minimax theorem requires mixed strategies in general but a pure saddle suffices in games with a dominant strategy. Specifically: identify which hypothesis of LP strong duality fails in the pure-strategy formulation of Matching Pennies, and explain how the lift to mixed strategies restores it.

Game Theory · Module 1 — Strategic Form Games · Node 1.5
Concepts: nash-equilibrium, dominance, convexity, constrained-optimization

Prerequisites: 1.1, 1.2, 1.4
Enables: 1.6, 2.1, 2.4, 9.2

hbar.university — own.your.cognition