

Separable and Low-Rank Continuous Games

Noah D. Stein, Asuman Ozdaglar, and Pablo A. Parrilo^{*†}

July 23, 2007

Abstract

In this paper, we study nonzero-sum separable games, which are continuous games whose payoffs take a sum-of-products form. Included in this subclass are all finite games and polynomial games. We investigate the structure of equilibria in separable games. We show that these games admit finitely supported Nash equilibria. Motivated by the bounds on the supports of mixed equilibria in two-player finite games in terms of the ranks of the payoff matrices, we define the notion of the rank of an n -player continuous game and use this to provide bounds on the cardinality of the support of equilibrium strategies. We present a general characterization theorem that states that a continuous game has finite rank if and only if it is separable. Using our rank results, we present an efficient algorithm for computing approximate equilibria of two-player separable games with fixed strategy spaces in time polynomial in the rank of the game.

1 Introduction

The structure and computation of equilibria in games with infinite strategy spaces have long been recognized as complex. Even seemingly “good” games may possess only “bad” equilibria; Gross has constructed an example of a zero-sum game with rational utility functions whose unique Nash equilibrium is for each player to play the Cantor measure [11, 13]. To avoid such pathologies, Dresher, Karlin, and Shapley introduced the class of zero-sum separable games [9, 8, 14, 13]. These are games in which each player’s payoff can be written as a sum of products of functions in each player’s strategy separately (e.g. as polynomials), and it is known that every separable game admits a mixed strategy Nash equilibrium that is finitely supported, i.e. all players randomize over a finite number of pure strategies. Parrilo has shown that Nash equilibria of zero-sum games with polynomial utility functions can be computed efficiently [19].

^{*}Department of Electrical Engineering, Massachusetts Institute of Technology: Cambridge, MA 02139. nstein@mit.edu, asuman@mit.edu, and parrilo@mit.edu.

[†]This research was funded in part by National Science Foundation grants DMI-0545910 and ECCS-0621922 and AFOSR MURI subaward 2003-07688-1.

In this paper, we study nonzero-sum separable games. We show that even in the nonzero-sum case, separable games have the property that an equilibrium exists in finitely supported mixed strategies. We then characterize the structure of these games and their Nash equilibria, and also propose methods for computing (exact and approximate) equilibria.

Our first major contribution is to define and develop the concept of the *rank of a continuous game*, which we use to construct bounds on the number of strategies played in equilibrium. In two-player finite games, Lipton et al. have recently shown that the ranks of the payoff matrices provide such a bound [16]. Our new definition of rank generalizes this one to allow for an arbitrary finite number of players (a problem explicitly left open in [16]) and infinite strategy spaces. We define the rank of a continuous game by introducing an equivalence relation between mixed strategies called *almost payoff equivalence*. The rank is the dimension of the mixed strategy space modulo this equivalence relation. Loosely speaking, low-rank continuous games are those where the variation in each player's payoff depends only on a low-dimensional projection of the mixed strategies of the players. For example in games with polynomial payoffs, the rank depends on the dimension of a projection of the moment space.

We also show that a continuous game has finite rank if and only if it is separable. This means that little can be said about the structure of equilibria in non-separable continuous games and highlights the importance of separable games. We provide simple techniques for evaluating the rank of separable games, which we specialize to games with polynomial payoffs and finite games.

Our second set of results concern efficient computation of mixed strategy equilibria in separable games. For n -player games, we provide a nonlinear optimization formulation, and show that the optimal solutions of this problem correspond to the (generalized) moments of exact Nash equilibria. This formulation generalizes the optimization formulation of Nash equilibria of finite games presented by Başar and Olsder [1].

While the nonlinear optimization formulation for the computation of equilibria is tractable for certain classes of separable games such as zero-sum polynomial games, the potential non-convexity of the optimization problem makes it impractical in other cases. We therefore supplement our computational results for exact mixed strategy equilibria with new methods for computing approximate mixed strategy equilibria. Using our rank results described above, we can link the computation of mixed strategy equilibria to the appropriate discretization of the strategy space into finite actions. For two-person separable games, this yields an efficient algorithm for computing approximate mixed strategy equilibria in time polynomial in the rank of the game. This algorithm searches for finitely supported approximate equilibria by enumerating all possible supports, and relies on the fact that the search can be limited to supports of size bounded by the rank of the game. For two-players, the set of equilibria for a given support can be described by linear equations and inequalities¹.

Our work is related to a number of literatures. There has been considerable work on the computation of equilibria in finite games. Lemke and Howson give a path-following al-

¹With more than two players this description is no longer linear (or convex), so this algorithm does not generalize to n -player games with $n > 2$.

gorithm for two-player finite games which can be viewed as the simplex method for linear programming operating with a different pivoting rule [15]. To find equilibria of games with more players, Scarf constructs a simplicial subdivision algorithm which also works for more general fixed point problems [22]. These methods rely on the polyhedral structure of the mixed strategy spaces of finite games, therefore they seem unlikely to generalize to continuous/separable games. For a survey of algorithms which compute equilibria of finite games, see [17].

Another growing literature has been investigating the complexity of computing mixed strategy Nash equilibria of finite games. Daskalakis, Goldberg, and Papadimitriou settle this question for finite normal form games with four or more players, showing that the problem of computing a single Nash equilibrium is PPAD-complete [6]. In essence this means that it is computationally equivalent to a number of other fixed point problems which are believed to be computationally difficult. These problems share the feature that a solution can be proven to exist, but the known proofs of existence are inefficient; for more about the complexity class PPAD, see [18]. Daskalakis and Papadimitriou later improve this result by proving PPAD-completeness in the case of three players [7]. Chen and Deng also prove this independently [4] and complete this line of work by proving PPAD-completeness for two players [5]. In this literature, there has been no work on continuous games.

Our work is most closely related to the work of Lipton et al. [16], who consider two-player finite games and provide bounds on the support of equilibrium strategies using the ranks of the payoff matrices of the players. Since finite games are a special case of separable games, our results on the cardinality of the support of equilibrium strategies generalize theirs by allowing for an arbitrary finite number of players as well as infinite strategy sets and a broader class of payoff functions.

Lipton et al. also investigate computing approximate equilibria in two-player finite games and present the first algorithm for computing approximate equilibria which is quasi-polynomial in the number of strategies [16]. The rank of a separable game measures the complexity of the payoffs, and in the case of a finite game it is bounded by the number of strategies. Therefore in finite games the complexity of the payoffs and the complexity of the strategy spaces do not vary independently; the effects of these parameters on the running time of algorithms cannot be distinguished. However, for games with infinite strategy sets the rank can be varied arbitrarily while the strategy set is held fixed. This allows us to construct an algorithm for computing approximate equilibria in two-player separable games with fixed (infinite) strategy spaces and show that it has a polynomial dependence on the rank. Since we assume the strategy space is fixed and consider general separable games rather than only finite games, our algorithm is not directly comparable to that of Lipton et al. [16]. Nonetheless it is interesting on its own since no algorithm for computing approximate equilibria of finite games is known which has polynomial dependence on the complexity of the game.

Our work is also related to that of Kannan and Theobald, who study a different notion of rank in two-player finite games [12]. They take an algorithmic perspective and view zero-sum games as the simplest type of games. To generalize these, they propose a hierarchy of classes of two-player finite games in which the rank of the sum $R+C$ of the payoff matrices is bounded

by a constant k ; the case $k = 0$ corresponds to zero-sum games. For fixed k , they show that approximate Nash equilibria of two-player finite games can be computed in time polynomial in the description length for the game. This algorithm relies on an approximation result for quadratic programming due to Vavasis [25] which depends on the polyhedral structure of the problem. It is conceivable that this technique may apply to polynomial games if the approximation technique can be extended to this more general algebraic setting, but we do not do so here.

The rest of this paper is organized as follows. In Section 2 we define separable games and prove some of their basic properties. Then in Section 3 we define the rank of a continuous game and use this definition to bound the cardinality of the support of Nash equilibria. We present a characterization theorem for separable games which in particular shows that within the class of continuous games, the low-rank results in this paper cannot be extended past separable games. We provide a simple formula for computing the rank of arbitrary separable games, which we specialize for finite and polynomial games. In Section 4 we discuss computation of Nash equilibria and approximate equilibria. We close with conclusions and directions for future work.

2 Basic Theory of Separable Games

We first introduce the notational conventions and the basic terminology used throughout the paper. Subscripts refer to players, while superscripts are reserved for other indices, rather than exponents. If S_j are sets for $j = 1, \dots, n$ then $S = \prod_{j=1}^n S_j$ and $S_{-i} = \prod_{j \neq i} S_j$. The n -tuple s and the $(n-1)$ -tuple s_{-i} are formed from the points s_j similarly. Given a subset S of a vector space, we use the notation $\text{span } S$, $\text{aff } S$, $\text{conv } S$, and \bar{S} to denote the span, affine hull, convex hull, and closure of the set S , respectively. We denote the transpose of a matrix M by M' .

Definition 2.1. An n -player **continuous game** consists of a **pure strategy** space C_i which is a nonempty compact metric space and a continuous **utility** or **payoff function** $u_i : C \rightarrow \mathbb{R}$ for each player $i = 1, \dots, n$. Throughout, Δ_i will denote the space of Borel probability measures σ_i over C_i , referred to as **mixed strategies**, and the u_i will be extended from C to Δ by expectation, defining

$$u_i(\sigma) = \int_C u_i(s) d\sigma,$$

where the n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ is identified with the product measure $\sigma_1 \times \dots \times \sigma_n$.

Definition 2.2. An n -player **separable game** is an n -player continuous game with utility functions $u_i : C \rightarrow \mathbb{R}$ taking the form

$$u_i(s) = \sum_{j_1=1}^{m_1} \dots \sum_{j_n=1}^{m_n} a_i^{j_1 \dots j_n} f_1^{j_1}(s_1) \dots f_n^{j_n}(s_n), \quad (1)$$

where $a_i^{j_1 \dots j_n} \in \mathbb{R}$ and the $f_i^j : C_i \rightarrow \mathbb{R}$ are continuous. Two important special cases are the **finite games** in which the C_i are finite and the u_i are arbitrary and the **polynomial games** in which the C_i are compact intervals in \mathbb{R} and the u_i are polynomials in the s_j .

When it is convenient to do so, and always for polynomial games, we will begin the summations in (1) at $j_i = 0$. For polynomial games we can then use the convention that $f_i^j(s_i) = s_i^j$, where the superscript on the right hand side denotes an exponent rather than an index.

Example 2.3. Consider a two player game with $C_1 = C_2 = [-1, 1] \subset \mathbb{R}$. Letting x and y denote the pure strategies of players 1 and 2, respectively, we define the utility functions

$$\begin{aligned} u_1(x, y) &= 2xy + 3y^3 - 2x^3 - x - 3x^2y^2 \quad \text{and} \\ u_2(x, y) &= 2x^2y^2 - 4y^3 - x^2 + 4y + x^2y. \end{aligned} \tag{2}$$

This is a polynomial game, and we will return to it periodically to apply the results presented. In particular, we will show using classical techniques that this game must have a Nash equilibrium in which each player randomizes over a set of cardinality at most 5. We will then apply our new rank results (see Theorem 3.3) to reduce this bound to 2 for the first player and 4 for the second player.

Let V_i denote the space of all finite-valued signed measures (henceforth simply called measures) on C_i , which can be identified with the dual of the Banach space $C(C_i)$ of all continuous real-valued functions on C_i endowed with the sup norm. Throughout, we will use the **weak*** topology on V_i , which is the weakest topology such that whenever $f : C_i \rightarrow \mathbb{R}$ is a continuous function, $\sigma \mapsto \int f d\sigma$ is a continuous linear functional on V_i .

We can extend the utility functions of a continuous game to all of V in the same way they are extended from C to Δ , yielding a multilinear functional on V . For a fixed separable game we can extend the f_i^j from C_i to V_i similarly, yielding the so-called **generalized moment functionals**, so (1) holds with s replaced by $\sigma \in V$. In polynomial games $f_i^j(s_i) = s_i^j$ so the generalized moment functionals are just the classical moment functionals. We will abuse notation and identify the elements of C_i with the atomic measures in Δ_i , so $C_i \subseteq \Delta_i \subset V_i$. Note that $\text{conv } C_i$ and $\text{span } C_i$ are the set of all finitely supported probability measures and the set of all finitely supported finite signed measures, respectively. The following are standard results (see [20] and [21]).

Proposition 2.4.

- (a) *The sets C_i and Δ_i are weak* compact.*
- (b) *The weak* closures of $\text{conv } C_i$ and $\text{span } C_i$ are Δ_i and V_i , respectively.*
- (c) *The weak* topology makes V_i a Hausdorff topological vector space.*

We next define three notions of equivalence between two measures, which allow us to exhibit simplifications in the structure of Nash equilibria in separable games.

Definition 2.5. Two measures $\sigma_i, \tau_i \in V_i$ are

- **moment equivalent** if $f_i^j(\sigma_i) = f_i^j(\tau_i)$ for all j (representation-dependent and only defined for separable games).
- **payoff equivalent** if $u_j(\sigma_i, s_{-i}) = u_j(\tau_i, s_{-i})$ for all j and all $s_{-i} \in C_{-i}$.
- **almost payoff equivalent** if $u_j(\sigma_i, s_{-i}) = u_j(\tau_i, s_{-i})$ for all $j \neq i$ and all $s_{-i} \in C_{-i}$.

Note that in separable games moment equivalence implies payoff equivalence and in all continuous games payoff equivalence implies almost payoff equivalence. Since the f_i^j and u_j are linear and multilinear functionals on V_i and V , respectively, these equivalence relations can be expressed in terms of (potentially infinitely many) linear constraints on $\sigma_i - \tau_i$.

Definition 2.6. Let 0 denote the zero measure in V_i and define

- $W_i = \{\text{measures moment equivalent to } 0\}$,
- $X_i = \{\text{measures payoff equivalent to } 0\}$,
- $Y_i = \{\text{measures almost payoff equivalent to } 0\}$.

Then $W_i \subseteq X_i \subseteq Y_i$, and $\sigma_i - \tau_i \in X_i$ if and only if σ_i is payoff equivalent to τ_i , etc. Furthermore, the subspaces X_i and Y_i are representation-independent and well-defined for any continuous game, separable or not. Note that these subspaces are given by the intersection of the kernels of (potentially infinitely many) continuous linear functionals, hence they are closed.

We will analyze separable games by considering the quotients of V_i by these subspaces, i.e. $V_i \text{ mod these three equivalence relations}$. To avoid defining excessively many symbols let Δ_i/W_i denote the image of Δ_i in V_i/W_i and so forth. For a concrete example of these sets see the discussion at end of this section, and in particular Figure 1.

The following theorem presents a fundamental result about separable games. It shows that regardless of the choices of the other players, each player is free to restrict his choice of strategies to a particularly simple class of mixed strategies, namely those which only assign positive probability to a finite number of pure strategies. Furthermore, a bound on the number of strategies needed can be easily computed in terms of the structure of the game. This theorem can be proven by a separating hyperplane argument as applied to zero-sum separable games by Karlin [13], but here we give a new topological argument.

Theorem 2.7. *In a separable game every mixed strategy σ_i is moment equivalent to a finitely-supported mixed strategy τ_i with $|\text{supp}(\tau_i)| \leq m_i + 1$. Moreover, if σ_i is countably-supported τ_i can be chosen with $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$.*

Proof. Note that the map

$$f_i : \sigma_i \mapsto (f_i^1(\sigma_i), \dots, f_i^{m_i}(\sigma_i)) \tag{3}$$

is linear and continuous. Therefore

$$f_i(\Delta_i) = f_i(\overline{\text{conv } C_i}) \subseteq \overline{f_i(\text{conv } C_i)} = \overline{\text{conv } f_i(C_i)} = \text{conv } f_i(C_i) = f_i(\text{conv } C_i) \subseteq f_i(\Delta_i).$$

The first three steps follow from Proposition 2.4, continuity of f_i , and linearity of f_i , respectively. The next equality holds because $\text{conv } f_i(C_i)$ is compact, being the convex hull of a compact subset of a finite-dimensional space. The final two steps follow from the linearity of f_i and the containment $\text{conv } C_i \subseteq \Delta_i$. Hence, we have

$$f_i(\Delta_i) = \text{conv } f_i(C_i) = f_i(\text{conv } C_i).$$

This shows that any mixed strategy is moment equivalent to a finitely-supported mixed strategy. Applying Carathéodory's theorem [2] to the set $\text{conv } f_i(C_i) \subset \mathbb{R}^{m_i}$ yields the uniform bound. Since a countable convex combination of points in a bounded subset of \mathbb{R}^{m_i} can always be written as a finite convex combination of at most $m_i + 1$ of those points, the final claim follows. \square

For the rest of the paper we will study the Nash equilibria of (nonzero-sum) separable games, which are defined for arbitrary continuous games as follows.

Definition 2.8. A mixed strategy profile σ is a **Nash equilibrium** if $u_i(\tau_i, \sigma_{-i}) \leq u_i(\sigma)$ for all i and all $\tau_i \in \Delta_i$.

Combining Theorem 2.7 with Glicksberg's result [10] that every continuous game has a Nash equilibrium yields the following:

Corollary 2.9. *Every separable game has a Nash equilibrium in which player i mixes among at most $m_i + 1$ pure strategies.*

Example 2.3 (cont'd). Apply the standard definition of the f_i^j to the polynomial game with payoffs given in (2). The set of moments $\Delta_i/W_i \cong f_i(\Delta_i)$ as described in Theorem 2.7 is shown in Figure 1 with the zeroth moment omitted (it is identically unity). In this case the set of moments is the same for both players. The space $V_i/W_i = f_i(V_i)$ is a four-dimensional real vector space, and Figure 1 shows a subset of the projection of V_i/W_i onto its final three coordinates. The set $C_i/W_i = f_i(C_i)$ of moments due to pure strategies is the curve traced out by the vector $(1, x, x^2, x^3)$ as x varies from -1 to 1 . Since the first coordinate is omitted in the figure, this can be seen as the sharp edge which goes from $(-1, 1, -1)$ through $(0, 0, 0)$ to $(1, 1, 1)$. As shown in the proof of Theorem 2.7, $f_i(\Delta_i)$ is exactly the convex hull of this curve.

For each player the range of the indices in (1) is $0 \leq j_i \leq 3$, so by Corollary 2.9, this game has an equilibrium in which each player mixes among at most $4 + 1 = 5$ pure strategies. To produce this bound, we have not used any information about the payoffs except for the degree of the polynomials. However, notice that there is extra structure here to be exploited. For example, u_2 depends on the expected value $\int x^2 d\sigma_1(x)$, but not on $\int x d\sigma_1(x)$ or $\int x^3 d\sigma_1(x)$. In particular, player 2 is indifferent between the two strategies $\pm x$ of player 1




Figure 1: The space $f_i(\Delta_i) \cong \Delta_i/W_i$ of possible moments for either player's mixed strategy under the payoffs given in (2) due to a measure σ_i on $[-1, 1]$. The zeroth moment, which is identically unity, has been omitted.

for all x , insofar as this choice does not affect his payoff (though it does affect what strategy profiles are equilibria). In the following section, we will present improved bounds on the number of strategies played in an equilibrium which take these simplifications into account in a systematic manner.

3 The Rank of a Continuous Game

In this section, we introduce the notion of the *rank of a continuous game*. We use this notion to provide improved bounds on the cardinality of the support of the equilibrium strategies for a separable game. We also provide a characterization theorem which states that a continuous game has finite rank if and only if it is separable, thus showing that separable games are the largest class of continuous games to which low-rank arguments apply (see subsection 3.2). We conclude this section by showing how to compute the rank of a separable game in Subsection 3.3, which may be read independently of the section on the characterization theorem.

Our main result on bounding the support of equilibrium strategies generalizes the following theorem of Lipton et al. [16] to arbitrary separable games, thereby removing the two-player assumption and weakening the restriction that the strategy spaces be finite. The proof given in [16] is essentially an algorithmic version of Carathéodory's theorem. Here we provide a shorter nonalgorithmic proof which illustrates some of the ideas to be used in

establishing the extended version of the theorem.

Theorem 3.1 (Lipton et al. [16]). *Consider a two-player finite (i.e. bimatrix) game defined by matrices R and C of payoffs to the row and column player, respectively. Let σ be a Nash equilibrium of the game. Then there exists a Nash equilibrium τ which yields the same payoffs to both players as σ , but in which the column player mixes among at most rank $R + 1$ pure strategies and the row player mixes among at most rank $C + 1$ pure strategies.*

Proof. Let r and c be probability column vectors corresponding to the mixed strategies of the row and column players in the given equilibrium σ . Then the payoffs to the row and column players are $r'Rc$ and $r'Cc$. Since c is a probability vector, we can view Rc as a convex combination of the columns of R . These columns all lie in the column span of R , which is a vector space of dimension rank R . By Carathéodory's theorem, we can therefore write any convex combination of these vectors using only rank $R + 1$ terms [2]. That is to say, there is a probability vector d such that $Rd = Rc$, d has at most rank $R + 1$ nonzero entries, and a component of d is nonzero only if the corresponding component of c was nonzero.

Since r was a best response to c and $Rc = Rd$, r is a best response to d . On the other hand, since (r, c) was a Nash equilibrium c must have been a mixture of best responses to r . But d only assigns positive probability to strategies to which c assigned positive probability. Thus d is a best response to r , so (r, d) is a Nash equilibrium which yields the same payoffs to both players as (r, c) , and d only assigns positive probability to rank $R + 1$ pure strategies. Applying the same procedure to r we could find an s which only assigns positive probability to rank $C + 1$ pure strategies and such that (s, d) is a Nash equilibrium with the same payoffs to both players as (r, c) . \square

We will show in the following subsection that our extension of this theorem provides a slightly tighter bound of rank R instead of rank $R + 1$ (respectively for C) in some cases, depending on the structure of R and C . This can be seen directly from the proof given here. We considered convex combinations of the columns of R and noted that these all lie in the column span of R . In fact, they all lie in the affine hull of the set of columns of R . If this affine hull does not include the origin, then it will have dimension rank $R - 1$. The rest of the proof goes through using this affine hull instead of the column span, and so in this case we get a bound of rank R on the number of strategies played with positive probability by the row player. Alternatively, the dimension of the affine hull can be computed directly as the rank of the matrix given by subtracting a fixed column of R from all the other columns of R .

3.1 Bound on the Support of Equilibrium Strategies

By comparing the two-player case of (1) with the singular value decomposition for matrices, separable games can be thought of as games of “finite rank.” Now we will define the rank of a continuous game precisely and use it to give a bound on the cardinality of the support of equilibrium strategies, generalizing Corollary 2.9 and Theorem 3.1. The primary tool will be

the notion of almost payoff equivalence from Definition 2.5. In what follows, the dimension of a set will refer to the dimension of its affine hull.

Definition 3.2. The **rank** of a continuous game is defined to be the n -tuple $\rho = (\rho_1, \dots, \rho_n)$ where $\rho_i = \dim \Delta_i/Y_i$. A game is said to have **finite rank** if $\rho_i < \infty$ for all i .

Since $W_i \subseteq Y_i$ and V_i/W_i is finite-dimensional for any separable game (see the proof of Theorem 2.7), it is clear that separable games have finite rank, and furthermore that $\rho_i \leq m_i$ for all i . In subsection 3.2 we show that all games of finite rank are also separable, so the two conditions are equivalent. Unlike the m_i , the rank is defined for all continuous games and is representation-independent.

We define the rank of a game in terms of almost payoff equivalence of mixed strategies by means of the subspace Y_i . This notion of equivalence between two strategies for a given player means that no matter what the other players do, their payoffs will not enable them to distinguish these two strategies. The use of this equivalence relation reflects the fact that at a (mixed strategy) Nash equilibrium, the payoff to a player is equalized among all pure strategies which he plays with positive probability. Therefore the player is free to switch to any other mixed strategy supported on the same pure strategies, as long as the change does not affect the other players, and the resulting strategy profile will remain a Nash equilibrium.

Using the rank of a game, Corollary 2.9 and Theorem 3.1 can now be improved as follows:

Theorem 3.3. *Given an equilibrium σ of a separable game with rank ρ , there exists an equilibrium τ such that each player i mixes among at most $\rho_i + 1$ pure strategies and $u_i(\sigma) = u_i(\tau)$.*

If $\dim \Delta_i/X_i = 1$ and the metric space C_i is connected, then this bound can be improved so that τ_i is a pure strategy.

Proof. By Theorem 2.7, we can assume without loss of generality that each player's mixed strategy σ_i is finitely supported. Fix i , let $\psi_i : V_i \rightarrow V_i/Y_i$ denote the canonical projection transformation and let $\sigma_i = \sum_j \lambda^j s_i^j$ be a finite convex combination of pure strategies. By linearity of ψ_i we have

$$\psi_i(\sigma_i) = \sum_j \lambda^j \psi_i(s_i^j).$$

Carathéodory's theorem states that (renumbering the s_i^j and adding some zero terms if necessary) we can write

$$\psi_i(\sigma_i) = \sum_{j=0}^{\rho_i} \mu^j \psi_i(s_i^j),$$

a convex combination potentially with fewer terms. Let $\tau_i = \sum_{j=0}^{\rho_i} \mu^j s_i^j$. Then $\psi_i(\sigma_i) = \psi_i(\tau_i)$. Since σ was a Nash equilibrium, and σ_i is almost payoff equivalent to τ_i , σ_j is a best response to $(\tau_i, \sigma_{-i,j})$ for all $j \neq i$. On the other hand σ_i was a mixture among best responses to the mixed strategy profile σ_{-i} , so the same is true of τ_i , making it a best response to σ_{-i} . Thus (τ_i, σ_{-i}) is a Nash equilibrium.

If $\dim \Delta_i/X_i = 1$ and C_i is connected, then C_i/X_i is connected, compact, and one-dimensional, i.e. it is an interval. Therefore it is convex, so $\Delta_i/X_i = \text{conv}(C_i/X_i) = C_i/X_i$. This implies that there exists a pure strategy s_i which is payoff equivalent to σ_i , so we may take $\tau_i = s_i$ and (τ_i, σ_{-i}) is a Nash equilibrium.

Beginning with this equilibrium and repeating the above steps for each player in turn completes the construction of τ and the final statement of the theorem is clear. \square

While the preceding theorem was the original reason for our choice of the definition of ρ , the definition turns out to have other interesting properties which we study below. The following alternative characterization of the rank of a continuous game is more concrete than the definition given above. This theorem simplifies the proofs of many rank-related results and will be applied to the problem of computing the rank of separable games in Subsection 3.3.

Theorem 3.4. *The rank ρ_i for player i in a continuous game is given by the smallest r_i such that there exist continuous functions $g_i^k : C_i \rightarrow \mathbb{R}$ and $h_{i,j}^k : C_{-i} \rightarrow \mathbb{R}$ which satisfy*

$$u_j(s) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{r_i} g_i^k(s_i) h_{i,j}^k(s_{-i}) \quad (4)$$

for all $s \in C$ and $j \neq i$ ($\rho_i = \infty$ if and only if no such representation exists). Furthermore, the minimum value of $r_i = \rho_i$ is achieved by functions $g_i^k(s_i)$ of the form $u_j(s_i, s_{-i})$ for some $s_{-i} \in C_{-i}$ and $j \neq i$ and functions $h_{i,j}^k(s_{-i})$ of the form $\int u_j(\cdot, s_{-i}) d\sigma_i$ for some $\sigma_i \in V_i$.

Proof. Throughout the proof we will automatically extend any functions $g_i^k : C_i \rightarrow \mathbb{R}$ to all of V_i in the canonical way. Suppose we are given a representation of the form (4). Let $g_i : C_i \rightarrow \mathbb{R}^{r_i}$ be defined by $g_i(s_i) = (g_i^1(s_i), \dots, g_i^{r_i}(s_i))$. By definition, ρ_i is the dimension of Δ_i/Y_i . Let Z_i denote the subspace of V_i parallel to Δ_i , i.e. the space of all signed measures σ_i such that $\int \sigma_i = 0$. Then $\rho_i = \dim Z_i / (Z_i \cap Y_i)$. By (4) any signed measure which is in Z_i and in $\ker g_i$ is almost payoff equivalent to the zero measure, so $Z_i \cap \ker g_i \subseteq Z_i \cap Y_i$ and therefore

$$\rho_i = \dim Z_i / (Z_i \cap Y_i) \leq \dim Z_i / (Z_i \cap \ker g_i) = \dim g_i(Z_i) \leq r_i.$$

It remains to show that if $\rho_i < \infty$ then there exists a representation of the form (4) with $r_i = \rho_i$. Recall that Y_i is defined to be

$$Y_i = \bigcap_{\substack{j \neq i \\ s_{-i} \in C_{-i}}} \ker u_j(\cdot, s_{-i})$$

where $u_j(\cdot, s_{-i})$ is interpreted as a linear functional on V_i . Since $\rho_i = \dim Z_i / (Z_i \cap Y_i)$, we can choose ρ_i linear functionals, call them $g_i^1, \dots, g_i^{\rho_i}$, from the collection of functionals whose intersection forms Y_i , such that $Z_i \cap Y_i = Z_i \cap \ker g_i$, where $g_i = (g_i^1, \dots, g_i^{\rho_i})$ as above. We cannot choose a smaller collection of linear functionals and achieve $Z_i \cap Y_i = Z_i \cap \ker g_i$, because $\rho_i = \dim Z_i / (Z_i \cap Y_i)$. Note that $Z_i \cap \ker g_i = \ker(1, g_i^1, \dots, g_i^{\rho_i})$ where 1 is the linear

functional $1(\sigma_i) = \int d\sigma_i$. Therefore no functional can be removed from the list $(1, g_i) = (1, g_i^1, \dots, g_i^{\rho_i})$ without affecting the kernel of the transformation $(1, g_i)$, so the functionals $1, g_i^1, \dots, g_i^{\rho_i}$ are linearly independent.

This means that any of the linear functionals $u_j(\cdot, s_{-i})$ (the intersection of whose kernels yields Y_i) can be written uniquely as a linear combination of the functionals $1, g_i^1, \dots, g_i^{\rho_i}$. That is to say, there are unique functions $h_{i,j}^k$ such that (4) holds with the functions g_i^k constructed here and $r_i = \rho_i$. The g_i^k are continuous by construction, so to complete the proof we must show that the functions $h_{i,j}^k$ are continuous as well. Since the functionals $1, g_i^1, \dots, g_i^{\rho_i}$ are linearly independent, we can choose a measure $\sigma_i^k \in V_i$ which makes the k^{th} of these functionals evaluate to unity and all the others zero. Substituting these values into (4) shows that $h_{i,j}^k(s_{-i}) = \int u_j(\cdot, s_{-i}) d\sigma_i^k$. Since u_j is continuous, $h_{i,j}^k$ is therefore also continuous. \square

Note that in the statement of Theorem 3.4 we have distinguished the component $h_{i,j}^0(s_{-i})$ in u_j . We have shown that this distinction follows from the definition of ρ_i , but there is also an intuitive game theoretic reason why this separation is natural. As mentioned above, ρ_i is intended to capture the number of essential degrees of freedom that player i has in his choice of strategy when playing a Nash equilibrium. Theorems 3.3 and 3.4 together show that player i only needs to take the other players' utilities into account to compute this number, and not his own. But player i is only concerned with the other players' utilities insofar as his own strategic choice affects them. The function $h_{i,j}^0(s_{-i})$ captures the part of player j 's utility which does not depend on player i 's strategy, so whether this function is zero or not it has no effect on the rank ρ_i .

Also note that while Theorem 3.4 gives decompositions of the utilities in terms of the ρ_i , it is not in general possible to summarize all these decompositions by writing the utilities in the form (1) with $m_i = \rho_i$ for all i . For a trivial counterexample, consider any game in which $u_i(s) = h_i^0(s_{-i})$ is independent of s_i for all i . Then Theorem 3.4 implies that $\rho_i = 0$ for all i , but the utilities are nonzero so we cannot take $m_i = 0$ in (1) for any i .

We close this subsection with an application. If a submatrix is formed from a matrix by "sampling," i.e. selecting a subset of the rows and columns, the rank of the submatrix is bounded by the rank of the original matrix. Theorem 3.4 shows that the same is true of continuous games, because a factorization of the form (4) for a game immediately provides a factorization for any smaller game produced by restricting the players' choices of strategies.

Corollary 3.5. *Let $(\{C_i\}, \{u_i\})$ be a continuous game with rank ρ and \tilde{C}_i be a nonempty compact subset of C_i for each i , with $\tilde{u}_i = u_i|_{\tilde{C}_i}$. Let $\tilde{\rho}$ denote the rank of the game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$. Then we have $\tilde{\rho}_i \leq \rho_i$ for all i .*

Definition 3.6. The game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$ in Corollary 3.5 is called a **sampled game** or a **sampled version** of $(\{C_i\}, \{u_i\})$.

Note that if we take \tilde{C}_i to be finite for each i , then the sampled game is a finite game. If the original game is separable and hence has finite rank, then Corollary 3.5 yields a uniform bound on the complexity of finite games which can arise from this game by sampling. This

fact is applied to the problem of computing approximate equilibria in Section 4 below. Finally, note that there are other kinds of bounds on the cardinality of the support of equilibria (e.g., for special classes of polynomial games as studied by Karlin [13]) which do not share this sampling property.

3.2 Characterizations of separable games

In this section we present a characterization theorem for separable games. We also provide an example that shows that the assumptions of this theorem cannot be weakened.

Theorem 3.7. *For a continuous game, the following are equivalent:*

1. *The game is separable.*
2. *The game has finite rank.*
3. *For each player i , every countably supported mixed strategy σ_i is almost payoff equivalent to a finitely supported mixed strategy τ_i with $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$.*

To prove that finite rank implies separability we repeatedly apply Theorem 3.4. The proof that the technical condition (3) implies (2) uses a linear algebraic argument to show that $\text{span } C_i/Y_i$ is finite dimensional and then a topological argument along the lines of the proof of Theorem 2.7 to show that V_i/Y_i is also finite dimensional.

After the proof of Theorem 3.7 we will give an explicit example of a game in which all mixed strategies are payoff equivalent to pure strategies, but for which the containment $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$ in condition (3) fails. In light of Theorem 3.7 this will show that the constructed game is nonseparable and that the containment $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$ cannot be dropped from condition (3), even if the other assumptions are strengthened.

Proof. (1 \Rightarrow 3) This was proven in Theorem 2.7.

(1 \Rightarrow 2) This follows from the proof of Theorem 2.7.

(2 \Rightarrow 1) We will prove this by induction on the number of players n . When $n = 1$ the statement is trivial and the case $n = 2$ follows immediately from Theorem 3.4. Suppose we have an n -player continuous game with $\rho_i < \infty$ for all i and that we have proven that $\rho_i < \infty$ for all i implies separability for $(n - 1)$ -player games. By fixing any signed measure $\sigma_n \in V_n$ we can form an $(n - 1)$ -player continuous game from the given game by removing the n^{th} player and integrating all payoffs of players $i < n$ with respect to σ_n , yielding a new game with payoffs $\tilde{u}_i(s_{-n}) = \int u_i(s_n, s_{-n}) d\sigma_n(s_n)$.

From the definition of Y_i , it is clear that $Y_i \subseteq \tilde{Y}_i$ for all $1 \leq i < n$. Therefore $\tilde{\rho}_i = \dim \Delta_i/\tilde{Y}_i \leq \dim \Delta_i/Y_i = \rho_i < \infty$ for $1 \leq i < n$ so the $(n - 1)$ -player game has finite rank. By the induction hypothesis, that means that the function \tilde{u}_1 is a separable function of the strategies s_1, \dots, s_{n-1} . Theorem 3.4 states that there exist continuous functions g_n^k and $h_{n,1}^k$ such that

$$u_1(s) = h_{n,1}^0(s_{-n}) + \sum_{k=1}^{\rho_n} g_n^k(s_n) h_{n,1}^k(s_{-n}) \quad (5)$$

where $h_{n,1}^k = \int u_1(s) d\sigma_n^k$ for some $\sigma_n^k \in V_n$. Therefore by choosing σ_n appropriately we can make $\tilde{u}_1 = h_{n,1}^k$ for any k , so $h_{n,1}^k(s_{-n})$ is a separable function of s_1, \dots, s_{n-1} for all k . By (5) u_1 is a separable function of s_1, \dots, s_n . The same argument works for all the u_i so the given game is separable and the general case is true by induction.

(3 \Rightarrow 2) Let $\psi_i : V_i \rightarrow V_i/Y_i$ be the canonical projection transformation. First we will prove that $\text{span } \psi_i(C_i)$ is finite dimensional. It suffices to prove that for every countable subset $\tilde{C}_i = \{s_i^1, s_i^2, \dots\} \subseteq C_i$, the set $\psi_i(\tilde{C}_i)$ is linearly dependent. Let $\{p^k\}$ be a sequence of positive reals summing to unity. Define the mixed strategy

$$\sigma_i = \sum_{k=1}^{\infty} p^k s_i^k.$$

By assumption there exists an M and $q^1, \dots, q^M \geq 0$ summing to unity such that

$$\psi_i(\sigma_i) = \psi_i \left(\sum_{k=1}^M q^k s_i^k \right) = \sum_{k=1}^M q^k \psi_i(s_i^k).$$

Let $\alpha = \sum_{k=M+1}^{\infty} p^k > 0$ and define the mixed strategy

$$\tau_i = \sum_{k=M+1}^{\infty} \frac{p^k}{\alpha} s_i^k.$$

Applying the assumption again shows that there exist N and r^{M+1}, \dots, r^N such that

$$\psi_i(\tau_i) = \psi_i \left(\sum_{k=M+1}^N r^k s_i^k \right) = \sum_{k=M+1}^N r^k \psi_i(s_i^k).$$

Therefore

$$\begin{aligned} \sum_{k=1}^M p^k \psi_i(s_i^k) &= \psi_i \left(\sum_{k=1}^M p^k s_i^k \right) = \psi_i(\sigma_i - \alpha \tau_i) = \psi_i(\sigma_i) - \alpha \psi_i(\tau_i) \\ &= \sum_{k=1}^M q^k \psi_i(s_i^k) - \sum_{k=M+1}^N \alpha r^k \psi_i(s_i^k), \end{aligned}$$

and rearranging terms shows that $\sum_{k=1}^M (p^k - q^k) \psi_i(s_i^k) + \sum_{k=M+1}^N \alpha r^k \psi_i(s_i^k) = 0$. Also $\sum_{k=1}^M (p^k - q^k) = -\alpha < 0$, so $p^k - q^k \neq 0$ for some k . Therefore $\psi_i(\tilde{C}_i)$ is linearly dependent, so $\text{span } \psi_i(C_i)$ is finite dimensional.

Since Y_i is closed, V_i/Y_i is a Hausdorff topological vector space under the quotient topology and ψ_i is continuous with respect to this topology [21]. Being finite dimensional, the subspace $\text{span } \psi_i(C_i) \subseteq V_i/Y_i$ is also closed [21]. Thus we have

$$V_i/Y_i = \psi_i(V_i) = \psi_i(\overline{\text{span } C_i}) \subseteq \overline{\psi_i(\text{span } C_i)} = \overline{\text{span } \psi_i(C_i)} = \text{span } \psi_i(C_i) \subseteq V_i/Y_i$$

where the first step is by definition, the second follows from Proposition 2.4, the next two are by continuity and linearity of ψ_i , and the final two are because $\text{span } \psi_i(C_i)$ is a closed subspace of V_i/Y_i . Therefore $\rho_i = \dim \Delta_i/Y_i \leq \dim V_i/Y_i = \dim \text{span } \psi_i(C_i) < \infty$. \square

The following counterexample shows that the containment $\text{supp } \tau_i \subset \text{supp } \sigma_i$ is a necessary part of condition 3 in Theorem 3.7 by showing that there exists a nonseparable continuous game in which every mixed strategy is payoff equivalent to a pure strategy.

Example 3.8. Consider a two-player game with $C_1 = C_2 = [0, 1]^\omega$, the set of all infinite sequences of reals in $[0, 1]$, which forms a compact metric space under the metric

$$d(x, x') = \sup_i \frac{|x_i - x'_i|}{i}.$$

Define the utilities

$$u_1(x, y) = u_2(x, y) = \sum_{i=1}^{\infty} 2^{-i} x_i y_i.$$

To show that this is a continuous game we must show that u_1 is continuous. Assume $d(x, x'), d(y, y') \leq \delta$. Then $|x_i - x'_i| \leq \delta i$ and $|y_i - y'_i| \leq \delta i$, so

$$\begin{aligned} |u_1(x, y) - u_1(x', y')| &= \left| \sum_{i=1}^{\infty} 2^{-i} (x_i y_i - x'_i y'_i) \right| \\ &= \left| \sum_{i=1}^{\infty} 2^{-i} (x_i y_i - x'_i y_i + x'_i y_i - x'_i y'_i) \right| \\ &\leq \sum_{i=1}^{\infty} 2^{-i} (y_i |x_i - x'_i| + x'_i |y_i - y'_i|) \\ &\leq \sum_{i=1}^{\infty} 2^{-i} (2\delta i) = \left(2 \sum_{i=1}^{\infty} 2^{-i} i \right) \delta. \end{aligned}$$

Thus $u_1 = u_2$ is continuous (in fact Lipschitz), making this a continuous game.

Let σ and τ be mixed strategies for the two players. By the Tonelli-Fubini theorem,

$$u_1(\sigma, \tau) = \int u_1 d(\sigma \times \tau) = \sum_{i=1}^{\infty} 2^{-i} \left(\int x_i d\sigma \right) \left(\int y_i d\tau \right).$$

Thus σ is payoff equivalent to the pure strategy $(\int x_1 d\sigma, \int x_2 d\sigma, \dots) \in C_1$ and similarly for τ , so this game has the property that every mixed strategy is payoff equivalent to a pure strategy.

Finally we will show that this game is nonseparable. Let $e^i \in C_1$ be the element having component i equal to unity and all other components zero. Let $\{p_i\}$ be a sequence of positive reals summing to unity and define the probability distribution $\sigma = \sum_{i=1}^{\infty} p_i e^i \in \Delta_1$. Suppose

σ were almost payoff equivalent to some mixture among finitely many of the e^i , call it $\tau = \sum_{i=1}^{\infty} q_i e^i$ where $q_i = 0$ for i greater than some fixed N . Let e_{N+1} be the strategy for player 2. Then the payoff if player 1 plays σ is

$$u_2(\sigma, e_{N+1}) = \int 2^{-(N+1)} x_{N+1} d\sigma = 2^{-(N+1)} p_{N+1}.$$

Similarly, if he chooses τ the payoff is $2^{-(N+1)} q_{N+1}$. Since $p_{N+1} > 0$ and $q_{N+1} = 0$, this contradicts the assumption that σ and τ are almost payoff equivalent. Thus condition 3 of Theorem 3.7 does not hold, so this game is not separable.

Therefore the condition that all mixed strategies be payoff equivalent to finitely supported strategies does not imply separability, even if a uniform bound on the size of the support is assumed. Hence the containment $\text{supp } \tau_i \subset \text{supp } \sigma_i$ cannot be removed from condition 3 of Theorem 3.7.

3.3 Computing the rank of a separable game

In this subsection we construct a formula for the rank of an arbitrary separable game and then specialize it to get formulas for the ranks of polynomial and finite games. For clarity of presentation we first prove a bound on the rank of a separable game which uses an argument that is similar to but simpler than the argument for the exact formula. While it is possible to prove all the results in this section directly from the definition $\rho_i = \dim \Delta_i / Y_i$, we will give proofs based on the alternative characterization in Theorem 3.4 because they are easier to understand and provide more insight into the structure of the problem.

Given a separable game in the standard form (1), construct a matrix $S_{i,j}$ for players i and j which has m_i columns and $\prod_{k \neq i} m_k$ rows and whose elements are defined as follows. Label each row with an $(n-1)$ -tuple $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_n)$ such that $1 \leq l_k \leq m_k$; the order of the rows is irrelevant. Label the columns $l_i = 1, \dots, m_i$. Each entry of the matrix then corresponds to an n -tuple (l_1, \dots, l_n) . The entry itself is given by the coefficient $a_j^{l_1 \dots l_n}$ in the utility function u_j .

Let $f_i(s_i)$ denote the column vector whose components are $f_i^1(s_i), \dots, f_i^{m_i}(s_i)$ and $f_{-i}(s_{-i})$ denote the row vector whose components are the products $f_1^{l_1}(s_1) \cdots f_{i-1}^{l_{i-1}}(s_{i-1}) f_{i+1}^{l_{i+1}}(s_{i+1}) \cdots f_n^{l_n}(s_n)$ ordered in the same way as the $(n-1)$ -tuples $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_n)$ were ordered above. Then $u_j(s) = f_{-i}(s_{-i}) S_{i,j} f_i(s_i)$.

Example 3.9. We introduce a new example game to clarify the subtleties of computing ranks when there are more than two players; we will return to Example 2.3 later. Consider the three player polynomial game with strategy spaces $C_1 = C_2 = C_3 = [-1, 1]$ and payoffs

$$\begin{aligned} u_1(x, y, z) &= 1 + 2x + 3x^2 + 2yz + 4xyz + 6x^2yz \\ &\quad + 3y^2z^2 + 6xy^2z^2 + 9x^2y^2z^2 \\ u_2(x, y, z) &= 7 + 2x + 3x^2 + 2y + 4xy + 6x^2y \\ &\quad + 3z^2 + 6xz^2 + 9x^2z^2 \\ u_3(x, y, z) &= -z - 2xz - 3x^2z - 2yz - 4xyz - 6x^2yz \\ &\quad - 3yz^2 - 6xyz^2 - 9x^2yz^2 \end{aligned} \tag{6}$$

where x , y , and z are the strategies of player 1, 2, and 3, respectively. Order the functions f_k^l so that $f_1(x) = [1 \ x \ x^2]^l$ and similarly for f_2 and f_3 with x replaced by y and z , respectively. If we wish to write down the matrices $S_{1,2}$ and $S_{1,3}$ we must choose an order for the pairwise products of the functions f_2^l and f_3^l . Here we will choose the order $f_{-1}(y, z) = [1 \ y \ y^2 \ z \ yz \ y^2z \ z^2 \ yz^2 \ y^2z^2]^l$. We can write down the desired matrices immediately from the given utilities.

$$S_{1,2} = \begin{bmatrix} 7 & 2 & 3 \\ 2 & 4 & 6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 3 & 6 & 9 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad S_{1,3} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1 & -2 & -3 \\ -2 & -4 & -6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ -3 & -6 & -9 \\ 0 & 0 & 0 \end{bmatrix}$$

This yields $u_2(x, y, z) = f_{-1}(y, z)S_{1,2}f_1(x)$ and $u_3(x, y, z) = f_{-1}(y, z)S_{1,3}f_1(x)$ as claimed.

Define S_i to be the matrix with m_i columns and $(n-1)\prod_{j \neq i} m_j$ rows which consists of all the matrices $S_{i,j}$ for $j \neq i$ stacked vertically (in any order). In the example above, S_1 would be the 18×3 matrix obtained by placing $S_{1,2}$ above $S_{1,3}$ on the page.

Theorem 3.10. *The rank of a separable game is bounded by $\rho_i \leq \text{rank } S_i$.*

Proof. Using any of a variety of matrix factorization techniques (e.g. the singular value decomposition), we can write S_i as

$$S_i = \sum_{k=1}^{\text{rank } S_i} v^k w^k$$

for some column vectors v^k and row vectors w^k . The vectors v^k will have length $(n-1)\prod_{j \neq i} m_j$ since that is the number of rows of S_i . Because of the definition of S_i , we can break each v^k into $n-1$ vectors of length $\prod_{j \neq i} m_j$, one for each player except i , and let v_j^k be the vector corresponding to player j . Then we have

$$S_{i,j} = \sum_{k=1}^{\text{rank } S_i} v_j^k w^k$$

for all $j \neq i$. Define the linear combinations $g_i^k(s_i) = w^k f_i(s_i)$ and $h_{i,j}^k = f_{-i}(s_{-i})v_j^k$, which are obviously continuous functions. Then

$$u_j(s) = f_{-i}(s_{-i})S_{i,j}f_i(s_i) = \sum_{k=1}^{\text{rank } S_i} g_i^k(s_i)h_{i,j}^k(s_{-i})$$

for all $s \in C$ and $j \neq i$, so $\rho_i \leq \text{rank } S_i$ by Theorem 3.4. □

Example 3.11. To demonstrate the power of the bound in Theorem 3.10 we will use it to give an immediate proof of Theorem 3.1. Consider any two-player finite game, where the first player chooses rows and the second player chooses columns. Let $C_i = \{1, \dots, m_i\}$ for $i = 1, 2$ and let R and C be the matrices of payoffs to the row and column players, respectively. We can then define $f_i^l(s_i)$ to be unity if $s_i = l$ and zero otherwise. This gives

$$\begin{aligned} u_1(s_1, s_2) &= f_1(s_1)' R f_2(s_2) \\ u_2(s_1, s_2) &= f_2(s_2) C' f_1(s_1) \end{aligned}$$

so $S_1 = C'$ and $S_2 = R$. Therefore by Theorem 3.10, $\rho_1 \leq \text{rank } S_1 = \text{rank } C$ and $\rho_2 \leq \text{rank } S_2 = \text{rank } R$. Substituting these bounds into Theorem 3.3 yields Theorem 3.1, so we have in fact generalized the results of Lipton et al. [16].

It is easy to see that there are cases in which the bound in Theorem 3.10 is not tight. For example, this will be the case (for generic coefficients $a_i^{j_1 \dots j_n}$) if $m_i \geq 2$ for each i and f_i^k is the same function for all k .

Fortunately we can use a technique similar to the one used above to compute ρ_i exactly instead of just computing a bound. To do so we need to write the utilities in a special form. First we add the new function $f_i^1(s_i) \equiv 1$ to the list of functions for player i appearing in the separable representation of the game if this function does not already appear, relabeling the other f_i^k as necessary. Next we consider the set of functions $\{f_j^k\}$ for each player j in turn and choose a maximal linearly independent subset. For players $j \neq i$ any such subset will do; for player i we must include the function which is identically unity in the chosen subset. Finally we rewrite the utilities in terms of these linearly independent sets of functions. This is possible because all of the f_j^k are linear combinations of those which appear in the maximal linearly independent sets.

From now on we will assume the utilities are in this form and that $f_i^1(s_i) \equiv 1$. Let $\bar{S}_{i,j}$ and \bar{S}_i be the matrices $S_{i,j}$ and S_i defined above, where the bar denotes the fact that we have put the utilities in this special form. Let T_i be the matrix \bar{S}_i with its first column removed. Note that this column corresponds to the function $f_i^1(s_i) \equiv 1$ which we have distinguished above, and therefore represents the components of the utilities of players $j \neq i$ which do not depend on player i 's choice of strategy. As mentioned in the note following Theorem 3.4, these components don't affect the rank. This is exactly the reason that we must remove the first column from \bar{S}_i in order to compute ρ_i . We will prove that $\rho_i = \text{rank } T_i$, but first we need a lemma.

Lemma 3.12. *If the functions $f_j^1(s_j), \dots, f_j^{m_j}(s_j)$ are linearly independent for all j , then the set of all $\prod_{j=1}^n m_j$ product functions of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$ is a linearly independent set.*

Proof. It suffices to prove this in the case $n = 2$, because the general case follows by induction. We prove the $n = 2$ case by contradiction. Suppose the set were linearly dependent. Then there would exist $\lambda_{k_1 k_2}$ not all zero such that

$$\sum_{k_1=1}^{m_1} \sum_{k_2=1}^{m_2} \lambda_{k_1 k_2} f_1^{k_1}(s_1) f_2^{k_2}(s_2) = 0 \quad (7)$$

for all $s \in C$. Choose l_1 and l_2 such that $\lambda_{l_1 l_2} \neq 0$. By the linear independence assumption there exists a finitely supported signed measure σ_2 such that $\int f_2^k d\sigma_2$ is unity for $k = l_2$ and zero otherwise. Integrating (7) with respect to σ_2 yields

$$\sum_{k_1=1}^{m_1} \lambda_{k_1 l_2} f_1^{k_1}(s_1) = 0,$$

contradicting the linear independence assumption for $f_1^1, \dots, f_1^{m_1}$. \square

Theorem 3.13. *If the representation of a separable game satisfies $f_i^1 \equiv 1$ and the set $\{f_j^1, \dots, f_j^{m_j}\}$ is linearly independent for all j then the rank of the game is $\rho_i = \text{rank } T_i$.*

Proof. The proof that $\rho_i \leq \text{rank } T_i$ follows essentially the same argument as the proof of Theorem 3.10. We use the singular value decomposition to write T_i as

$$T_i = \sum_{k=1}^{\text{rank } T_i} v^k w^k$$

for some column vectors v^k and row vectors w^k . The vectors v^k will have length $(n-1)\prod_{j \neq i} m_j$ since that is the number of rows of S_i . Let v^0 be the first column of \bar{S}_i , which was removed from \bar{S}_i to form T_i . Because of the definition of T_i and \bar{S}_i , we can break each v^k into $n-1$ vectors of length $\prod_{j \neq i} m_j$, one for each player except i , and let v_j^k be the vector corresponding to player j . Putting these definitions together we get

$$\bar{S}_{i,j} = v_j^0 [1 \ 0 \ \dots \ 0] + \sum_{k=1}^{\text{rank } T_i} v_j^k [0 \ w^k].$$

Define the linear combinations $g_i^k(s_i) = [0 \ w^k] f_i(s_i)$ and $h_{i,j}^k(s_{-i}) = f_{-i}(s_{-i}) v_j^k$, which are obviously continuous functions. Then

$$u_j(s) = f_{-i}(s_{-i}) \bar{S}_{i,j} f_i(s_i) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{\text{rank } T_i} g_i^k(s_i) h_{i,j}^k(s_{-i})$$

for all $s \in C$ and $j \neq i$, so $\rho_i \leq \text{rank } T_i$ by Theorem 3.4.

To prove the reverse inequality, choose continuous functions $g_i^k(s_i)$ and $h_{i,j}^k(s_{-i})$ such that

$$u_j(s) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{\rho_i} g_i^k(s_i) h_{i,j}^k(s_{-i})$$

holds for all $s \in C$ and $j \neq i$. By Theorem 3.4 we can choose these so that $g_i^k(s_i)$ is of the form $u_j(s_i, s_{-i})$ for some $s_{-i} \in C_{-i}, j \neq i$ and $h_{i,j}^k(s_{-i})$ is of the form $\int u_j(\cdot, s_{-i}) d\sigma_i$ for some $\sigma_i \in V_i$. Substituting these conditions into equation (1) defining the form of a separable

game shows that $g_i^k(s_i) = w^k f_i(s_i)$ for some row vectors w^k and $h_{i,j}^k = f_{-i}(s_{-i})v_j^k$ for some column vectors v_j^k . Define $w^0 = [1 \ 0 \ \dots \ 0]$. Then

$$u_j(s) = \sum_{k=0}^{\rho_i} f_{-i}(s_{-i})' v_j^k w^k f_i(s_i)$$

for all $s \in C$ and $j \neq i$.

This expresses $u_j(s)$ as a linear combination of products of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$. By assumption the sets $\{f_j^1, \dots, f_j^{m_j}\}$ are linearly independent for all j , and therefore the set of products of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$ is linearly independent by Lemma 3.12. Thus the expression of $u_j(s)$ as a linear combination of these products is unique.

But we also have $u_j(s) = f_{-i}(s_{-i})' \bar{S}_{i,j} f_i(s_i)$ by definition of $\bar{S}_{i,j}$, so uniqueness implies that $\bar{S}_{i,j} = \sum_{k=0}^{\rho_i} v_j^k w^k$. Let v^k be the vector of length $(n-1)\prod_{j \neq i} m_j$ formed by concatenating the v_j^k in the obvious way. Then $\bar{S}_i = \sum_{k=0}^{\rho_i} v^k w^k$. Let \tilde{w}^k be w^k with its first entry removed. By definition of T_i we have $T_i = \sum_{k=0}^{\rho_i} v^k \tilde{w}^k$. But w^0 is the standard unit vector with a 1 in the first coordinate, so \tilde{w}^0 is the zero vector and we may therefore remove the $k=0$ term from the sum. Thus $T_i = \sum_{k=1}^{\rho_i} v^k \tilde{w}^k$, which proves that $\text{rank } T_i \leq \rho_i$. \square

As corollaries of Theorem 3.13 we obtain formulas for the ranks of polynomial and finite games.

Corollary 3.14. *Consider a game with polynomial payoffs*

$$u_i(s) = \sum_{j_1=0}^{m_1-1} \cdots \sum_{j_n=0}^{m_n-1} a_i^{j_1 \cdots j_n} s_1^{j_1} \cdots s_n^{j_n} \quad (8)$$

and compact strategy sets $C_i \subset \mathbb{R}$ which satisfy the cardinality condition $|C_i| > m_i$ for all i . Then T_i is S_i with its first column removed and $\rho_i = \text{rank } T_i$.

Proof. Linear independence of the f_i^l follows from the cardinality condition and we have $f_i^0 \equiv 1$, so Theorem 3.13 applies. \square

Example 2.3 (cont'd). Applying this formula to the utilities in (2) shows that $\rho_1 = 1$ and $\rho_2 = 3$.

Example 3.9 (cont'd). Applying this formula to the utilities in (6) shows that in this case $\rho_1 = 1$ and $\rho_2 = \rho_3 = 2$.

Corollary 3.15. *Consider an n -player finite game with strategy sets $C_i = \{1, \dots, m_i\}$ and payoff $a_i^{s_1 \cdots s_n}$ to player i if the players play strategy profile (s_1, \dots, s_n) . The utilities can be written as*

$$u_i(s) = \sum_{j_1=1}^{m_1} \cdots \sum_{j_n=1}^{m_n} a_i^{j_1 \cdots j_n} f_1^{j_1}(s_1) \cdots f_n^{j_n}(s_n)$$

where $f_i^l(s_i)$ is unity if $s_i = l$ and zero otherwise. Let S_i be the matrix for player i as defined above and let c_1, \dots, c_{m_i} be the columns of S_i . Then we may take $T_i = [c_2 - c_1 \ \cdots \ c_{m_i} - c_1]$ and $\rho_i = \text{rank } T_i$.

Proof. If we replace f_i^1 with the function which is identically unity then the linear independence assumption on the f_k^l will still be satisfied, so we can apply Theorem 3.13. After this replacement, the coefficients in the new separable representation for the game are

$$\bar{a}_k^{j_1 \dots j_n} = \begin{cases} a_k^{j_1 \dots j_n} & \text{if } j_i = 1, \\ a_k^{j_1 \dots j_n} - a_k^{j_1 \dots j_{i-1} 1 j_{i+1} \dots j_n} & \text{if } j_i \neq 1. \end{cases}$$

Therefore if c_1, \dots, c_{m_i} are the columns of S_i from the original representation of the game we get $\bar{S}_i = [c_1 \quad c_2 - c_1 \quad \dots \quad c_{m_i} - c_1]$, so T_i is as claimed and an application of Theorem 3.13 completes the proof. \square

4 Computation of Nash Equilibria and Approximate Equilibria

In this section, we study computation of exact and approximate Nash equilibria. We first present an optimization formulation for the computation of (exact) Nash equilibria of general separable games. We show that for two-player polynomial games, this formulation has a biaffine objective function and linear matrix inequality constraints. We then present an algorithm for computing approximate equilibria of two-player separable games with infinite strategy sets which follows directly from the results on the rank of games given in Section 3 and compare it with known algorithms for finite games.

4.1 Computing Nash equilibria

The moments of an equilibrium can in principle be computed by nonlinear programming techniques using the following generalization of the Nash equilibrium formulation presented by Başar and Olsder [1]:

Proposition 4.1. *The following optimization problem has optimal value zero and the variables x in any optimal solution are the moments of a Nash equilibrium strategy profile with payoff p_i to player i :*

$$\begin{aligned} \max \quad & \sum_{i=1}^n [v_i(x) - p_i] \\ \text{s.t.} \quad & x_i \in \Delta_i / W_i = f_i(\Delta_i) \text{ for all } i \\ & v_i(f_i(s_i), x_{-i}) \leq p_i \text{ for all } i, \text{ all } s_i \in C_i \end{aligned}$$

The function f_i is the moment function defined in (3) and v_i is the payoff function on the moment spaces defined by $v_i(f_1(\sigma_1), \dots, f_n(\sigma_n)) = u_i(\sigma)$.

Proof. The constraints imply that $v_i(x) - p_i \leq 0$ for all i , so the objective function is bounded above by zero. Given any n -tuple of moments x which form a Nash equilibrium, let $p_i = v_i(x)$ for all i . Then the objective function evaluates to zero and all the constraints are satisfied, by definition of a Nash equilibrium. Therefore the optimal objective function value is zero and it is attained at all Nash equilibria.

Conversely suppose some feasible x and p achieve objective function value zero. Then the condition $v_i(x) - p_i \leq 0$ implies that $v_i(x) = p_i$ for all i . Also, the final constraint implies that player i cannot achieve a payoff of more than p_i by unilaterally changing his strategy. Therefore the moments x form a Nash equilibrium. \square

To compute equilibria by this method, we require an explicit description of the spaces of moments Δ_i/W_i . We also require a method for describing the payoff p_i to player i if he plays a best response to an m_{-i} -tuple of moments for the other players.

While it seems doubtful that such descriptions could be found for arbitrary f_i^j , they do exist for two-player polynomial games in which the pure strategy sets are intervals. In this case they can be written in terms of linear matrix inequalities as in Parrilo's treatment of the zero-sum case [19]. This yields a problem with biaffine objective and linear matrix inequality constraints.

Example 2.3 (cont'd). Directly solving this nonconvex problem with MATLAB's `fmincon` has proven error-prone, as there appear to be many local minima which are not global. However, we were able to compute the equilibrium measures

$$\begin{aligned}\sigma_1 &= 0.5532\delta(x + 1) + 0.4468\delta(x - 0.1149), \\ \sigma_2 &= \delta(y - 0.7166)\end{aligned}$$

(i.e. player 1 plays the pure strategy $x = -1$ with probability 0.5532 and so on) for the payoffs in (2) by this method.

4.2 Computing ϵ -equilibria

The difficulties in computing equilibria by general nonconvex optimization techniques suggest the need for more specialized systematic methods. As a step toward this, we present an algorithm for computing approximate Nash equilibria of two-player separable games. There are several possible definitions of approximate equilibrium, but here we will use:

Definition 4.2. A mixed strategy profile $\sigma \in \Delta$ is an ϵ -equilibrium ($\epsilon \geq 0$) if

$$u_i(s_i, \sigma_{-i}) \leq u_i(\sigma) + \epsilon$$

for all $s_i \in C_i$ and $i = 1, \dots, n$.

For $\epsilon = 0$, the definition of an ϵ -equilibrium reduces to that of a Nash equilibrium. We consider computing an ϵ -equilibrium of separable games that satisfy the following assumption:

Assumption 4.3.

- There are two players.
- The game is separable.

- The utilities can be evaluated efficiently.

To simplify the presentation of our algorithm, we also adopt the following assumption:

Assumption 4.4.

- The strategy spaces are $C_1 = C_2 = [-1, 1]$.
- The utility functions are Lipschitz.

In the description of the algorithm we will emphasize why Assumption 4.3 is needed for our analysis. After presenting the algorithm we will discuss how Assumption 4.4 could be relaxed.

Theorem 4.5. *For $\epsilon > 0$, the following algorithm computes an ϵ -equilibrium of a game of rank ρ satisfying Assumptions 4.3 and 4.4 in time polynomial in $\frac{1}{\epsilon}$ for fixed ρ and time polynomial in the components of ρ for fixed ϵ (for the purposes of asymptotic analysis of the algorithm with respect to ρ the Lipschitz condition is assumed to be satisfied uniformly by the entire class of games under consideration).*

Algorithm 4.6. By the Lipschitz assumption there are real numbers L_1 and L_2 such that

$$|u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})| \leq L_i |s_i - s'_i|$$

for all $s_{-i} \in C_{-i}$ and $i = 1, 2$. Clearly this is equivalent to requiring the same inequality for all $\sigma_{-i} \in \Delta_{-i}$. Divide the interval C_i into equal subintervals of length no more than $2\frac{\epsilon}{L_i}$; at most $\lceil \frac{L_i}{\epsilon} \rceil$ such intervals are required. Let \tilde{C}_i be the set of center points of these intervals, and construct a finite sampled game by restricting the strategy sets to the \tilde{C}_i . Call the resulting payoff matrices U_1 and U_2 . Compute a Nash equilibrium of the sampled game. To do so, iterate over all pairs of nonempty subsets $S_1 \subseteq \tilde{C}_1$ and $S_2 \subseteq \tilde{C}_2$ such that the cardinality of S_i is at most $\rho_i + 1$ for $i = 1, 2$. Let x_1 and x_2 be probability vectors indexed by the elements of \tilde{C}_1 and \tilde{C}_2 , respectively. For each such pair (S_1, S_2) we use the fact that linear programs are polynomial-time solvable to find x_1 and x_2 such that the following linear constraints are satisfied, or prove that no such vectors exist [3].

$$\begin{aligned}
[x_1 U_2]_{s_2} &\geq [x_1 U_2]_{t_2} \text{ for all } s_2 \in S_2, t_2 \in \tilde{C}_2 \\
[U_1 x_2]_{s_1} &\geq [U_1 x_2]_{t_1} \text{ for all } s_1 \in S_1, t_1 \in \tilde{C}_1 \\
x_i(s_i) &\geq 0 \text{ for all } s_i \in S_i, i = 1, 2 \\
x_i(s_i) &= 0 \text{ for all } s_i \in \tilde{C}_i \setminus S_i, i = 1, 2 \\
\sum_{s_i \in S_i} x_i(s_i) &= 1 \text{ for } i = 1, 2
\end{aligned} \tag{9}$$

(There are many redundant constraints here which could easily be removed, but we have presented the constraints in this form for simplicity.) Any feasible point for any pair (S_1, S_2) is a Nash equilibrium of the sampled game and an ϵ -equilibrium of the original game. The algorithm will find at least one such point.

Proof. For the purpose of analyzing the complexity of the algorithm we will view the Lipschitz constants as fixed, even as ρ varies. Let \tilde{u}_i be the payoffs of the sampled game and suppose σ is a Nash equilibrium of the sampled game. Choose any $s_i \in C_i$ and let s'_i be an element of \tilde{C}_i closest to s_i , so $|s_i - s'_i| \leq \frac{\epsilon}{L_i}$. Then

$$\begin{aligned} & u_i(s_i, \sigma_{-i}) - u_i(\sigma) \\ & \leq u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i}) + u_i(s'_i, \sigma_{-i}) - u_i(\sigma) \\ & \leq |u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i})| + \tilde{u}_i(s'_i, \sigma_{-i}) - \tilde{u}_i(\sigma) \\ & \leq L_i \frac{\epsilon}{L_i} + 0 = \epsilon \end{aligned}$$

so σ is automatically an ϵ -equilibrium of the original separable game. Thus it will suffice to compute a Nash equilibrium of the finite sampled game.

To do so, first compute or bound the rank ρ of the original separable game using Theorem 3.13 or 3.10. By Theorem 3.3 and Corollary 3.5, the sampled game has a Nash equilibrium in which player i mixes among at most $\rho_i + 1$ pure strategies, independent of how large $|\tilde{C}_i|$ is. The separability assumption is fundamental because without it we would not obtain this uniform bound independent of $|\tilde{C}_i|$. The number of possible choices of at most $\rho_i + 1$ pure strategies from \tilde{C}_i is

$$\sum_{k=1}^{\rho_i+1} \binom{|\tilde{C}_i|}{k} \leq \binom{|\tilde{C}_i| + \rho_i}{1 + \rho_i} = \binom{|\tilde{C}_i| + \rho_i}{|\tilde{C}_i| - 1},$$

which is a polynomial in $|\tilde{C}_i| \propto \frac{1}{\epsilon}$ for fixed ρ and a polynomial in the components of ρ for fixed ϵ . This leaves the step of checking whether there exists an equilibrium σ for a given choice of $S_i = \text{supp}(\sigma_i) \subseteq \tilde{C}_i$ with $|S_i| \leq \rho_i + 1$ for each i , and if so, computing such an equilibrium. Since the game has two players, the set of such equilibria for given supports is described by a number of linear equations and inequalities which is polynomial in $\frac{1}{\epsilon}$ for fixed ρ and polynomial in the components of ρ for fixed ϵ ; these equations and inequalities are given by (9). Since linear programs are polynomial-time solvable, we can find a feasible solution to such inequalities or prove infeasibility in polynomial time. The two player assumption is key at this step, because with more players the constraints would fail to be linear or convex and we could no longer use a polynomial time linear programming algorithm.

Thus we can check all supports and find an ϵ -equilibrium of the sampled game in polynomial time as claimed. \square

We will now consider weakening Assumption 4.4. The Lipschitz condition could be weakened to a Hölder condition and the same proof would work, but it seems that we must require some quantitative bound on the speed of variation of the utilities in order to bound the running time of the algorithm. Also, the strategy space could be changed to any compact set which can be efficiently sampled, e.g. a box in \mathbb{R}^n . However, for the purpose of asymptotic analysis of the algorithm, the proof here only goes through when the Lipschitz constants and strategy space are fixed. A more complex analysis would be required if the strategy space were allowed to vary with ρ , for example.

It should be noted that the requirement that the strategy space be fixed for asymptotic analysis means that Theorem 4.5 does not apply to finite games, at least not if the number of strategies is allowed to vary. For the sake of comparison and completeness we state the best known ϵ -equilibrium algorithm for finite games below.

Theorem 4.7 (Lipton et al. [16]). *There exists an algorithm to compute an ϵ -equilibrium of an m -player finite game with n strategies per player which is polynomial in $\frac{1}{\epsilon}$ for fixed m and n , polynomial in m for fixed n and ϵ , and quasipolynomial in n for fixed ϵ and m (assuming the payoffs of the games are uniformly bounded).*

In the case of two-player separable games which we have considered, the complexity of the payoffs is captured by ρ , which is bounded by the cardinality of the strategy spaces in two-player finite games. Therefore in finite games the complexity of the payoffs and the complexity of the strategy spaces are intertwined, whereas in games with infinite strategy spaces they are decoupled. The best known algorithm for finite games stated in Theorem 4.7 has quasipolynomial dependence on the complexity of the game. Our algorithm is interesting because it has polynomial dependence on the complexity of the payoffs when the strategy spaces are held fixed. In finite games this type of asymptotic analysis is not possible due to the coupling between the two notions of complexity of a game, so a direct comparison between Theorem 4.5 and Theorem 4.7 cannot be made.

5 Conclusions

We have shown that separable games provide a natural setting for the study of games with payoffs satisfying a low-rank condition. This level of abstraction allows the low-rank results of Lipton et al. [16] to be extended to infinite strategy spaces and multiple players. Since the rank of a separable game gives a bound on the cardinality of the supports of equilibria for any sampled version of that separable game, approximate equilibria can be computed in time polynomial in $\frac{1}{\epsilon}$ by discretizing the strategy spaces and applying standard computational techniques for low-rank games.

Other types of low-rank conditions have been studied for finite games, for example Kannan and Theobald have considered the condition that the sum of the payoff matrices be low-rank [12]. It is likely that that the discretization techniques used here can be applied in an analogous way to yield results about computing approximate equilibria of continuous games when the sum of the payoffs of the players is a separable function.

There also exist many computational techniques for finite games which do not make low-rank assumptions. It may be possible to extend some of these techniques directly to separable games to yield algorithms for computing exact equilibria of separable games. Such an extension would likely require an explicit description of the moment spaces in terms of inequalities rather than the description given above as the convex hull of the set of moments due to pure strategies. In the case of two-player polynomial games, such an explicit description is known to be possible using linear matrix inequalities and has been applied to zero-sum polynomial games by Parrilo [19]. While the lack of polyhedral structure in the

moment spaces would most likely prohibit the use of a Lemke-Howson type algorithm, a variety of other finite game algorithms may be extendable to this setting; see McKelvey and McLennan for a survey of such algorithms [17].

Finally, there exist a variety of other solution concepts for strategic form games which may be amenable to analysis and computation in the case of separable games, and in particular in polynomial games. Preliminary results on computation of correlated equilibria appear in [23, 24]. For a correlated equilibrium version of the rank bounds on Nash equilibria of separable games presented above, see [23]. We leave the extension to other solution concepts, in particular iterated elimination of strictly dominated strategies, for future work.

References

- [1] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999.
- [2] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, 2003.
- [3] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA, 1997.
- [4] X. Chen and X. Deng. 3-Nash is PPAD-complete. *Electronic Colloquium on Computational Complexity*, TR05-134, 2005.
- [5] X. Chen and X. Deng. Settling the complexity of 2-player Nash equilibrium. *Electronic Colloquium on Computational Complexity*, TR05-140, 2005.
- [6] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th annual ACM symposium on theory of computing*, pages 71 – 78, New York, NY, 2006. ACM Press.
- [7] C. Daskalakis and C. H. Papadimitriou. Three-player games are hard. *Electronic Colloquium on Computational Complexity*, TR05-139, 2005.
- [8] M. Drescher and S. Karlin. Solutions of convex games as fixed points. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, number 28 in Annals of Mathematics Studies, pages 75 – 86. Princeton University Press, Princeton, NJ, 1953.
- [9] M. Drescher, S. Karlin, and L. S. Shapley. Polynomial games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games I*, number 24 in Annals of Mathematics Studies, pages 161 – 180. Princeton University Press, Princeton, NJ, 1950.
- [10] I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170 – 174, February 1952.
- [11] O. Gross. A rational payoff characterization of the Cantor distribution. Technical Report D-1349, The RAND Corporation, 1952.
- [12] R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. [arXiv:cs.GT/0511021](https://arxiv.org/abs/cs.GT/0511021), 2005.
- [13] S. Karlin. *Mathematical Methods and Theory in Games, Programming, and Economics*, volume 2: Theory of Infinite Games. Addison-Wesley, Reading, MA, 1959.

- [14] S. Karlin and L. S. Shapley. *Geometry of Moment Spaces*. American Mathematical Society, Providence, RI, 1953.
- [15] C. E. Lemke and J. T. Howson, Jr. Equilibrium points in bimatrix games. *SIAM Journal on Applied Math*, 12:413 – 423, 1964.
- [16] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 36 – 41, New York, NY, 2003. ACM Press.
- [17] R.D. McKelvey and A. McLennan. Computation of equilibria in finite games. In H. M. Amman, D. A. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, volume 1, pages 87 – 142. Elsevier, Amsterdam, 1996.
- [18] C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498 – 532, June 1994.
- [19] P. A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control (CDC)*, 2006.
- [20] K. R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, New York, NY, 1967.
- [21] W. Rudin. *Functional Analysis*. McGraw-Hill, New York, 1991.
- [22] H. Scarf. The approximation of fixed points of a continuous mapping. *SIAM Journal of Applied Mathematics*, 15:1328 – 1343, 1967.
- [23] N. D. Stein. Characterization and computation of equilibria in infinite games. Master’s thesis, Massachusetts Institute of Technology, May 2007.
- [24] N. D. Stein, P. A. Parrilo, and A. Ozdaglar. Characterization and computation of correlated equilibria in infinite games. Submitted to the 46th IEEE Conference on Decision and Control (CDC), 2007.
- [25] S. Vavasis. Approximation algorithms for indefinite quadratic programming. *Mathematical Programming*, 57:279 – 311, 1992.