A simple proof of Von Neumann's minimax theorem for the existence of Nash equilibrium in a two-players zero-sum game.

Rémi Lajugie

mail prenom.nom@ac-limoges.fr

April 2024

Abstract

In this note, we present a simple demonstration (using only elementary mathematical tools which are accessible to undergraduate students) of Von Neumann's minimax theorem for two-player zero-sum games. This is an important special case of the theorem asserting the existence of a Nash equilibrium for a game.

1 Introduction

Von Neumann's minimax theorem forms the foundation of much of modern game theory. It demonstrates the existence of a Nash equilibrium in the context of two-player zero-sum games. There are several proofs of this theorem. The original proof relies on a clever interpretation of the study set [von Neumann, 1928], while most modern approaches are based on linear programming duality [Boyd and Vandenberghe, 2004] or on approximating game strategies using reinforcement learning algorithms [Cesa-Bianchi and Lugosi, 2006].

However, these proofs are beyond the reach of undergraduate tools, significantly hindering the dissemination of game theory in student curricula and necessitating the acceptance of the most fundamental results without proof.

In this note, we present a proof of Von Neumann's minimax theorem that uses only arguments accessible to undergraduate students. Our proof does not claim to be the only elementary proof of this result and has interesting connections to the proof proposed by Jean Ville in the 1930s, cf. [Ville, 1938] or [Sion, 1958].

2 Setting

We consider two players, Alice and Bob. They compete in a game with the following properties:

- Alice has a set $\mathcal{A}$ of n_1 actions at her disposal, and Bob has a set $\mathcal{B}$ of n_2 actions at his disposal.
- They play simultaneously and immediately receive a reward, the sum of which is zero. This means that what Bob gains is lost by Alice. Thus, we can naturally introduce the gain function $G : \mathcal{A} \times \mathcal{B} \mapsto \mathbb{R}$ which associates the reward obtained from Alice's point of view for each pair of possible actions.

Such a game can be represented by a gain matrix $A \in \mathbb{R}^{n_1 \times n_2}$. In position $A_{i,j}$ we represent the gain value corresponding to Alice's i-th action and Bob's j-th action.

For example, in a game where each player has 2 actions, we can imagine the following gain matrix:

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

Obviously, each player seeks to maximize their average gain.

Let's introduce the notion of a *strategy*.

Definition 1. *A strategy for Alice (resp. Bob) is a vector $u \in \mathbb{R}^{n_1}$ (resp. $\in \mathbb{R}^{n_2}$), such that $u_i \geq 0$ for all i and $\sum_{i=1}^{n_1} u_i = 1$ (resp. $\sum_{i=1}^{n_2} u_i = 1$).*

We denote Δ_{n_1} as the set of Alice's strategies and Δ_{n_2} as the set of Bob's strategies.

Fact 1. *The set of strategies is a convex set.*

Fact 2. *Given two strategies $p \in \mathbb{R}^{n_1}$ and $q \in \mathbb{R}^{n_2}$, the average gain for Alice and Bob if they play according to these strategies is $p^\top Aq$.*

Definition 2. *A pair of strategies $(p^*, q^*) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ is called a* Nash equilibrium *if neither player has an incentive to change their strategy.*

In other words, ${p^*}^\top Aq^* = \max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_2}} p^\top Aq$ (Alice has no incentive to change her strategy assuming Bob tries to minimize the opponent's gain) AND ${p^*}^\top Aq^* = \min_{q\in\Delta_{n_2}} \max_{p\in\Delta_{n_1}} p^\top Aq$ (Bob also has no incentive to change his strategy).

3 The demonstration

We will demonstrate the following result, known as Von Neumann's minimax theorem.

Theorem 1. *If $A \in \mathbb{R}^{n_1\times n_2}$, Δ_{n_1}, and Δ_{n_1} are the convex sets corresponding to the possible strategies (probability distributions), then:*

$$\max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq = \min_{q\in\Delta_{n_1}} \max_{p\in\Delta_{n_1}} p^\top Aq$$

3.1 Some Preliminary Facts

Fact 3. *For fixed q, the function $f(p) = p^\top Aq$ has a maximum on Δ_{n_1}.*

Proof. The set Δ_{n_1} is closed and bounded in finite dimension. Hence, it is compact and the function $p \mapsto p^\top Aq$ reaches a maximum on Δ_{n_1}. □

Fact 4. *Let $p \in \Delta_{n_1}$, then the mapping $\tau : \eta \in \mathbb{R}^{n_2} \mapsto p^\top A\eta$ is K-Lipschitz with K independent of the choice of p.*

Proof. We endow the space with the infinity norm. $p^\top A$ is a row vector such that $(p^\top A)_i = \sum p_k A_{ki}$. Since each $0 \leq p_k \leq 1$, we can crudely bound $(p^\top A)_i$ by $K = \max_{kl}(A_{kl})$. Therefore, $|\tau(\eta)| \leq K\|\eta\|_\infty$. □

Fact 5. *The function $f : q \mapsto \max_{p\in\Delta_{n_1}} p^\top Aq$ is continuous on Δ_{n_1}.*

Proof. Let $q_0 \in \Delta_{n_1}$, then by Fact 3, there exists $p_0 \in \Delta_{n_1}$ such that $f(q_0) = {p_0}^\top Aq_0$.

- Note that ${p_0}^\top Aq_0 = \lim_{q\to q_0} {p_0}^\top Aq$, and also ${p_0}^\top Aq \leq f(q)$ for all q, so by preserving the wide limits, $f(q_0) \leq \lim_{q\to q_0} f(q)$. We still need to verify the reverse inequality.
- Let $q_1 = q_0 + \eta$ where $\eta \in \mathbb{R}^{n_2}$ is a perturbation. Again, by Fact 3, $f(q_1) = {p_1}^\top Aq_1$ for some $p_1 \in \Delta_{n_1}$. We note that $|{p_1}^\top A\eta| \leq K\|\eta\|_\infty$ by Fact 4 with K **independent** of the choice of p_1.

  Thus, for $\epsilon > 0$, we can find a ball centered at 0 such that for any $p_1 \in \Delta_{n_1}$, $|{p_1}^\top A\eta| \leq K\|\eta\|_\infty \leq \epsilon$.

  This ensures that $f(q_1) = {p_1}^\top Aq_1 = {p_1}^\top Aq_0 + {p_1}^\top A\eta \leq {p_0}^\top Aq_0 + \epsilon = f(q_0) + \epsilon$.

  By passing to the limit in the inequalities, we thus obtain $f(q_0) + \epsilon \geq \lim_{q_1\to q_0} f(q_1)$.

  We can conclude by letting ϵ tend to 0.

□

Fact 6. *The quantity $\min_{q\in\Delta_{n_1}} \max_{p\in\Delta_{n_1}} p^\top Aq$ exists. The same holds for $\max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq$.*

3.2 The Core of the Proof

Fact 7. *We have:*

$$\max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq \leq \min_{q\in\Delta_{n_1}} \max_{p\in\Delta_{n_1}} p^\top Aq.$$

Proof. This is a universal result. Indeed, we always have for any $\tilde{q} \in \Delta_{n_1}$, $\max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq \leq \max_{p\in\Delta_{n_1}} p^\top A\tilde{q}$ and this remains true when taking the min over $\tilde{q}$ on the right-hand side. □

Let $\alpha = \max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq$ and $\beta = \min_{q\in\Delta_{n_1}} \max_{p\in\Delta_{n_1}} p^\top Aq$. The difference $\beta - \alpha$ is called the *duality gap*. We will demonstrate that it is zero.

Fact 8. *There exist $p_1, p_2 \in \Delta_{n_1}$ and $q_1, q_2 \in \Delta_{n_1}$ such that*

$$\alpha = {p_1}^\top Aq_1 \quad \textit{and} \quad \beta = {p_2}^\top Aq_2.$$

Proof. This is a consequence of the extreme value theorem. □

We can now complete the proof of Theorem 1.

Proof. Using the notations from Fact 8, we see that

$${p_1}^\top Aq_1 \leq {p_2}^\top Aq_2.$$

Introduce the function

$$f : \lambda \in [0,1] \mapsto {p_2}^\top A\left(\lambda q_2 + (1-\lambda)q_1\right).$$

- It is a continuous function in λ.
- We have $f(0) = {p_2}^\top Aq_1 \leq \max_{p\in\Delta_{n_1}} p^\top Aq_1 = {p_1}^\top Aq_1 = \alpha$.
- On the other hand, $f(1) = {p_2}^\top Aq_2 = \beta \geq \alpha$

Therefore, by the intermediate value theorem, there exists $\lambda \in [0,1]$ such that $\alpha = {p_2}^\top A\left(\lambda q_2 + (1-\lambda)q_1\right)$.
Furthermore, by the convexity of Δ_{n_1}, $\lambda q_2 + (1-\lambda)q_1 \in \Delta_{n_1}$. Thus, the previous point assures us that there exists $\tilde{q} \in \Delta_{n_1}$ such that $\alpha = {p_2}^\top A\tilde{q}$.
By definition, ${p_2}^\top A\tilde{q} \geq {p_2}^\top Aq_2 = \beta$.
We thus deduce that $\alpha \geq \beta$.
In conclusion: $\alpha = \beta$ and there is no duality gap. In other words:

$$\max_{p\in\Delta_{n_1}} \min_{q\in\Delta_{n_1}} p^\top Aq = \min_{q\in\Delta_{n_1}} \max_{p\in\Delta_{n_1}} p^\top Aq.$$

□

References

Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 1958.

Jean Ville. Sur la théorie générale des jeux ou intervient l'habileté des joueurs. *Traité du Calcul des Probabilités et des ses Applications', Paris, Gauthiers-Villars*, 171, 1938.

John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.