



Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Minimax Property in Infinite Two-Person Win-Lose Games

Ron Holzman

To cite this article:

Ron Holzman (2025) The Minimax Property in Infinite Two-Person Win-Lose Games. Mathematics of Operations Research 50(3):2287-2300. <https://doi.org/10.1287/moor.2023.0352>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2024, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Minimax Property in Infinite Two-Person Win-Lose Games

Ron Holzman^a

^aDepartment of Mathematics, Technion – Israel Institute of Technology, Technion City, Haifa 3200003, Israel

Contact: holzman@technion.ac.il,  <https://orcid.org/0000-0002-4629-410X> (RH)

Received: November 13, 2023

Revised: April 24, 2024

Accepted: June 28, 2024

Published Online in Articles in Advance:
September 2, 2024

Subject Classifications: Primary: 91A05;
Secondary: 91A44

<https://doi.org/10.1287/moor.2023.0352>

Copyright: © 2024 INFORMS

Abstract. We explore a version of the minimax theorem for two-person win-lose games with infinitely many pure strategies. In the countable case, we give a combinatorial condition on the game which implies the minimax property. In the general case, we prove that a game satisfies the minimax property along with all its subgames if and only if none of its subgames is isomorphic to the “larger number game.” This generalizes a recent theorem of Hanneke, Livni, and Moran. We also propose several applications of our results outside of game theory.

Keywords: minimax theorem • infinite games • larger number game • convergent series • infinite matrices • hypergraph duality • stable theories • online learning

1. Introduction

The minimax theorem is historically considered to be the starting point of game theory, and plays a fundamental role in it. Its standard formulation is for *finite two-person zero-sum games*. In such a game, written as $G = (S, T, \pi)$, players 1 and 2 simultaneously choose their pure strategies $s \in S$ and $t \in T$, respectively, where S and T are finite, nonempty sets. The function $\pi : S \times T \rightarrow \mathbb{R}$ determines the payoffs: $\pi(s, t)$ to player 1 and $-\pi(s, t)$ to player 2. In the *mixed extension* $G^{\text{mix}} = (\Delta(S), \Delta(T), \pi^{\text{mix}})$ of the game G , players 1 and 2 simultaneously choose their mixed strategies $\mathbf{p} \in \Delta(S)$ and $\mathbf{q} \in \Delta(T)$, respectively, where $\Delta(X)$ is the set of probability distributions over X . The payoff function $\pi^{\text{mix}} : \Delta(S) \times \Delta(T) \rightarrow \mathbb{R}$ extends π by taking the expectation: $\pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \mathbb{E}_{(s,t) \sim \mathbf{p} \times \mathbf{q}} \pi(s, t)$.

The Minimax Theorem (von Neumann [12]). *Let $G = (S, T, \pi)$ be a finite two-person zero-sum game. Then in its mixed extension $G^{\text{mix}} = (\Delta(S), \Delta(T), \pi^{\text{mix}})$ we have*

$$\max_{\mathbf{p} \in \Delta(S)} \min_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \min_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}). \quad (1)$$

The left-hand side of Equation (1) is what player 1 can expect to get in a scenario where player 1 first commits to a mixed strategy $\mathbf{p}$ and then player 2, knowing $\mathbf{p}$ (but not its realization), chooses a best response to it. The right-hand side is what player 1 can expect to get in the opposite scenario, where player 2 first commits to a mixed strategy $\mathbf{q}$ and then player 1 best responds to it. In the pure-strategy setting, the two scenarios typically lead to different outcomes: knowing the strategy of one’s opponent before choosing one’s own strategy is a clear advantage. The significance of the minimax theorem is that it shows that this advantage disappears when the players have access to mixed strategies (i.e., flipping coins to choose their pure strategy). In the mixed extension, the two opposite scenarios lead to the same expected payoff v to player 1, namely, the common value of the two sides of Equation (1). This v is called the *value* of G in mixed strategies. It is a payoff level v so that each player can unilaterally guarantee that the expected outcome will be (from that player’s perspective) at least as good as v , regardless of the opponent’s choice of strategy.

Actually, von Neumann proved a more general version of the minimax theorem: whereas the formulation above looks at a bilinear function $\pi^{\text{mix}} : \Delta(S) \times \Delta(T) \rightarrow \mathbb{R}$, he considered a function $f : X \times Y \rightarrow \mathbb{R}$ where X and Y are compact and convex subsets of Euclidean spaces, and f is continuous, concave in the first variable and convex in the second. There is a large body of literature extending the minimax theorem in this form, by allowing X and Y to be subsets of infinite-dimensional spaces and/or relaxing some of the compactness, continuity, or concavity/convexity conditions. See, for example, Fan [5], Sion [11], or the survey by Simons [10]. Another line of literature, exemplified by Yanovskaya [13], has considered the mixed extension of games with infinitely many pure strategies, allowing the use of mixed strategies which are finitely additive but not necessarily σ -additive.

In this paper, we consider games with infinitely many pure strategies, in which every outcome is a win for one player and a loss for the other. Focusing on win-lose games will allow us to find conditions for the minimax theorem to hold, which are based on the combinatorial structure of the payoff function, rather than on continuity and concavity/convexity properties as in most of the literature mentioned above. We will, however, briefly comment on how our results may be extended from win-lose games to real-valued games.

Definition 1.1. A two-person win-lose game is a tuple $G = (S, T, \pi)$, where

- S is the set of pure strategies of player 1,
- T is the set of pure strategies of player 2,
- $\pi : S \times T \rightarrow \{0, 1\}$ is the payoff function.

The sets S, T are nonempty and may be infinite. The game G is *countable* if $|S| = |T| = \aleph_0$.

We will use the term *game* for brevity, to mean a two-person win-lose game. Note that the payoff 1 stands for a win of player 1, whereas the payoff 0 stands for a win of player 2 (a choice of ± 1 for these two outcomes would be more symmetric, but less convenient). As is standard for finite games, a game G can be represented by a 0/1-matrix M_G with rows indexed by $s \in S$, columns indexed by $t \in T$, and $\pi(s, t)$ appearing where row s and column t meet, but such a matrix may have infinitely many rows and columns.

In order to speak about the mixed extension of a game G , we need to specify the class of mixed strategies that are allowed. We will work mostly with countable games. In this case, we will naturally call a *mixed strategy* of player 1 any assignment of probabilities $\mathbf{p} = (p_s)_{s \in S}$ to the pure strategies (equivalently, to the rows), such that $p_s \geq 0$ and $\sum_{s \in S} p_s = 1$, and similarly for player 2. When S (or T) is uncountable, we keep the same definition of a mixed strategy, with the understanding that $p_s > 0$ for at most countably many $s \in S$. (We could allow more general probability distributions over S or T , but then the definition would need to cope with measurability issues. Whenever the minimax property defined below holds with our definition of the class of mixed strategies, it remains valid for any extension of that class, as long as the payoff function is well defined for that extension. Thus, allowing more general mixed strategies would not affect our results.)

Thus, for a set X of pure strategies, we let the set $\Delta(X)$ of mixed strategies consist of the probability distributions over X having at most countable support. Given mixed strategies $\mathbf{p} \in \Delta(S)$ and $\mathbf{q} \in \Delta(T)$, we define $\pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \mathbb{E}_{(s,t) \sim \mathbf{p} \times \mathbf{q}} \pi(s, t)$. Note that this is the probability that player 1 wins under $\mathbf{p} \times \mathbf{q}$.

Definition 1.2. We say that the game $G = (S, T, \pi)$ satisfies the *minimax property* if the following equality holds:

$$\sup_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \inf_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}). \quad (2)$$

Equation (2) is the same as Equation (1), except that we do not insist that the maximum or minimum be attained. For reasons of tradition, we still call this the minimax (rather than the infsup) property. When we can actually assert this equality with max instead of sup or min instead of inf, we will do so. Note that the left-hand side of Equation (2) is always less than or equal to the right-hand side. Equality means that the game G has an ε -value v in mixed strategies, namely, each player can unilaterally guarantee that the expected outcome will be, up to any $\varepsilon > 0$, at least as good as v (from that player's perspective).

Not all games satisfy the minimax property. A standard counterexample is the following.

Example 1.3 (The Larger Number Game (LNG)). In this game, each of the two players chooses a positive integer, and the one who chose the larger number wins. Specifying how ties are broken is immaterial for the analysis, but for concreteness we will assume that they are broken in favor of player 1. Thus, $G = (\mathbb{N}, \mathbb{N}, \pi)$ where

$$\pi(s, t) = \begin{cases} 1 & \text{if } s \geq t, \\ 0 & \text{otherwise.} \end{cases}$$

This game is represented by the following matrix M_G :

	1	2	3	4	5	...
1	1	0	0	0	0	
2	1	1	0	0	0	
3	1	1	1	0	0	...
4	1	1	1	1	0	
5	1	1	1	1	1	
⋮			⋮			⋱

Intuitively, this game violates the minimax property because it is clearly advantageous to know your opponent's strategy before choosing your own. Formally, the left-hand side of Equation (2) equals zero in this example, because given any $\mathbf{p} = (p_s)_{s \in \mathbb{N}} \in \Delta(\mathbb{N})$ and $\varepsilon > 0$, player 2 can hold the winning probability of player 1 below ε by choosing a number $t \in \mathbb{N}$ for which $\sum_{s=1}^{t-1} p_s > 1 - \varepsilon$. A symmetric argument shows that the right-hand side of Equation (2) equals one. The gap between the values of the two sides is as large as it can possibly be.

In view of the above example, our goal in this paper is to identify combinatorial conditions on the game G under which the minimax property does hold. Our first main result gives a sufficient condition for a countable game G to satisfy the minimax property. To formulate the condition, we first extend the game $G = (S, T, \pi)$ by adding to the matrix M_G all possible new rows which are less than some old row (in the natural partial order on $\{0, 1\}^T$). This adds new pure strategies for player 1, each of which is weakly dominated by some strategy in S . Hence, this extension is innocuous in regard to what the players can achieve in the game. Our sufficient condition for G to have the minimax property is that for any infinite sequence $s_1, s_2, \dots$ of player 1's strategies in the extended game, in which s_i weakly dominates s_j for all $i > j$, there is a strategy s that weakly dominates every s_i . Note that the LNG (Example 1.3) violates this condition for the strategy sequence $1, 2, \dots$. In fact, to satisfy the condition, player 1 would need to have ∞ as well in the strategy set. The formal statement of this result, its proof, and some consequences appear in Section 2.

The above condition is sufficient but not necessary for the game G to satisfy the minimax property. One reason for that is its asymmetric treatment of the two players. Another reason, which pertains also to other conceivable conditions that forbid a certain substructure in the matrix M_G , is that a forbidden substructure may appear in one part of the matrix, but G may still satisfy the minimax property due to another part of the matrix. Therefore, it makes more sense to try to characterize those games G which intrinsically satisfy the minimax property, in the sense of the following definition.

Definition 1.4. Let $G = (S, T, \pi)$ be a game. A *subgame* of G is a game of the form $G' = (S', T', \pi')$, where $S' \subseteq S$, $T' \subseteq T$, and π' is the restriction of π to $S' \times T'$. We say that the game G is *totally minimax* if every subgame G' of G satisfies the minimax property.

The “totally minimax” terminology is inspired by cooperative game theory. Recall that a coalitional game is balanced if its core is nonempty, and totally balanced if the core of every subgame is nonempty. Although the contexts (and the notions of subgame) are different, in both cases the total version of the property is interesting because it is more robust. In our case, the robustness is with respect to any deletion of pure strategies.

Clearly, for G to be totally minimax, it must not contain LNG (or any game isomorphic to it) as a subgame. Our second main result shows that the absence of any LNG subgame is also a sufficient condition for G to be totally minimax. (This applies to all, not just countable, games.) In other words, when G is not totally minimax, this can always be attributed to a simple reason, namely, the presence of an LNG subgame. This strengthens a recent result of Hanneke et al. [7], who proved the same under an additional condition on the game G : they required that the VC-dimension of the matrix M_G be finite (the VC-dimension is a fundamental complexity measure in statistical learning theory). The characterization result for totally minimax games is formulated and proved in Section 3.

The final Section 4 describes connections and applications of our results to several topics outside of game theory. These include convergence of number series, row and column densities in infinite binary matrices, fractional matchings and covers in infinite hypergraphs, stability in first-order logic, and online learning theory.

2. A Sufficient Condition for the Minimax Property in Countable Games

Let $G = (S, T, \pi)$ be a game. We assume w.l.o.g. that the rows in the matrix M_G are all distinct, that is, no two strategies of player 1 are equivalent. Thus, each pure strategy $s \in S$ corresponds bijectively to its row $(\pi(s, t))_{t \in T}$ in $\{0, 1\}^T$, or equivalently to the subset B_s of T defined by

$$B_s = \{t \in T : \pi(s, t) = 1\}.$$

In words, for a strategy s of player 1, B_s is the set of player 2's strategies that s beats. Thus, the containment relation $B_s \subset B_{s'}$ corresponds to the game theoretic property that s is weakly dominated by s' . We denote the family of all these sets by $\mathcal{B}^1(G)$, that is,

$$\mathcal{B}^1(G) = \{B_s : s \in S\}.$$

Clearly, the set T and the family of subsets $\mathcal{B}^1(G)$ fully describe the game G . For example, when G is the larger number game, $\mathcal{B}^1(G)$ is the family $\{\{1\}, \{1, 2\}, \dots\}$. We also consider the downward closure $\mathcal{B}_{\text{down}}^1(G)$ of the family

$\mathcal{B}^1(G)$, namely, the family

$$\mathcal{B}_{\text{down}}^1(G) = \{A \subseteq T : \exists s \in S \text{ such that } A \subseteq B_s\}.$$

The family $\mathcal{B}_{\text{down}}^1(G)$ corresponds to an extension of the game G , obtained by adding to player 1's strategy set all possible new strategies that are weakly dominated by some old strategy. This extension has no effect on what the players can achieve in the game, but it facilitates the statement of our sufficient condition for the minimax property. This condition is based on the following set theoretic property.

Definition 2.1. Let $\mathcal{F}$ be a family of sets. We say that $\mathcal{F}$ is *ascending-union closed* if

$$A_i \in \mathcal{F}, i = 1, 2, \dots \text{ and } A_1 \subset A_2 \subset \dots \subset A_i \subset \dots \Rightarrow \bigcup_{i=1}^{\infty} A_i \in \mathcal{F}.$$

In words, the definition requires that the family $\mathcal{F}$ be closed with respect to taking unions of countable ascending chains. Our sufficient condition for the minimax property states that the family $\mathcal{B}_{\text{down}}^1(G)$ should be ascending-union closed. Recalling the correspondence between containment and domination, we can interpret the condition as follows: given any infinite ascending (in the sense of weak domination) sequence of pure strategies of player 1 in the extended game, there should be a strategy which weakly dominates all of them.

The main result of this section is the following.

Theorem 2.2. Let $G = (S, T, \pi)$ be a countable game, such that the family $\mathcal{B}_{\text{down}}^1(G)$ is ascending-union closed. Then G satisfies the minimax property in the following form:

$$\max_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \inf_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}). \quad (3)$$

Proof. By identifying T with the set $\mathbb{N}$ of positive integers, and each subset A with its characteristic function $\mathbb{1}_A : \mathbb{N} \rightarrow \{0, 1\}$, we view $\mathcal{B}_{\text{down}}^1(G)$ as a subset of $\{0, 1\}^{\mathbb{N}}$. Henceforth, we move freely between sets and their characteristic functions. Equipping $\{0, 1\}$ with the discrete topology and $\{0, 1\}^{\mathbb{N}}$ with the corresponding product topology, the latter is a compact topological space by Tychonoff's theorem. Moreover, as $\mathbb{N}$ is countable, the space $\{0, 1\}^{\mathbb{N}}$ is metrizable. Treating $\mathcal{B}_{\text{down}}^1(G)$ as a subset of the compact metric space $\{0, 1\}^{\mathbb{N}}$ is the key idea of our proof, together with the observation made in the first step below.

Step 1. $\mathcal{B}_{\text{down}}^1(G)$ is a closed subset of $\{0, 1\}^{\mathbb{N}}$, hence is itself a compact metric space.

Proof. Let $\mathbf{z}^{(i)} = (z_1^{(i)}, z_2^{(i)}, \dots) \in \mathcal{B}_{\text{down}}^1(G)$, $i = 1, 2, \dots$, and assume that the sequence $\mathbf{z}^{(i)}$ converges to $\mathbf{z} = (z_1, z_2, \dots) \in \{0, 1\}^{\mathbb{N}}$ in the product topology (i.e., pointwise). We have to show that $\mathbf{z} \in \mathcal{B}_{\text{down}}^1(G)$. Let $A = \{j \in \mathbb{N} : z_j = 1\}$. If A is finite, then for large-enough i we have $A \subseteq \{j \in \mathbb{N} : z_j^{(i)} = 1\}$. Because $\mathbf{z}^{(i)}$ is in $\mathcal{B}_{\text{down}}^1(G)$, so is $\mathbf{z}$, and we are done. Thus, we may assume that A is an infinite subset of $\mathbb{N}$, and view it as the ascending union of $A_1 \subset A_2 \subset \dots \subset A_i \subset \dots$, where A_i consists of the first i elements of A . Repeating the argument just made, we see that each A_i is in $\mathcal{B}_{\text{down}}^1(G)$. Because $\mathcal{B}_{\text{down}}^1(G)$ is ascending-union closed, it follows that $A \in \mathcal{B}_{\text{down}}^1(G)$, as required.

Step 2. Let $G_n = (S, T_n, \pi_n)$ be the subgame of G obtained by restricting player 2 to the first n pure strategies. Then for every n , the game G_n has a value v_n , and the sequence $(v_n)_{n=1,2,\dots}$ is nonincreasing and converges to some $v \in [0, 1]$.

Proof. In G_n player 1 has at most 2^n pure strategies, up to equivalence. Thus, G_n is equivalent to a finite game, so by the finite minimax theorem it has a value $v_n \in [0, 1]$. Clearly, the sequence $(v_n)_{n=1,2,\dots}$ is nonincreasing, and therefore it converges to some $v \in [0, 1]$.

In the following two steps, we show that the two sides of Equation (3) are well defined, and both are equal to the v of Step 2.

Step 3. For each $\mathbf{q} \in \Delta(T)$, $\max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q})$ is attained at some pure strategy in S , and we have $\inf_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq v$.

Proof. Fix $\mathbf{q} = (q_j)_{j \in \mathbb{N}} \in \Delta(T)$, and consider the mapping $f : \mathcal{B}_{\text{down}}^1(G) \rightarrow [0, 1]$ which maps $\mathbf{z} = (z_j)_{j \in \mathbb{N}}$ to $\sum_{j \in \mathbb{N}} q_j z_j$. It is easy to check that f is continuous with respect to the product topology. By Step 1, $\mathcal{B}_{\text{down}}^1(G)$ is compact; hence, f attains its maximum at some $\mathbf{z} \in \mathcal{B}_{\text{down}}^1(G)$. From the way that $\mathcal{B}_{\text{down}}^1(G)$ was derived from $\mathcal{B}^1(G)$, the maximum of f over $\mathcal{B}^1(G)$ is the same. This means that there is a pure strategy $s \in S$ which maximizes $\pi^{\text{mix}}(s', \mathbf{q})$ over all $s' \in S$. Hence, this $s \in S$, viewed as a mixed strategy, also maximizes $\pi^{\text{mix}}(\mathbf{p}, \mathbf{q})$ over all $\mathbf{p} \in \Delta(S)$.

To prove the second part of the statement of Step 3, consider for each n an optimal mixed strategy $\mathbf{q}_n \in \Delta(T_n)$ of player 2 in G_n . This $\mathbf{q}_n$ can also be viewed as a mixed strategy of player 2 in G , and as such it satisfies $\max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}_n) = v_n$. Thus, we have

$$\inf_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq \inf_{n \in \mathbb{N}} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}_n) = \inf_{n \in \mathbb{N}} v_n = v.$$

Step 4. There exists $\mathbf{p} \in \Delta(S)$ such that $\inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \geq v$.

Proof. We start this step by recalling the notion of weak convergence of probability measures and a useful fact about it. Let X be a metric space, and let $(P_n)_{n=1,2,\dots}$ be an infinite sequence of Borel probability measures on X . We say that the sequence *converges weakly* to a Borel probability measure P on X if $\int_X f dP_n \rightarrow \int_X f dP$ for every bounded, continuous function $f: X \rightarrow \mathbb{R}$. We will use the following fact: if X is compact, then every such sequence has a weakly convergent subsequence (this follows from Prokhorov's theorem; see, for example, Billingsley [3]).

Now, consider for each n an optimal mixed strategy $\mathbf{p}_n \in \Delta(S)$ of player 1 in G_n . We view $\mathbf{p}_n$ as a Borel probability measure on $X = \mathcal{B}_{\text{down}}^1(G)$, whose support is contained in $\mathcal{B}^1(G)$. By Step 1, X is a compact metric space. By the fact above, we can find a subsequence $(\mathbf{p}_{n_i})_{i=1,2,\dots}$ that converges weakly to a Borel probability measure P on X . Fixing a pure strategy $t \in T$, we consider the characteristic function $\mathbb{1}_{[z_t=1]}: X \rightarrow \{0,1\}$ of those $\mathbf{z} = (z_1, z_2, \dots) \in \mathcal{B}_{\text{down}}^1(G)$ having $z_t = 1$. This function is clearly continuous with respect to the product topology. Hence, by weak convergence, we have

$$\int_X \mathbb{1}_{[z_t=1]} d\mathbf{p}_{n_i} \rightarrow \int_X \mathbb{1}_{[z_t=1]} dP. \quad (4)$$

These integrals compute the probability, under $\mathbf{p}_{n_i}$ and P , respectively, of player 1 winning against the pure strategy t of player 2. For i large enough so that $t \in T_{n_i}$, we know that this probability, under $\mathbf{p}_{n_i}$, is at least v_{n_i} . As $v_{n_i} \rightarrow v$, it follows from (4) that the probability under P of player 1 winning against t is at least v . This being true for every pure strategy $t \in T$, it is also true for every mixed strategy $\mathbf{q} \in \Delta(T)$.

The above almost establishes the statement of Step 4, except that P is a Borel probability measure on $\mathcal{B}_{\text{down}}^1(G)$ rather than on $\mathcal{B}^1(G)$. To correct this, we replace P by a probability measure $\mathbf{p} = (p_{B_s})_{B_s \in \mathcal{B}^1(G)}$ on $\mathcal{B}^1(G)$ as follows. We fix an arbitrary enumeration $B_{s_1}, B_{s_2}, \dots$ of the sets in $\mathcal{B}^1(G)$. We define a mapping $f: \mathcal{B}_{\text{down}}^1(G) \rightarrow \mathcal{B}^1(G)$ by $f(A) = B_{s_i}$ where i is the least index so that $A \subseteq B_{s_i}$. Note that f is Borel measurable, because

$$f^{-1}(B_{s_i}) = \{A \subseteq \mathbb{N} : A \subseteq B_{s_i} \text{ and } A \not\subseteq B_{s_j} \text{ for all } j < i\},$$

and for each $s \in S$, the set $\{A \subseteq \mathbb{N} : A \subseteq B_s\}$ is closed in the product topology. This allows us to define the desired $\mathbf{p}$ by $p_{B_s} = P[f^{-1}(B_s)]$ for each $B_s \in \mathcal{B}^1(G)$. This $\mathbf{p}$ can be viewed as a mixed strategy in $\Delta(S)$. As $A \subseteq B_s$ for every $A \in f^{-1}(B_s)$, player 1's winning probability cannot decrease when replacing P by $\mathbf{p}$. Thus, $\mathbf{p}$ satisfies the statement of Step 4.

Steps 3 and 4 together yield the equality (3). $\square$

A few remarks about Theorem 2.2 are in order.

Remark 2.3. The condition that the family $\mathcal{B}_{\text{down}}^1(G)$ be ascending-union closed is sufficient but not necessary for the minimax property to hold. As an example, we can extend the LNG (Example 1.3) by allowing player 2 to choose ∞ , which beats all of player 1's strategies. After this extension, $\mathcal{B}_{\text{down}}^1(G)$ is still not ascending-union closed, but the value of the extended game is trivially zero.

Remark 2.4. In any countable game $G = (S, T, \pi)$ where all sets $B_s \in \mathcal{B}^1(G)$ are finite, the left-hand side of Equality (3) equals zero. Indeed, let $\mathbf{p} = (p_s)_{s \in S} \in \Delta(S)$ and let $\varepsilon > 0$. Then we can find a finite subset $S' \subset S$ such that $\sum_{s \in S'} p_s > 1 - \varepsilon$. By our assumptions, the set $\cup_{s \in S'} B_s$ is finite and T is infinite; hence, we can find a pure strategy $t \in T \setminus (\cup_{s \in S'} B_s)$. For this t we have $\pi^{\text{mix}}(\mathbf{p}, t) < \varepsilon$, implying that $\inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) < \varepsilon$. As $\mathbf{p} \in \Delta(S)$ and $\varepsilon > 0$ were arbitrary, the left-hand side of Equality (3) equals zero.

Remark 2.5. Requiring that the family $\mathcal{B}^1(G)$, not $\mathcal{B}_{\text{down}}^1(G)$, be ascending-union closed is not sufficient for the minimax property to hold. As an example, let T be the disjoint union of two copies C_1 and C_2 of $\mathbb{N}$. The sets B_s in $\mathcal{B}^1(G)$ are of one of two forms: either B_s consists of the first n elements of C_1 and the n -th element of C_2 for some $n \in \mathbb{N}$, or vice versa. There are no containments between the sets in $\mathcal{B}^1(G)$; hence, $\mathcal{B}^1(G)$ is vacuously ascending-union closed. Yet the minimax property does not hold. On the one hand, by Remark 2.4, no mixed strategy $\mathbf{p}$ of player 1 guarantees a positive winning probability. On the other hand, given any mixed strategy $\mathbf{q}$ of player 2

and $\varepsilon > 0$, player 1 has a pure strategy s which wins with probability at least $(1/2) - \varepsilon$. Indeed, we may assume w.l.o.g. that $\mathbf{q}$ assigns a probability of at least $1/2$ to C_1 , hence at least $(1/2) - \varepsilon$ to the first n elements of C_1 , for some $n \in \mathbb{N}$. Choosing s such that B_s contains the first n elements of C_1 works for player 1.

Remark 2.6. The infimum over $\mathbf{q} \in \Delta(T)$, on either side of Equation (3), cannot be replaced by a minimum. As an example, consider the following game G , an infinite version of matching pennies. Each player chooses a positive integer, and player 1 wins if and only if they chose the same number. This game is represented by a diagonal matrix M_G with 1's on the diagonal. The family $\mathcal{B}_{\text{down}}^1(G)$ consists of the singletons and the empty set, and is vacuously ascending-union closed. The ε -value of G is clearly zero, but player 2 has no exactly optimal mixed strategy, nor an exactly best response to a full-support mixed strategy of player 1.

Remark 2.7. The proof of Step 3 above shows that player 2 can achieve arbitrarily good approximations of v by using the mixed strategies $\mathbf{q}_n$, which have finite support and w.l.o.g. assign rational probabilities to the pure strategies. The same is true of player 1: arbitrarily good approximations of v can be achieved by starting from the mixed strategy $\mathbf{p}$ obtained in Step 4, and making small-enough perturbations which render the support finite and the assigned probabilities rational. However, *exact* optimality for player 1 cannot always be achieved under these restrictions. By translating example 4.1 of Aharoni and Holzman [1] to the present context, one obtains a game G satisfying the condition of Theorem 2.2, in which player 1 has a unique optimal mixed strategy $\mathbf{p}$, that has full support and assigns irrational probabilities to some pure strategies. (See Subsection 4.3 below on how to translate the terminology of Aharoni and Holzman [1] to ours.)

Remark 2.8. Theorem 2.2 may be extended beyond the class of win-lose games, at the cost of losing the combinatorial/set theoretic nature of its condition. Instead of requiring the range of the payoff function π to be $\{0, 1\}$, we may allow it to be a bounded subset of $\mathbb{R}$, say, the interval $[0, 1]$. The discrete topology on $\{0, 1\}$ is replaced by the standard topology on $[0, 1]$. Instead of the combinatorial condition on $\mathcal{B}_{\text{down}}^1(G)$ in Theorem 2.2, which is used in the proof to show that it is a closed subset of $\{0, 1\}^{\mathbb{N}}$, we may require directly that the analog of $\mathcal{B}_{\text{down}}^1(G)$ be a closed subset of $[0, 1]^{\mathbb{N}}$. The rest of the proof works as above, with minor adaptations. We omit the details.

In the special case when all sets in $\mathcal{B}^1(G)$ are finite (equivalently, every row of the matrix M_G has finitely many 1's), the condition of Theorem 2.2 is directly related to the absence of LNG subgames in G (recall Example 1.3 and Definition 1.4). We show this after introducing the following definition. We say that two games $G_1 = (S_1, T_1, \pi_1)$ and $G_2 = (S_2, T_2, \pi_2)$ are *isomorphic* if there exist bijections $\sigma: S_1 \rightarrow S_2$ and $\tau: T_1 \rightarrow T_2$ such that $\pi_2(\sigma(s), \tau(t)) = \pi_1(s, t)$ for every $s \in S_1, t \in T_1$.

Definition 2.9. Let $G = (S, T, \pi)$ be a game. We say that G is *LNG-free* if no subgame of G is isomorphic to the LNG.

A subgame of G which is isomorphic to LNG appears in the matrix M_G as a submatrix which, possibly after reordering its rows and columns, has the full lower-triangular form shown in Example 1.3. It may be checked that the notion of LNG-freeness is invariant to changes in the tie-breaking rule used in LNG.

Proposition 2.10. Let $G = (S, T, \pi)$ be a game such that all sets in the family $\mathcal{B}^1(G)$ are finite. Then the following three conditions are equivalent:

1. $\mathcal{B}_{\text{down}}^1(G)$ is ascending-union closed.
2. $\mathcal{B}_{\text{down}}^1(G)$ contains no countable ascending chain $A_1 \subset A_2 \subset \dots \subset A_i \subset \dots$.
3. G is LNG-free.

Proof. The implication $2 \Rightarrow 1$ is trivially true. Thus, it suffices to prove the implications $1 \Rightarrow 2$, $2 \Rightarrow 3$, and $3 \Rightarrow 2$. We prove those in turn.

$1 \Rightarrow 2$. We know that all sets in $\mathcal{B}^1(G)$, and hence all sets in $\mathcal{B}_{\text{down}}^1(G)$, are finite. If $\mathcal{B}_{\text{down}}^1(G)$ does contain a countable ascending chain $A_1 \subset A_2 \subset \dots \subset A_i \subset \dots$, then $\bigcup_{i=1}^{\infty} A_i$ is infinite and therefore not in $\mathcal{B}_{\text{down}}^1(G)$, so $\mathcal{B}_{\text{down}}^1(G)$ is not ascending-union closed.

$2 \Rightarrow 3$. Suppose, for the sake of contradiction, that $G' = (S', T', \pi')$ is a subgame of G which is isomorphic to LNG. Then the sets $\{t \in T' : \pi(s, t) = 1\}$, $s \in S'$, are in $\mathcal{B}_{\text{down}}^1(G)$ and form a countable ascending chain, contradicting 2.

$3 \Rightarrow 2$. Suppose, for the sake of contradiction, that $A_1 \subset A_2 \subset \dots \subset A_i \subset \dots$ is a countable ascending chain of sets in $\mathcal{B}_{\text{down}}^1(G)$. We will show that G has a subgame $G' = (S', T', \pi')$ which is isomorphic to LNG, in contradiction to 3. This requires us to construct an infinite sequence of pure strategies, alternating between the two players, of the form $t_1, s_1, t_2, s_2, \dots, t_k, s_k, \dots$, in which each strategy beats all the preceding strategies of the opponent. Then, taking $S' = \{s_1, s_2, \dots\}$ and $T' = \{t_1, t_2, \dots\}$, we will be done, because the corresponding subgame will be isomorphic to LNG under the bijections $\sigma(s_i) = i$ and $\tau(t_j) = j$. Our construction of the sequence is inductive, starting with an

arbitrary $t_1 \in \cup_{i=1}^{\infty} A_i$ and maintaining the condition that all t_k are in $\cup_{i=1}^{\infty} A_i$. Now suppose that we have constructed the sequence up to a certain point, and we need to find the next strategy to be added. If that strategy is t_k , we choose it in $(\cup_{i=1}^{\infty} A_i) \setminus (\cup_{j < k} B_{s_j})$. Note that $\cup_{j < k} B_{s_j}$ is finite by assumption whereas $\cup_{i=1}^{\infty} A_i$ is infinite, so this choice is feasible. By the definition of the sets B_{s_j} , our t_k beats all s_j , $j < k$, as required. If the strategy to be added is s_k , we look at the set $\{t_1, \dots, t_k\}$. Each element of this set lies in some A_i , and as the A_i form a chain, there exists an i so that $\{t_1, \dots, t_k\} \subseteq A_i$. Because $A_i \in \mathcal{B}_{\text{down}}^1(G)$, we can choose s_k so that $A_i \subseteq B_{s_k}$. This s_k beats all t_j , $j \leq k$, as required. This completes the inductive construction. $\square$

The above proposition leads to the following special case of Theorem 2.2.

Corollary 2.11. *Let $G = (S, T, \pi)$ be a countable game such that all sets in the family $\mathcal{B}^1(G)$ are finite. Assume that G satisfies one (hence all) of the equivalent conditions 1–3 in Proposition 2.10. Then G satisfies the minimax property in the following form:*

$$\max_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = 0 = \inf_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}).$$

Proof. This follows immediately from Theorem 2.2 and Remark 2.4. $\square$

The condition in Theorem 2.2 focuses on player 1's strategies. We can interchange the roles of the players and get a dual result which focuses on player 2's strategies. We briefly go through the dual notations.

Let $G = (S, T, \pi)$ be a game. We assume w.l.o.g. that the columns in the matrix M_G are all distinct. Thus, each pure strategy $t \in T$ corresponds bijectively to its column $(\pi(s, t))_{s \in S}$ in $\{0, 1\}^S$, or equivalently to the subset B_t of S defined by

$$B_t = \{s \in S : \pi(s, t) = 0\}.$$

We denote the family of all these sets by $\mathcal{B}^2(G)$, that is,

$$\mathcal{B}^2(G) = \{B_t : t \in T\}.$$

For example, when G is the larger number game, $\mathcal{B}^2(G)$ is the family $\{\emptyset, \{1\}, \{1, 2\}, \dots\}$. We also consider the downward closure $\mathcal{B}_{\text{down}}^2(G)$ of the family $\mathcal{B}^2(G)$, namely, the family

$$\mathcal{B}_{\text{down}}^2(G) = \{A \subseteq S : \exists t \in T \text{ such that } A \subseteq B_t\}.$$

The dual result to Theorem 2.2 is the following.

Theorem 2.12. *Let $G = (S, T, \pi)$ be a countable game, such that the family $\mathcal{B}_{\text{down}}^2(G)$ is ascending-union closed. Then G satisfies the minimax property in the following form:*

$$\sup_{\mathbf{p} \in \Delta(S)} \min_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \min_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}).$$

Dualizing Corollary 2.11, we get the following.

Corollary 2.13. *Let $G = (S, T, \pi)$ be a countable game such that all sets in the family $\mathcal{B}^2(G)$ are finite. Assume that G is LNG-free. Then G satisfies the minimax property in the following form:*

$$\sup_{\mathbf{p} \in \Delta(S)} \min_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = 1 = \min_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}).$$

3. A Characterization of Totally Minimax Games

As explained in the introduction, we should not expect a structural characterization of the games that satisfy the minimax property. A more realistic goal, that we achieve in this section, is a characterization of those games which are totally minimax (Definition 1.4). Our result shows that the presence of a copy of LNG (Example 1.3) is the only impediment to being totally minimax. This improves an earlier result of Hanneke et al. [7], who proved the same under the additional condition that the VC-dimension of the game be finite. This condition requires the existence of a finite bound on the size of subsets $T' \subseteq T$ which are shattered, in the sense that for every subset $A \subseteq T'$ there exists $s \in S$ for which $\{t \in T' : \pi(s, t) = 1\} = A$. The following example shows that this condition of Hanneke et al. [7] is indeed restrictive.

Example 3.1 (An LNG-Free Game of Infinite VC-Dimension). The $\mathbb{N}$ by $\mathbb{N}$ binary matrix M_G of our game G looks as follows:

	1	2	3	4	5	6	...
1	1	0	1	0	0	1	
2	0	1	0	0	1	0	
3	0	1	1	0	1	1	
4	$\vdots$	0	0	1	0	0	
5		0	0	1	0	1	
6		$\vdots$	$\vdots$	1	1	0	
7				1	1	1	
8				0	0	0	
9				0	0	0	
$\vdots$				$\vdots$	$\vdots$	$\vdots$	

Formally, the index set of the columns is partitioned as $\mathbb{N} = \bigcup_{j=1}^{\infty} T_j$, where $|T_j| = j$ for each j (in the illustration above $T_1 = \{1\}$, $T_2 = \{2, 3\}$, $T_3 = \{4, 5, 6\}$, etc.). In the columns corresponding to T_j , there are $2^j - 1$ rows displaying all nonzero binary vectors of length j , and the remaining infinitely many rows are identically zero in those columns. Thus, each set T_j is shattered, and as the sizes of the sets T_j are finite but arbitrarily large, the VC-dimension of G is infinite. Yet G is LNG-free, because each column contains finitely many 1 entries.

Our characterization result is the following.

Theorem 3.2. Let $G = (S, T, \pi)$ be a game. Then G is totally minimax if and only if G is LNG-free.

Proof. As LNG violates the minimax property, the “only if” direction is trivial. In the opposite direction, we prove below that if G is LNG-free, then G satisfies the minimax property. This will suffice, because all subgames of an LNG-free game are themselves LNG-free.

From now on, we assume that G is LNG-free. For most of the proof, we also assume that G is countable (in the last step we will show how to reduce the result for arbitrary games to the countable case). Our basic approach is the same as in the proof of Theorem 2.2. Recall that $\mathcal{B}^1(G)$ is the family of sets B_s , $s \in S$, which correspond to the rows of the matrix M_G . We treat $\mathcal{B}^1(G)$ as a subset of the compact metric space $\{0, 1\}^{\mathbb{N}}$. In the proof of Theorem 2.2 we extended $\mathcal{B}^1(G)$ to $\mathcal{B}_{\text{down}}^1(G)$, which was shown, under the condition of that theorem, to be a closed subset of $\{0, 1\}^{\mathbb{N}}$. Here, instead, we extend $\mathcal{B}^1(G)$ to its closure in $\{0, 1\}^{\mathbb{N}}$, which we denote $\text{cl}(\mathcal{B}^1(G))$. Thus, $\text{cl}(\mathcal{B}^1(G))$ is itself a compact metric space.

We want to show that G satisfies the minimax property. By Corollaries 2.11 and 2.13, we already know that this follows from LNG-freeness in the special case when $\mathcal{B}^1(G)$ or $\mathcal{B}^2(G)$ consists of finite sets only. We use this fact in the first step of our proof, where we show that for any $\varepsilon > 0$, every pure strategy of player 1 in the extended game corresponding to $\text{cl}(\mathcal{B}^1(G))$ can be replaced by a mixed strategy in the original game G , while losing at most ε against any pure strategy of player 2.

Step 1. Let $\varepsilon > 0$, and let $\mathbf{z} = (z_1, z_2, \dots) \in \text{cl}(\mathcal{B}^1(G))$. Then there exists $\mathbf{p} \in \Delta(S)$ such that for all $t \in T$:

$$z_t = 1 \Rightarrow \pi^{\text{mix}}(\mathbf{p}, t) \geq 1 - \varepsilon. \quad (5)$$

Moreover, the mixed strategy $\mathbf{p} = (p_s)_{s \in S}$ may be chosen so that its support is finite and all probabilities p_s , $s \in S$, are rational.

Proof. As $\mathbf{z} \in \text{cl}(\mathcal{B}^1(G))$, we can find a sequence $\mathbf{z}^{(i)} = (z_1^{(i)}, z_2^{(i)}, \dots) \in \mathcal{B}^1(G)$, $i = 1, 2, \dots$, which converges to $\mathbf{z}$ in the product topology (i.e., pointwise). We consider the subgame $G' = (S', T', \pi')$ of G , where the strategies $s \in S'$ are those corresponding to $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \dots$, and the strategies $t \in T'$ are those for which $z_t = 1$. If there exists $s \in S'$ such that $\pi(s, t) = 1$ for all $t \in T'$, then we can take $\mathbf{p}$ to be the pure strategy s . Note that such $s \in S'$ must exist if either S' or T' is finite. Indeed, if S' is finite then some $s \in S'$ corresponds to $\mathbf{z}^{(i)}$ for infinitely many superscripts (i); these $\mathbf{z}^{(i)}$ form a constant infinite subsequence, whose convergence to $\mathbf{z}$ means that they are all equal to $\mathbf{z}$. If T' is finite, then for sufficiently large i , we have $z_t^{(i)} = z_t = 1$ for all $t \in T'$. Thus, we may assume that both S' and T' are infinite, so G' is a countable game. Because $(\mathbf{z}^{(i)})_{i=1,2,\dots}$ converges pointwise to $\mathbf{z}$, for every $t \in T'$, all but finitely many i satisfy $z_t^{(i)} = z_t = 1$. By definition, each set $B_t \in \mathcal{B}^2(G')$ is of the form $B_t = \{s \in S' : \pi(s, t) = 0\}$ for some $t \in T'$, and we have just observed that these sets are finite. Thus, G' satisfies the assumptions of Corollary 2.13 (it is

LNG-free as a subgame of G). Therefore, its ε -value is one, so we can find a mixed strategy $\mathbf{p}$ of player 1 (in G' , hence also in G) so that $\pi^{\text{mix}}(\mathbf{p}, t) \geq 1 - \varepsilon$ for all $t \in T'$. Such $\mathbf{p}$ satisfies (5), and moreover can be chosen to have finite support and assign rational probabilities (see Remark 2.7).

Step 2. Let $G_n = (S, T_n, \pi_n)$ be the subgame of G obtained by restricting player 2 to the first n pure strategies. Then for every n , the game G_n has a value v_n , and the sequence $(v_n)_{n=1,2,\dots}$ is nonincreasing and converges to some $v \in [0, 1]$.

Proof. This is proved just like the corresponding step in the proof of Theorem 2.2.

We go on to show that the v of Step 2 is the ε -value of G .

Step 3. We have $\inf_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq v$.

Proof. This is again similar to the second part of the corresponding step in the proof of Theorem 2.2. Except for having a supremum, not a maximum over $\mathbf{p} \in \Delta(S)$, the proof is the same.

Step 4. We have $\sup_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \geq v$.

Proof. Repeating the arguments in the corresponding step in the proof of Theorem 2.2, with $X = \text{cl}(\mathcal{B}^1(G))$ instead of $\mathcal{B}_{\text{down}}^1(G)$, we find a Borel probability measure P on $\text{cl}(\mathcal{B}^1(G))$, under which player 1 wins with probability at least v against every strategy of player 2. We show below, using Step 1, how to replace P by a mixed strategy $\mathbf{p} \in \Delta(S)$, while incurring an arbitrarily small loss in this guarantee.

Let $\varepsilon > 0$ be given. We denote by $\Delta_{\mathbb{Q}}^0(S)$ the set of all mixed strategies in $\Delta(S)$ that have a finite support and assign only rational probabilities (this set is countable). We fix an arbitrary enumeration $\mathbf{p}^{(1)}, \mathbf{p}^{(2)}, \dots$ of $\Delta_{\mathbb{Q}}^0(S)$. We define a mapping $f : \text{cl}(\mathcal{B}^1(G)) \rightarrow \Delta_{\mathbb{Q}}^0(S)$ by $f(\mathbf{z}) = \mathbf{p}^{(i)}$ where i is the least index so that $\mathbf{p}^{(i)}$ satisfies Condition (5) with respect to ε and $\mathbf{z} = (z_1, z_2, \dots)$. By Step 1, this is well defined. The mapping f is Borel measurable, because $f^{-1}(\mathbf{p}^{(i)})$ consists of those $\mathbf{z} = (z_1, z_2, \dots) \in \text{cl}(\mathcal{B}^1(G))$ that satisfy: $z_t = 0$ for all $t \in T$ such that $\pi^{\text{mix}}(\mathbf{p}^{(i)}, t) < 1 - \varepsilon$ (this defines a closed set in the product topology) and the negation of this for every $j < i$ (this defines an open set).

Consider the probability distribution $\mathbf{p}$ over S defined by the following two-step random process. First, $\mathbf{z} = (z_1, z_2, \dots)$ is chosen in $\text{cl}(\mathcal{B}^1(G))$ according to the Borel probability measure P obtained above. Next, a pure strategy s is chosen in S according to the mixed strategy $f(\mathbf{z})$. We fix $t \in T$ and lower bound the probability that $\pi(s, t) = 1$, when s is chosen according to $\mathbf{p}$ as explained. The known property of P ensures that in the first step $z_t = 1$ with probability at least v . By (5), the conditional probability that $\pi(s, t) = 1$ given that $z_t = 1$ is at least $1 - \varepsilon$. Thus, the probability that $\pi(s, t) = 1$ is at least $v(1 - \varepsilon)$. This is true for every pure strategy $t \in T$, hence, also for every mixed strategy $\mathbf{q} \in \Delta(T)$. So, our $\mathbf{p} \in \Delta(S)$ satisfies $\inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \geq v(1 - \varepsilon)$. As $\varepsilon > 0$ was arbitrary, the claim of Step 4 is proved.

Steps 3 and 4 together show that G satisfies the minimax property, completing the proof in the countable case. In the next step, we extend the result from the countable case to the general case.

Step 5. Let $G = (S, T, \pi)$ be a game that violates the minimax property. Then G has a countable subgame $G' = (S', T', \pi')$ that violates the minimax property.

Proof. Here we use a construction from Hanneke et al. [7]. We denote the two sides of Equation (2) by $\underline{v}$ and $\bar{v}$, respectively, and assume that

$$\sup_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \underline{v} < \bar{v} = \inf_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}).$$

We construct an infinite sequence of mixed strategies with finite support, alternating between the two players, of the form $\mathbf{p}^{(1)}, \mathbf{q}^{(1)}, \mathbf{p}^{(2)}, \mathbf{q}^{(2)}, \dots, \mathbf{p}^{(k)}, \mathbf{q}^{(k)}, \dots$, in which each strategy gives the player using it an expected payoff at least as good as the better among $\underline{v}$ and $\bar{v}$, against any pure strategy in the union of the supports of the preceding strategies of the opponent. Namely, for every k we require:

$$\pi^{\text{mix}}(\mathbf{p}^{(k)}, t) \geq \bar{v} \text{ for all } t \in \bigcup_{j < k} \text{supp}(\mathbf{q}^{(j)}), \quad (6)$$

$$\pi^{\text{mix}}(s, \mathbf{q}^{(k)}) \leq \underline{v} \text{ for all } s \in \bigcup_{j \leq k} \text{supp}(\mathbf{p}^{(j)}). \quad (7)$$

The sequence is constructed inductively, starting with an arbitrary $\mathbf{p}^{(1)} \in \Delta(S)$ with finite support. Suppose that we have constructed the sequence up to a certain point, and we need to find the next strategy to be added. If that strategy is $\mathbf{p}^{(k)}$, we consider the subgame $G^{(k)} = (S, T^{(k)}, \pi^{(k)})$ obtained by restricting player 2 to the pure strategies $t \in \bigcup_{j < k} \text{supp}(\mathbf{q}^{(j)})$. This subgame is equivalent to a finite game, because the $\mathbf{q}^{(j)}$ have finite support. Hence, $G^{(k)}$ has a value $v^{(k)}$, and clearly $v^{(k)} \geq \bar{v}$. Taking $\mathbf{p}^{(k)}$ to be an optimal mixed strategy (with finite support) of player 1

in $G^{(k)}$ satisfies Condition (6). If the strategy to be added is $\mathbf{q}^{(k)}$, a similar argument allows us to choose it so as to satisfy Condition (7). This completes the inductive construction.

Now, consider the subgame $G' = (S', T', \pi')$ where $S' = \bigcup_{j=1}^{\infty} \text{supp}(\mathbf{p}^{(j)})$ and $T' = \bigcup_{j=1}^{\infty} \text{supp}(\mathbf{q}^{(j)})$. Clearly, S' and T' are at most countable. We claim that

$$\sup_{\mathbf{p} \in \Delta(S')} \inf_{\mathbf{q} \in \Delta(T')} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq \underline{v} < \bar{v} \leq \inf_{\mathbf{q} \in \Delta(T')} \sup_{\mathbf{p} \in \Delta(S')} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}). \quad (8)$$

This will show that G' also violates the minimax property, thereby completing the proof. We prove the first inequality in (8); the proof of the last inequality is similar. It suffices to show that for any fixed $\mathbf{p} \in \Delta(S')$, we have $\inf_{k \in \mathbb{N}} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}^{(k)}) \leq \underline{v}$. This follows from (7) and the fact that when $s \in S'$ is chosen according to $\mathbf{p}$, the probability that $s \in \bigcup_{j \leq k} \text{supp}(\mathbf{p}^{(j)})$ tends to one as $k \rightarrow \infty$. $\square$

Theorem 3.2 says that LNG-freeness is a necessary and sufficient condition for being totally minimax. However, checking whether a given game satisfies this condition may not be easy. The following corollary lists some simple-to-check counting conditions, each of which precludes an LNG subgame and therefore suffices for being totally minimax.

Corollary 3.3. *Let $G = (S, T, \pi)$ be a game. Assume that the representing matrix M_G satisfies one (or more) of the following conditions:*

1. *The number of 0 entries in each row (except maybe finitely many) is finite.*
2. *The number of 1 entries in each row (except maybe finitely many) is finitely bounded.*
3. *The number of 0 entries in each column (except maybe finitely many) is finitely bounded.*
4. *The number of 1 entries in each column (except maybe finitely many) is finite.*

Then G is totally minimax.

Remark 3.4. Steve Hanneke [6] has extended Theorem 3.2 to characterize those zero-sum games with a bounded payoff function which are totally minimax. Instead of our single forbidden binary submatrix that represents LNG, he forbids all real-valued submatrices of the same shape in which the supremum of the entries above the diagonal is strictly less than the infimum of the entries below the diagonal. His proof extends the techniques used here to the real case, and also applies Theorem 3.2 in the context of a nontrivial discretization argument. A similar extension to the real case of the above-mentioned result of Hanneke et al. [7] was obtained by Daskalakis and Golowich [4].

4. Connections and Applications

4.1. Enforcing the Convergence of Nonnegative Number Series

Here we address the question: Which families of finite sets of positive integers enforce the convergence of nonnegative number series, in the sense of the following definition?

Definition 4.1. Let $\mathcal{F}$ be a family of finite subsets of $\mathbb{N}$. We say that $\mathcal{F}$ is *convergence enforcing* if the following implication holds for every infinite sequence $(a_i)_{i=1,2,\dots}$ of nonnegative real numbers:

$$\sum_{i \in A} a_i \leq 1 \text{ for all } A \in \mathcal{F} \Rightarrow \sum_{i \in \mathbb{N}} a_i < \infty. \quad (9)$$

The standard definition of convergence of a series yields the simplest example of a convergence enforcing family: when $\mathcal{F}$ is the family of all finite initial segments of $\mathbb{N}$, the implication (9) holds with $\sum_{i \in \mathbb{N}} a_i \leq 1$. Another example is the family of all finite initial segments of the even positive integers or of the odd ones, which enforces $\sum_{i \in \mathbb{N}} a_i \leq 2$. There are many other, more sophisticated examples. Here is a preliminary observation.

Proposition 4.2. *Let $\mathcal{F}$ be a family of finite subsets of $\mathbb{N}$ such that $\bigcup \mathcal{F} = \mathbb{N}$. Assume that $\mathcal{F}$ is convergence enforcing. Then there exists a finite constant $C = C(\mathcal{F}) < \infty$ so that the following implication holds for every infinite sequence $(a_i)_{i=1,2,\dots}$ of nonnegative real numbers:*

$$\sum_{i \in A} a_i \leq 1 \text{ for all } A \in \mathcal{F} \Rightarrow \sum_{i \in \mathbb{N}} a_i \leq C. \quad (10)$$

Proof. Suppose, for the sake of contradiction, that no such constant C exists. Applying this to all C of the form 2^n , $n \in \mathbb{N}$, we can find for every n a sequence $(a_i^{(n)})_{i=1,2,\dots}$ as above, such that $\sum_{i \in A} a_i^{(n)} \leq 1$ for all $A \in \mathcal{F}$, yet $\sum_{i \in \mathbb{N}} a_i^{(n)} > 2^n$. Note that $a_i^{(n)} \leq 1$ for all n and i , because every i belongs to some $A \in \mathcal{F}$. Hence $a_i = \sum_{n \in \mathbb{N}} a_i^{(n)} / 2^n$ is a

finite nonnegative number for each i . However, the sequence $(a_i)_{i=1,2,\dots}$ violates (9): we have

$$\sum_{i \in A} a_i = \sum_{i \in A} \sum_{n \in \mathbb{N}} \frac{a_i^{(n)}}{2^n} = \sum_{n \in \mathbb{N}} \frac{1}{2^n} \sum_{i \in A} a_i^{(n)} \leq \sum_{n \in \mathbb{N}} \frac{1}{2^n} = 1$$

for all $A \in \mathcal{F}$, yet

$$\sum_{i \in \mathbb{N}} a_i = \sum_{i \in \mathbb{N}} \sum_{n \in \mathbb{N}} \frac{a_i^{(n)}}{2^n} = \sum_{n \in \mathbb{N}} \frac{1}{2^n} \sum_{i \in \mathbb{N}} a_i^{(n)} \geq \sum_{n \in \mathbb{N}} 1 = \infty.$$

This contradicts our assumption that $\mathcal{F}$ is convergence enforcing. $\square$

Another observation is that $\mathcal{F}$ is convergence enforcing if and only if its downward closure is. Therefore, when studying this property, we may assume that the family $\mathcal{F}$ is downward closed, that is, $A \subseteq B$ and $B \in \mathcal{F} \Rightarrow A \in \mathcal{F}$.

Although downward-closed convergence enforcing families may look very different from one another, we show here that they all have one thing in common: they must contain an infinite ascending chain. This is an application of Corollary 2.11.

Theorem 4.3. *Let $\mathcal{F}$ be a family of finite subsets of $\mathbb{N}$. Assume that $\mathcal{F}$ is downward closed and convergence enforcing. Then $\mathcal{F}$ contains an infinite ascending chain $A_1 \subset A_2 \subset \dots \subset A_i \subset \dots$.*

Proof. Suppose, for the sake of contradiction, that $\mathcal{F}$ contains no infinite ascending chain.

We may assume that $\cup \mathcal{F} = \mathbb{N}$. Indeed, the set $\mathbb{N} \setminus (\cup \mathcal{F})$ must be finite, otherwise convergence enforcing clearly fails. If $\mathbb{N} \setminus (\cup \mathcal{F})$ is nonempty, we may consider $\cup \mathcal{F}$ as our copy of $\mathbb{N}$.

Assuming that $\cup \mathcal{F} = \mathbb{N}$, we can find a constant $C = C(\mathcal{F}) < \infty$ so that (10) holds for every infinite sequence $(a_i)_{i=1,2,\dots}$ of nonnegative real numbers. Consider the game $G = (S, T, \pi)$ where $T = \mathbb{N}$ and $B^1(G) = \mathcal{F}$. As $\mathcal{F}$ is downward closed, we have $B_{\text{down}}^1(G) = \mathcal{F}$ as well. By our assumption, this family contains no infinite ascending chain. Thus, condition 2 in Proposition 2.10 holds. So Corollary 2.11 applies, showing that $\inf_{\mathbf{q} \in \Delta(T)} \max_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = 0$. Let $0 < \varepsilon < (1/C)$, and let $\mathbf{q} = (q_i)_{i \in \mathbb{N}}$ be a mixed strategy of player 2 such that $\pi^{\text{mix}}(s, \mathbf{q}) \leq \varepsilon$ for every pure strategy s of player 1. Now consider the sequence $(a_i)_{i=1,2,\dots}$ where $a_i = q_i/\varepsilon$. We get a contradiction by showing that this sequence violates (10). To check the premise of (10), let $A \in \mathcal{F}$ and let s be the corresponding pure strategy of player 1. Then $\sum_{i \in A} a_i = (1/\varepsilon) \sum_{i \in A} q_i = (1/\varepsilon) \pi^{\text{mix}}(s, \mathbf{q}) \leq 1$. On the other hand, $\sum_{i \in \mathbb{N}} a_i = (1/\varepsilon) > C$. $\square$

The contrapositive of the theorem says that if a downward-closed family $\mathcal{F}$ of finite subsets of $\mathbb{N}$ contains no infinite ascending chain, then there is a divergent series $\sum_{i \in \mathbb{N}} a_i$ that “fools” $\mathcal{F}$ in the sense that $\sum_{i \in A} a_i \leq 1$ for all $A \in \mathcal{F}$. A nice example is the family $\mathcal{F}$ of all sets A such that $|A| \leq \min A$, which is fooled by the harmonic series. But there are richer, more complicated examples of families without an infinite ascending chain, for which it is hard to explicitly find a fooling series.

4.2. Separation of Row and Column Densities in Infinite Binary Matrices

The notions of upper/lower asymptotic density of a subset of $\mathbb{N}$, or equivalently of an infinite binary sequence, play an important role in number theory. For a sequence $\mathbf{a} = (a_i)_{i \in \mathbb{N}} \in \{0, 1\}^{\mathbb{N}}$, they are respectively defined by

$$\bar{d}(\mathbf{a}) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n a_i \quad \text{and} \quad \underline{d}(\mathbf{a}) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n a_i.$$

If the two are equal, their common value $d(\mathbf{a})$ is called the asymptotic density of $\mathbf{a}$.

Here we look at an $\mathbb{N}$ by $\mathbb{N}$ binary matrix $A = (a_{ij})_{i \in \mathbb{N}, j \in \mathbb{N}} \in \{0, 1\}^{\mathbb{N} \times \mathbb{N}}$. We denote its i -th row by $\mathbf{r}^{(i)} = (a_{ij})_{j \in \mathbb{N}}$ and its j -th column by $\mathbf{c}^{(j)} = (a_{ij})_{i \in \mathbb{N}}$. How are the row and column upper/lower asymptotic densities related?

Definition 4.4. Let A be an $\mathbb{N}$ by $\mathbb{N}$ binary matrix. We say that the row and column upper/lower asymptotic densities are *separated* if there exist real numbers $0 \leq \alpha < \beta \leq 1$ such that

$$\bar{d}(\mathbf{r}^{(i)}) \leq \alpha \quad \text{and} \quad \underline{d}(\mathbf{c}^{(j)}) \geq \beta \quad \text{for all } i, j \in \mathbb{N}. \quad (11)$$

Clearly, the analog of (11) would be impossible in a finite matrix. But such separation can occur in an $\mathbb{N}$ by $\mathbb{N}$ matrix. Indeed, in the binary matrix from Example 1.3 (in which $a_{ij} = 1 \Leftrightarrow i \geq j$), every row has density 0 whereas every column has density 1. We call this matrix fully lower-triangular. We say that a matrix A contains a fully lower-triangular submatrix, if it has a submatrix which, possibly after reordering rows and columns, is fully

lower-triangular. Here we show, as an application of Theorem 3.2, that separation can only occur in the presence of a fully lower-triangular submatrix.

Theorem 4.5. *Let A be an $\mathbb{N}$ by $\mathbb{N}$ binary matrix in which the row and column upper/lower asymptotic densities are separated. Then A contains a fully lower-triangular submatrix.*

Proof. Consider the game $G = (S, T, \pi)$ represented by the matrix A (here $S = T = \mathbb{N}$). We claim that

$$\sup_{\mathbf{p} \in \Delta(S)} \inf_{\mathbf{q} \in \Delta(T)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq \alpha < \beta \leq \inf_{\mathbf{q} \in \Delta(T)} \sup_{\mathbf{p} \in \Delta(S)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}), \quad (12)$$

where α and β witness the separation condition (11) for A . This will show that G violates the minimax property. By Theorem 3.2, it will follow that G has an LNG subgame, or equivalently, A contains a fully lower-triangular submatrix, completing the proof.

We prove the first inequality in (12); the proof of the last inequality is similar. Let $\mathbf{p} = (p_i)_{i \in S} \in \Delta(S)$ and $\varepsilon > 0$ be given. We have to find some $\mathbf{q} \in \Delta(T)$ such that $\pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) \leq \alpha + \varepsilon$. First, we find a large-enough k so that $\sum_{i=1}^k p_i \geq 1 - (\varepsilon/2)$. Next, for each $i = 1, \dots, k$, we look at the corresponding row $\mathbf{r}^{(i)}$ in the matrix A . By (11) its upper density is at most α , so there exists $n_i \in \mathbb{N}$ such that $(1/n) \sum_{j=1}^n a_{ij} \leq \alpha + (\varepsilon/2)$ for all $n \geq n_i$. Now let $n = \max\{n_1, \dots, n_k\}$ and let $\mathbf{q} \in \Delta(T)$ be the uniform distribution over the first n columns. Then we have

$$\pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^{\infty} \sum_{j=1}^n \frac{p_i}{n} a_{ij} = \sum_{i=1}^k \frac{p_i}{n} \sum_{j=1}^n a_{ij} + \sum_{i=k+1}^{\infty} \frac{p_i}{n} \sum_{j=1}^n a_{ij} \leq \left(\alpha + \frac{\varepsilon}{2}\right) + \frac{\varepsilon}{2} = \alpha + \varepsilon,$$

as desired. $\square$

4.3. Duality of Fractional Matchings and Covers in Infinite Hypergraphs

A hypergraph is a pair $H = (V, \mathcal{E})$, where V is a set of vertices and $\mathcal{E}$ is a set of subsets of V , called edges. The sets V and $\mathcal{E}$, and the edges E in $\mathcal{E}$, may be finite or infinite. A fractional matching in H is a function $f : \mathcal{E} \rightarrow \mathbb{R}_+$ such that $\sum_{E \ni v} f(E) \leq 1$ for every vertex $v \in V$. Its size is $|f| = \sum_{E \in \mathcal{E}} f(E)$. A fractional cover in H is a function $g : V \rightarrow \mathbb{R}_+$ such that $\sum_{v \in E} g(v) \geq 1$ for every edge $E \in \mathcal{E}$. Its size is $|g| = \sum_{v \in V} g(v)$. Two important hypergraph invariants are

$$\nu^*(H) = \sup\{|f| : f \text{ is a fractional matching in } H\}$$

and

$$\tau^*(H) = \inf\{|g| : g \text{ is a fractional cover in } H\}.$$

Both invariants may be finite or infinite.

Fractional matchings and covers in finite hypergraphs have numerous applications in combinatorics and discrete optimization. A fundamental fact is the equality $\nu^*(H) = \tau^*(H)$ which holds for every finite hypergraph by the duality theorem of linear programming. What can be said for infinite hypergraphs?

Definition 4.6. We say that the hypergraph H satisfies the *weak duality property* if $\nu^*(H) = \tau^*(H)$.

A simple example of an infinite hypergraph where weak duality fails is the *tail hypergraph* $H = (\mathbb{N}, \mathcal{E})$, where the edges in $\mathcal{E}$ are all sets of the form $E_j = \{i \in \mathbb{N} : i \geq j\}$, $j \in \mathbb{N}$. Here $\nu^*(H) = 1$ whereas $\tau^*(H) = \infty$.

A partial subhypergraph of $H = (V, \mathcal{E})$ is obtained by restricting V and $\mathcal{E}$ to some subsets: it is of the form $H' = (V', \mathcal{E}')$ where $V' \subseteq V$ and $\mathcal{E}' = \{E \cap V' : E \in \mathcal{E}_0\}$ for some $\mathcal{E}_0 \subseteq \mathcal{E}$. As an application of Theorem 3.2, we characterize here those hypergraphs all of whose partial subhypergraphs satisfy the weak duality property.

Theorem 4.7. *Let $H = (V, \mathcal{E})$ be a hypergraph. Then every partial subhypergraph H' of H satisfies the weak duality property if and only if H has no partial subhypergraph H' which is isomorphic to the tail hypergraph.*

Proof. Given a hypergraph $H = (V, \mathcal{E})$, we associate with it a game $G = G(H) = (V, \mathcal{E}, \pi)$ in which player 1 chooses a vertex $v \in V$, player 2 chooses an edge $E \in \mathcal{E}$, and player 1 wins if and only if $v \in E$. Note that if H is the tail hypergraph, then $G(H)$ is isomorphic to the LNG. We claim that

$$\sup_{\mathbf{p} \in \Delta(V)} \inf_{\mathbf{q} \in \Delta(\mathcal{E})} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \frac{1}{\tau^*(H)} \quad \text{and} \quad \inf_{\mathbf{q} \in \Delta(\mathcal{E})} \sup_{\mathbf{p} \in \Delta(V)} \pi^{\text{mix}}(\mathbf{p}, \mathbf{q}) = \frac{1}{\nu^*(H)}. \quad (13)$$

This will show that H satisfies the weak duality property if and only if $G(H)$ satisfies the minimax property. Clearly, partial subhypergraphs correspond to subgames, so the present theorem is in fact a reformulation of Theorem 3.2.

We prove the first part of (13); the proof of the second part is similar. If g is a fractional cover in H of finite size $|g|$, then the mixed strategy $\mathbf{p} = (p_v)_{v \in V}$ in which $p_v = g(v)/|g|$ guarantees to player 1 in $G(H)$ an expected payoff of at least $1/|g|$. Conversely, if $\mathbf{p} = (p_v)_{v \in V}$ is a mixed strategy of player 1 in $G(H)$ that guarantees an expected payoff of at least $\alpha > 0$, then the function $g: V \rightarrow \mathbb{R}_+$ where $g(v) = p_v/\alpha$ is a fractional cover in H of size $|g| = 1/\alpha$. These two reciprocal maps together establish the first part of (13). $\square$

We note that Aharoni and Holzman [1] studied the duality of fractional matchings and covers in infinite hypergraphs. They introduced two versions of duality: the weak one that we use here, and a stronger one which is an infinite counterpart of the complementary slackness conditions of linear programming duality. The main positive results in Aharoni and Holzman [1] gave sufficient conditions for weak duality, which essentially correspond to conditions 2 and 4 in Corollary 3.3. However, no characterization result corresponding to Theorem 3.2 was obtained there.

4.4. A Characterization of Stable Theories in First-Order Logic

Stability is a central concept in model theory, with many applications to algebra and other areas of mathematics. Shelah [9] established several equivalent definitions of stability, of which we give here one that is related to our work.

Definition 4.8. Let $\mathcal{L}$ be a language in first-order logic, and let T be a complete theory in this language. Let $\varphi(x_1, \dots, x_n; y_1, \dots, y_m)$ be a formula in $\mathcal{L}$ with $n + m$ free variables. We say that φ has the *order property* over T if there exists a model $\mathcal{M}$ of T , and there exist two infinite sequences $(\mathbf{a}_i)_{i=1,2,\dots}$ and $(\mathbf{b}_j)_{j=1,2,\dots}$ where each $\mathbf{a}_i$ is an n -tuple of elements of $\mathcal{M}$ and each $\mathbf{b}_j$ is an m -tuple of elements of $\mathcal{M}$, such that

$$\mathcal{M} \models \varphi(\mathbf{a}_i; \mathbf{b}_j) \Leftrightarrow i \geq j.$$

The theory T is *stable* if no such formula φ has the order property over T .

Here we apply Theorem 3.2 to give a game theoretic characterization of stable theories. Fix a language $\mathcal{L}$ and a theory T as above. Given a formula $\varphi(x_1, \dots, x_n; y_1, \dots, y_m)$ in $\mathcal{L}$ and a model $\mathcal{M}$ of T , we associate with them a game $G = G(\mathcal{L}, T, \varphi, \mathcal{M})$ as follows: player 1 chooses an n -tuple $\mathbf{a}$ of elements of $\mathcal{M}$, player 2 chooses an m -tuple $\mathbf{b}$ of elements of $\mathcal{M}$, and player 1 wins if and only if $\mathcal{M} \models \varphi(\mathbf{a}; \mathbf{b})$.

Theorem 4.9. Let $\mathcal{L}$ be a language in first-order logic, and let T be a complete theory in $\mathcal{L}$. Then T is stable if and only if for every formula $\varphi(x_1, \dots, x_n; y_1, \dots, y_m)$ in $\mathcal{L}$ and every model $\mathcal{M}$ of T the associated game $G = G(\mathcal{L}, T, \varphi, \mathcal{M})$ is totally minimax.

Proof. Fixing $\mathcal{L}, T, \varphi$ and $\mathcal{M}$, we note that a pair of sequences $(\mathbf{a}_i)_{i=1,2,\dots}$ of n -tuples of elements of $\mathcal{M}$ and $(\mathbf{b}_j)_{j=1,2,\dots}$ of m -tuples of elements of $\mathcal{M}$ such that $\mathcal{M} \models \varphi(\mathbf{a}_i; \mathbf{b}_j) \Leftrightarrow i \geq j$ corresponds to an LNG subgame of $G = G(\mathcal{L}, T, \varphi, \mathcal{M})$. Thus, φ has the order property over T if and only if there exists a model $\mathcal{M}$ of T such that $G = G(\mathcal{L}, T, \varphi, \mathcal{M})$ has an LNG subgame. Therefore, by the definition above, T is stable if and only if for every φ and $\mathcal{M}$ the game $G = G(\mathcal{L}, T, \varphi, \mathcal{M})$ is LNG-free, or equivalently (by Theorem 3.2), G is totally minimax. $\square$

We remark that in the case of a countable language $\mathcal{L}$, it is known that it suffices to check the definition above for countable models $\mathcal{M}$ of T . In this case, the games G in Theorem 4.9 are all countable as well.

We note that there are several other examples where strategic games nicely capture concepts in logic and set theory. But we are not aware of such instances where mixed strategies are involved, as is the case here.

4.5. Online Learning with Simple Predictors

We briefly describe here a standard model of online learning. There is a set $\mathcal{X}$ of instances, each of which may be labeled by 0 or 1. A concept class $\mathbb{H} \subseteq \{0,1\}^{\mathcal{X}}$ is given; each hypothesis $h \in \mathbb{H}$ represents a possible labeling of all instances. A learner, who knows the class $\mathbb{H}$ but not the “true” hypothesis, is faced with a sequence of instances that are revealed in online fashion. In each round, the learner tries to guess the label of the given instance, and then learns whether the guess was correct or not, before proceeding to the next round. A fundamental result of Littlestone [8] characterizes those concept classes $\mathbb{H}$ that are online learnable, in the sense that the learner can keep the total number of mistakes below some finite constant. To this end, he defined a combinatorial complexity measure called the Littlestone dimension of $\mathbb{H}$, and showed that $\mathbb{H}$ is online learnable if and only if its Littlestone dimension is finite.

Hanneke et al. [7] showed recently that when $\mathbb{H}$ is online learnable, the learning may be achieved by a “simple predictor” of a certain particular type. Their analysis used the following game $G = G(\mathbb{H})$ associated with the concept class $\mathbb{H}$: the learner chooses a hypothesis $h \in \mathbb{H}$, an adversary chooses a labeled instance $(x, y) \in \mathcal{X} \times \{0,1\}$,

and the learner wins if and only if $h(x) = y$. A key step in their proof is showing that when $\mathbb{H}$ has finite Littlestone dimension, the game $G(\mathbb{H})$ satisfies the minimax property. This motivated them to prove, more generally, that a game G of finite VC-dimension (see explanation in Section 3) is totally minimax if and only if it is LNG-free. Our Theorem 3.2 generalizes their result by removing the assumption of finite VC-dimension. This has no consequence for their original application, though, because the VC-dimension is bounded above by the Littlestone dimension.

Incidentally, the argument in Hanneke et al. [7] involved yet another complexity measure from learning theory: the threshold dimension, which is related to LNG-freeness. In a previous paper, Alon et al. [2] had proved that the threshold dimension is finite if and only if the Littlestone dimension is (this was essentially a learning theory reformulation of a result of Shelah [9] about stability—see Subsection 4.4). Hanneke et al. [7] used this to show that if $\mathbb{H}$ has finite Littlestone dimension, then $G(\mathbb{H})$ is LNG-free, allowing them to deduce the minimax property in their application.

Acknowledgments

The research for this paper originated in discussions with Shay Moran, aiming to generalize his result with Hanneke and Livni. The author is grateful to him and to Ron Aharoni, Zachary Chase, Bogdan Chornomaz, Yuval Dagan, Yuval Filmus, and Steve Hanneke for many helpful discussions. In particular, Yuval Dagan suggested the argument in the last step of the proof of Theorem 3.2; Steve Hanneke communicated his result mentioned in Remark 3.4; the application presented in Subsection 4.2 came up in discussions among some of the above-mentioned people. The author thanks two anonymous reviewers for their useful comments.

References

- [1] Aharoni R, Holzman R (1992) Optimal fractional matchings and covers in infinite hypergraphs: Existence and duality. *Graphs Combin.* 8(1):11–22.
- [2] Alon N, Livni R, Malliaris M, Moran S (2019) Private PAC learning implies finite Littlestone dimension. Charikar M, Cohen E, eds. *STOC'19 Proc. 51st Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 852–860.
- [3] Billingsley P (1999) *Convergence of Probability Measures*, Wiley Series in Probability and Statistics (John Wiley & Sons, Inc., New York).
- [4] Daskalakis C, Golowich N (2022) Fast rates for nonparametric online learning: From realizability to learning in games. Leonardi S, Gupta A, eds. *STOC'22 Proc. 54th Annual AGM SIGACT Sympos. Theory Comput.* (ACM, New York), 846–859.
- [5] Fan K (1953) Minimax theorems. *Proc. Natl. Acad. Sci. USA* 39(1):42–47.
- [6] Hanneke S (2023) Personal communication, August 26.
- [7] Hanneke S, Livni R, Moran S (2021) Online learning with simple predictors and a combinatorial characterization of minimax in 0/1 games. *Proc. 34th Conf. Learn. Theory*, vol. 134 (PMLR, New York), 2289–2314.
- [8] Littlestone N (1987) Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learn.* 2(4):285–318.
- [9] Shelah S (1978) *Classification Theory and the Number of Nonisomorphic Models*, Studies in Logic and the Foundations of Mathematics, vol. 92 (North-Holland Publishing Co., Amsterdam).
- [10] Simons S (1995) Minimax theorems and their proofs. Du DZ, Pardalos PM, eds. *Minimax and Applications*, Nonconvex Optimization and Its Applications, vol. 4 (Kluwer Academic Publishing, Dordrecht, Netherlands), 1–23.
- [11] Sion M (1958) On general minimax theorems. *Pacific J. Math.* 8(1):171–176.
- [12] von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100(1):295–320.
- [13] Yanovskaya EB (1970) Solution of infinite antagonistic games in finitely additive strategies. *Teor. Veroyatn. Primen.* 15(1):162–168.