Two Elementary Proofs of the Minimax Theorem

Jonathan Weinstein*
Washington University in St. Louis

November 23, 2022

Abstract

We give a proof of the Minimax Theorem where the key step involves reducing the strategy sets. The proof is self-contained and elementary, avoiding appeals to theorems from geometry, analysis or algebra, such as the separating hyperplane theorem or linear-programming duality. The argument is valid with any ordered field in place of the usual field of real numbers. We give a second proof with similar merits which is closely modeled on an argument from Loomis (1946). This one is a bit simpler, but uses the existence of maxima on compact domains and hence is only valid in $\mathbb{R}$.

JEL code: C72

Keywords: Game theory; Minimax theorem

*Email: j.weinstein@wustl.edu. ORCID: 0000-0001-8245-2067. I thank the anonymous referees and associate editor Ben Brooks for their valuable advice.

1 Introduction

The Minimax Theorem was the first major existence theorem in game theory, and as such it paved the way for game theory as a modern mathematical discipline. Perhaps it can best be described to a modern audience as saying that a Nash Equilibrium exists for any finite two-player zero-sum game. For a statement closer to the original, which avoids the anachronistic reference to Nash, see the next section. Since the time the Minimax Theorem was first proved in von Neumann (1928), many other proofs have appeared. Perhaps best-known is the argument in von Neumann and Morgenstern (1947), which uses the separating hyperplane theorem. This approach, as well as von Neumann's original proof, used properties special to the real numbers, specifically the Weirstrass Theorem for continuous functions on compact sets. But the Minimax Theorem actually holds when payoffs lie in any ordered field $\mathbb{F}$, and probabilities in the mixed strategies are allowed to range over $\mathbb{F}$, as was shown in Kuhn and Tucker (1950). The proof there relies on a result from Stiemke (1915), which was a forerunner of duality results in linear programming. Indeed, linear-programming duality, or related results such as Farkas' Lemma, also leads to a quick proof of the Minimax Theorem – see for instance Vohra (2004).

I have often felt a bit unsatisfied by these proofs. Because the Minimax Theorem is so fundamental to our understanding of strategic interaction, I wanted to create a proof that reasons directly about strategies. That is, an argument that harnesses strategic intuition, avoiding the use of theorems from geometry, algebra or analysis as the primary engine. Such a proof is provided here, with some preliminaries in Section 2 and the proof of the workhorse lemma in Section 3. The proof is valid for any ordered field, because it avoids any arguments based on continuity or compactness. I can only hope that others feel as I do that it contributes understanding beyond the existing proofs. The most similar proof I have seen, pointed out by a referee, is in Peck (1958), which has a similar method of eliminating strategies. It relies on properties of the real numbers, however: the results are stated in terms of sup and inf over mixed strategy sets, which are guaranteed to exist only over the field of reals, and the proof begins with this existence.

If one is willing to use basic properties of the real numbers and compactness, there is a very simple argument which was first given by Loomis (1946). Loomis comments that von Neumann had given a recent lecture where he challenged the audience to give an elementary proof of the Minimax Theorem, or, more precisely, of an algebraic theorem which was a slight generalization of Minimax. Apparently von Neumann, too, thought there was value in finding a more elementary approach than his own proofs. Loomis succeeded at this admirably, giving a very simple proof. It relies on the Weierstrass Theorem to show that the max-min and min-max in (1) exist, which both calls into question how "elementary" it is, and means it is valid only for $\mathbb{R}$. Nonetheless, it is the simplest proof of the original theorem I have seen, subject to the constraint of using no advanced theorems other than Weierstrass. Certainly, it is far more elegant than von Neumann's 1928 proof. For purposes of comparison, I have included an adaptation of it here, in Section 4, which I believe has two advantages for the modern student of game theory over the original Loomis paper: I have stripped it of some extra generality, which simplifies the notation, and I have rewritten it in the notation of game theory, in a way which emphasizes strategic intuition. The resulting argument is simpler than my proof, but with the relative weakness of its appeal to compactness. The simplicity of this argument is evidenced by its rediscovery by Owen (1967) and Thomassen (2000). Binmore

(2004), inspired specifically by Owen, expands the argument, using transfinite induction, to prove a more general version of the minimax theorem with more general strategy spaces.

2 Definitions and Proof Outline

All of our games will be finite two-player zero-sum games, allowing a streamlining of notation. Such a game will be a triple $G = (A_1, A_2, u)$ where A_1, A_2 are finite sets and u is a function from $A_1 \times A_2$ to $\mathbb{F}$ for some ordered field $\mathbb{F}$.[1] We refer to elements of A_i as either *actions* or *pure strategies*. The function u is the utility of Player 1; Player 2's utility is $-u$. The game is *symmetric* if $A_1 = A_2$ and for all a_1 and a_2, $u(a_1, a_2) = -u(a_2, a_1)$.

For each player i, a *mixed strategy* is, formally, a non-negative-valued function from A_i to $\mathbb{F}$ which sums to 1, and Σ_i is the set of mixed strategies. In a slight overloading of notation, any $a_i \in A_i$ will sometimes be considered to be the mixed strategy assigning probability 1 to a_i. We extend the function u to $\Sigma_1 \times \Sigma_2$ by linearity as usual. For any $\sigma_i^1, \sigma_i^2 \in \Sigma_i$ and $\alpha \in [0, 1]$, the mixture $\alpha\sigma_i^1 + (1-\alpha)\sigma_i^2 \in \Sigma_i$ has the expected meaning, assigning probability $\alpha\sigma_i^1(a_i) + (1-\alpha)\sigma_i^2(a_i)$ to each a_i. [2]

The minimax theorem can then be stated as follows:

Theorem 1 (Minimax Theorem) *For any finite two-player zero-sum game G,*

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u(\sigma_1, \sigma_2) \tag{1}$$

Note that when we work in an arbitrary $\mathbb{F}$, there is no immediate reason that either side of (1) must be well-defined. It is part of the content of the theorem that they are. The common value in (1) is called the *value* of G. It is straightforward to see that the theorem is equivalent to the existence of a Nash equilibrium (σ_1^*, σ_2^*). Indeed, σ_1^* would optimize the left-hand-side, and σ_2^* the right-hand-side, with the common value $u(\sigma_1^*, \sigma_2^*)$.

We call a mixed strategy *bulletproof* if it guarantees non-negative utility. Our main lemma will be the following:

Lemma 1 (Bulletproof Lemma) *In any two-player zero-sum game, at least one player has a bulletproof strategy.*

From the Bulletproof Lemma, the Minimax Theorem follows in two simple and intuitive steps. First, the lemma easily implies the theorem for symmetric games. Indeed, the lemma along with symmetry tells us that *both* players have bulletproof strategies, and then both sides of (1) equal 0.

Next, we show that the symmetric case implies the general case.[3] Given $G = (A_1, A_2, u)$, we create a symmetric version where both players simultaneously play both roles and we

[1] The most familiar examples of ordered fields are $\mathbb{R}$ and subsets of $\mathbb{R}$ which are closed under the field operations, such as the rationals $\mathbb{Q}$. There are also ordered fields which cannot be embedded in $\mathbb{R}$, called *non-Archimedean*; such a field includes infinite and infinitesimal elements. One useful introduction to the topic is Propp (2013).

[2] The interval [0,1] should be interpreted as a subset of $\mathbb{F}$, not a subset of $\mathbb{R}$ as per usual.

[3] This reduction of the general case to the symmetric case appears in Kuhn and Tucker (1950), p. 81. The authors credit it to von Neumann. It is equivalent to flipping a coin to decide who plays what role. That article also uses a version of the Bulletproof Lemma on the way to the main result. The lemma is also very similar to Ville's Theorem, which has been used in many proofs of the minimax theorem. Our lemma is slightly weaker.

add the payoffs. That is, we let $\bar{G} = (A_1 \times A_2, A_1 \times A_2, \bar{u})$ where $\bar{u}((a_1, a_2), (a'_1, a'_2)) = u(a_1, a'_2) - u(a'_1, a_2)$. Then, a minimax strategy in $\bar{G}$ must have marginal distributions on each coordinate which solve the min-max and max-min problem of G and establish (1). Indeed, let σ be a max-min strategy for Player 1 in $\bar{G}$, which we know has value zero, and let its marginals be σ_1^*, σ_2^*. Player 2's best response, which is to best-respond in each coordinate, must give total value zero, i.e.

$$\min_{\sigma_2} u(\sigma_1^*, \sigma_2) - \max_{\sigma_1} u(\sigma_1, \sigma_2^*) = 0$$

and this implies that both players in G can guarantee value

$$v = \min_{\sigma_2} u(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u(\sigma_1, \sigma_2^*)$$

via the strategies (σ_1^*, σ_2^*).

So, the bulk of the work is in the proof of the lemma. It is here that I believe my approach is somewhat novel. The details are in the next section; here is the idea. We proceed by induction on the size of A_1, by an elimination step which shrinks both players' *mixed* strategy sets in a way which preserves the existence (or non-existence) of a bulletproof strategy for each player. Specifically, we remove one of P1's pure strategies a_1^*, and simultaneously remove those of P2's mixed strategies which lose to a_1^*. This is precisely the way to leave the "balance of power" unchanged: in P2's search for a bulletproof strategy, some options are foreclosed by a_1^*; if we remove a_1^*, we should remove these options as well. The reduction does not necessarily preserve the value of the game; it does preserve its sign. The remaining mixed strategies of P2 form a polytope; its corners are the new pure strategies for P2. This reduction may actually increase the number of pure strategies for P2, but always reduces those of P1, so that the induction is valid. Please note that while the validity of the proof does not rely on any knowledge of the properties of polytopes, such knowledge does help understand why it works. In particular, the set of P2's mixed strategies with non-negative payoff against a_1^* is the convex hull of finitely many points, and this helps the proof operate. Notice that the proof implies an algorithm for determining which player has a bulletproof strategy, and for constructing one, by iterating the elimination step.

Of course, the special role of utility level 0 in the lemma is arbitrary; by subtracting a constant from all payoffs, the lemma easily implies that for each $k \in \mathbb{F}$, either P1 can guarantee at least k or P2 can guarantee at least $-k$. If we knew that the max-min and min-max were well-defined, this would quickly yield the main result, but in absence of this knowledge we have relied on the symmetrization argument from Kuhn and Tucker (1950).

In most prominent examples of zero-sum games, utility level 0 does have a special significance. Sometimes, as in poker, it represents zero monetary transfer, i.e. the utility level a player would achieve by opting out of the game. In parlor games such as rock-paper-scissors, it represents a tie. In Borel (1921), the original article introducing the idea of mixed-strategy solutions to zero-sum games, Borel focused on games with only two final outcomes, winning and losing (but perhaps a role for chance after the players act). This avoided the need to think about expected utility, which wasn't invented yet. For Borel, players were simply maximizing the probability of a win. Here, utility 0 represents a .5 probability of winning. Thus, even though a bulletproof strategy certainly might not be the best strategy, it does in all these examples guarantee that one is not, on average, a "loser" in the game. Borel's work was focused exclusively on symmetric games, so his initial papers (though taken in retrospect as introducing the general minimax idea) actually defined the goal as guaranteeing non-negative value. He called such mixed strategies, which we call "bulletproof," simply "supérieures" (and claimed that they almost never existed, except with three or fewer strategies.) So, there is historical reason to give a flattering name to such strategies, though they may be suboptimal.

2.1 Further Comparison to Existing Proofs

The proof of the main lemma is reminiscent of the algorithm of Fourier-Motzkin (FM) elimination, which eliminates one variable while adding constraints, to preserve feasibility or infeasibility of a linear program. Consider the program of Player 1, trying to find a bulletproof strategy σ_1. This is a linear feasibility problem: the probability $\sigma_1(a_1)$ attached to each of his actions is a variable, and each of Player 2's actions creates a linear constraint, namely that $u(\sigma_1, a_2) \geq 0$. In FM elimination, at each step one variable is eliminated and a new set of constraints is created which may be larger in cardinality than before. This is similar to how P1's pure-strategy set shrinks by one element and P2's may grow in cardinality (though both sets of mixed strategies shrink). FM elimination can be used to prove Farkas' Lemma, and another proof by reduction in dimensionality appears in Gale (1960). The Farkas Lemma can be used to prove Ville's Theorem, which is essentially a slightly stronger version of the Bulletproof Lemma and has often been used in proofs of the minimax theorem.

The following proof, then, uses ideas similar to existing algebraic proofs to give an algorithm which, through natural steps of strategic elimination, always finds a bulletproof strategy for some player. I find it instructive to see how these ideas play out when the elimination is specialized and described in terms of strategies, and the elimination has the specific goal of preserving the existence of a bulletproof strategy for each player.

3 Proof of the Bulletproof Lemma

We will induct on $|A_1|$. Let $G = (A_1, A_2, u)$ be any two-player zero-sum game. If $|A_1| = 1$ the result is easy, so let $|A_1| \geq 2$. Pick any action $a_1^* \in A_1$; in order to apply the inductive hypothesis, we will construct a new game G' where Player 1 has action set $A_1' = A_1 - \{a_1^*\}$. Player 2's new action set A_2' will be a finite subset of Σ_2. To construct it, we first partition A_2 as follows:

$$A_2^g = \{a_2 \in A_2 : u(a_1^*, a_2) \leq 0\}, A_2^b = \{a_2 \in A_2 : u(a_1^*, a_2) > 0\}$$

That is, actions in A_2^g, the "good" actions, are those that defeat or tie the eliminated action a_1^*, and A_2^b, for "bad", lose to action a_1^*. If $A_2^g = \emptyset$, then a_1^* is bulletproof, so assume not. For each pair $(g, b) \in A_2^g \times A_2^b$, there is a unique mixture of g and b which has expected payoff 0 against a_1^*, which we denote $\langle g, b\rangle$, i.e.

$$\langle g, b\rangle := \frac{u(a_1^*, b)}{u(a_1^*, b) - u(a_1^*, g)} g + \frac{-u(a_1^*, g)}{u(a_1^*, b) - u(a_1^*, g)} b$$

Then we set

$$A_2' = A_2^g \cup \{\langle g, b\rangle : g \in A_2^g, b \in A_2^b\}$$

Notice that both players' action sets are subsets of their mixed strategy sets from G, so that payoffs in G' can be simply inherited from those in G. We now can apply the inductive hypothesis to $G' = (A_1', A_2', u)$.

The case that Player 2 has a bulletproof strategy σ_2 in G' is very simple. This same strategy σ_2 is bulletproof in G, because all actions in A_2' were chosen to have non-negative payoff against the missing action a_1^*.

If Player 1 has a bulletproof strategy σ_1 in G', there is a bit more work. If σ_1 is also bulletproof in G, we are done, so presume not; then let $A_2^{b'} = \{b \in A_2 : u(\sigma_1, b) < 0\} \neq \emptyset$. Note that $A_2^{b'} \subseteq A_2^b$ by definition of σ_1, because $A_2^g \subseteq A_2'$. We will construct an appropriate mixture $\sigma_1^\alpha := \alpha a_1^* + (1-\alpha)\sigma_1$ which is bulletproof in G. Note that by definition of "bad", a_1^* has positive payoff against bad actions, and since good actions are in A_2', σ_1 has non-negative payoff against good actions. Therefore, a sufficiently large α works against all bad actions, and a sufficiently small one against all good actions. The idea is to find an α which works against both, and an obvious candidate is the smallest α which works against all bad actions. We can restrict attention to $A_2^{b'}$; for other bad actions, any α suffices.

Indeed, for each action $b \in A_2^{b'}$, define α_b as the solution to the equation $u(\sigma_1^{\alpha_b}, b) = 0$, i.e.

$$\alpha_b = \frac{-u(\sigma_1, b)}{-u(\sigma_1, b) + u(a_1^*, b)}$$

Notice that $\alpha_b > 0$ by definition of $A_2^{b'}$, and also $\alpha_b < 1$.

Then, let $\alpha = \max_{b \in A_2^{b'}} \alpha_b = \alpha_{\bar{b}}$ for appropriate $\bar{b}$. By construction, $u(\sigma_1^\alpha, b) \geq 0$ for each $b \in A_2^b$. For any $g \in A_2^g$, recall that $u(a_1^*, \langle g, \bar{b} \rangle) = 0$, and also, because σ_1 is bulletproof in G', $u(\sigma_1, \langle g, \bar{b} \rangle) \geq 0$, which together imply (by linearity) $u(\sigma_1^\alpha, \langle g, \bar{b} \rangle) \geq 0$. But also $u(\sigma_1^\alpha, \bar{b}) = 0$, so we must have $u(\sigma_1^\alpha, g) \geq 0$, and the proof is complete. $\square$.

4 The proof of Loomis (1946) for $\mathbb{F} = \mathbb{R}$

If the theorem fails, let $G = (A_1, A_2, u)$ be a counterexample for which $|A_1|+|A_2|$ is smallest. Because Σ_1, Σ_2 are compact, the expressions in (1) are well-defined. Let σ_1^* optimize the left-hand-side (the max-min problem) and σ_2^* the right (the min-max problem), and call the optimal values $\underline{v}$ and $\bar{v}$. It is immediate that

$$\underline{v} \leq u(\sigma_1^*, \sigma_2^*) \leq \bar{v}$$

and by assumption, at least one inequality is strict; say the first. Then there is some a_2^b (a "bad" strategy for Player 2) with $\underline{v} < u(\sigma_1^*, a_2^b)$. Form G' from G by eliminating a_2^b. By inductive hypothesis, (1) holds for G', say with common value v. The min-max can only go up when we eliminate one of Player 2's strategies, so $\bar{v} \leq v$ and thus $\underline{v} < v$. Let σ_1^{**} solve the max-min problem in G'. It is then easy to check that for small enough $\varepsilon > 0$, the strategy $\sigma_1^\varepsilon := (1-\varepsilon)\sigma_1^* + \varepsilon\sigma_1^{**}$ improves on the alleged solution σ_1^* to the max-min problem of G. Indeed, the strict inequality $\underline{v} < u(\sigma_1^\varepsilon, a_2^b)$ holds for $\varepsilon = 0$, hence also in a neighborhood of 0, and for $a_2 \neq a_2^b$ we simply have, for *any* $\varepsilon \in (0,1)$,

$$\underline{v} < (1-\varepsilon)\underline{v} + \varepsilon v \leq (1-\varepsilon)u(\sigma_1^*, a_2) + \varepsilon u(\sigma_1^{**}, a_2) = u(\sigma_1^\varepsilon, a_2)$$

The contradiction proves the desired result. □

This proof requires that the underlying field be $\mathbb{R}$ in the first paragraph, in order to apply compactness and conclude that the max-min and min-max exist. The implicit use of continuity in the final paragraph is only for convenience and could be replaced by field operations.

References

Binmore, K. (2004): "Guillermo Owen's proof of the minimax theorem," in *Essays in Cooperative Games*, pp. 19–23. Springer.

Borel, E. (1921): "La théorie du jeu et les équations intégrale à noyau symétrique," *Comptes Rendu de l'Académie des Sciences*, 173, 1304–1308.

Gale, D. (1960): *The theory of linear economic models.* University of Chicago press.

Kuhn, H. W., and A. W. Tucker (1950): *Contributions to the Theory of Games, Volume I.* Princeton University Press.

Loomis, L. H. (1946): "On a theorem of von Neumann," *Proceedings of the National Academy of Sciences of the United States of America*, 32(8), 213.

Owen, G. (1967): "Communications to the Editor—An Elementary Proof of the Minimax Theorem," *Management Science*, 13(9), 765–765.

Peck, J. (1958): "Yet another proof of the minimax theorem," *Canadian Mathematical Bulletin*, 1(2), 97–100.

Propp, J. (2013): "Real analysis in reverse," *The American Mathematical Monthly*, 120(5), 392–408.

Stiemke, E. (1915): "Über positive Lösungen homogener linearer Gleichungen," *Mathematische Annalen*, 76(2), 340–342.

Thomassen, C. (2000): "The rendezvous number of a symmetric matrix and a compact connected metric space," *The American Mathematical Monthly*, 107(2), 163–166.

Vohra, R. V. (2004): *Advanced mathematical economics*. Routledge.

von Neumann, J. (1928): "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, 100, 295–320.

von Neumann, J., and O. Morgenstern (1947): *Theory of Games and Economic Behavior*. Princeton University Press, second edn.