Separable and Low-Rank Continuous Games

Noah D. Stein, Asuman Ozdaglar, and Pablo A. Parrilo

Abstract— Separable games are a structured subclass of continuous games whose payoffs take a sum-of-products form; the zero-sum case has been studied in earlier work. Included in this subclass are all finite games and polynomial games. Separable games provide a unified framework for analyzing and generating results about the structural properties of low rank games. This work extends previous results on low-rank finite games by allowing for multiple players and a broader class of payoff functions. We also discuss computation of exact and approximate equilibria in separable games. Finally we tie these results together with alternative characterizations of separability which show that separable games are the largest class of continuous games to which low-rank arguments apply.

I. INTRODUCTION

There has been much interest recently in the problem of computing mixed strategy Nash equilibria of games. The hardness of this problem for cases as simple as two-player finite games has been established by Chen and Deng [3], and this suggests the need to search for classes of games which are computationally tractable. Many such classes have been identified, but in this paper we will consider a generalization of the low-rank games studied by Lipton et al. [11]. Here the computational simplification arises because it can be proven that in an equilibrium, each player chooses a "simple strategy" without loss of generality, where simple refers to the fact that the players randomize among a small number of pure strategies with positive probability.

In games with infinite strategy sets, additional complications arise in computing equilibria. Without compactness of the strategy spaces and continuity of the payoff functions, an equilibrium may fail to exist at all. Even when these assumptions are made, as in so-called continuous games, the strategies played in equilibrium may be arbitrarily complicated. Specifically, two-player zero-sum games with compact intervals for strategy sets and rational payoff functions can be constructed in which the unique equilibrium strategies are the Cantor measure or worse, see Karlin [9] for an example. To compute equilibria of continuous games, we therefore require some additional structure which will ensure that an equilibrium exists which can be succinctly represented.

There are several possible structures which could be imposed. In this paper we study a separable structure in which the payoffs to each player can be written as a weighted sum of products of functions in each player's strategy separately, e.g. as polynomials. Zero-sum games possessing

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

such a separable structure have been studied by Karlin [9]. Efficient algorithms for computing equilibria of games with polynomial payoffs have been given by Parrilo [14]. Since continuous functions can be uniformly approximated by polynomials, a large class of games can be exactly or approximately represented as separable games.

Another advantage of the separable structure is that for games of this type, the complexity of the payoffs gives an immediate bound on the complexity of the strategies played in equilibrium. If the payoff functions are viewed as matrices whose indices range over compact metric spaces rather than merely finite sets, then the separable form of the payoffs can be viewed as a "finite rank" condition. Since a generic continuous function has "infinite rank," separable games are a natural generalization of games with low-rank payoff matrices. In this more general setting we extend the results of [11] on Nash equilibria of low-rank two-player finite games, allowing for an arbitrary finite number of players as well as infinite strategy sets and a broader class of payoff functions. We also extend the correlated equilibrium results in [5] to this setting.

The rest of this paper is organized as follows. In Section II, we present the basic definitions and theorems as well as introduce a running example. Then in Section III we define the rank of a continuous game, present a theorem demonstrating the relevance of this definition to the cardinality of the support of Nash equilibria, and show how to bound the rank for arbitrary separable games and compute it exactly for finite and polynomial games. In Section IV we discuss computation of Nash equilibria and approximate equilibria. Section V contains 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 close with conclusions and directions for future work.

II. BASIC THEORY

Some notational conventions used throughout are that subscripts refer to players, while superscripts are reserved for other indices, rather than exponents. If S_j are sets for $j = 1, \ldots, n$ then $S = \Pi_{j=1}^n S_j$ and $S_{-i} = \Pi_{j \neq i} S_j$. The n-tuple s and the $(n-1)$-tuple s_{-i} are formed from the points s_j similarly. We use the symbols $\operatorname{span} S$, $\operatorname{aff} S$, $\operatorname{conv} S$, and $\overline{S}$ to denote the span, affine hull, convex hull, and closure of the set S, respectively.

Definition 2.1: An n-player **continuous game** is defined by n **pure strategy** spaces C_i assumed to be nonempty compact metric spaces and n **utility** or **payoff functions** $u_i : C \to \mathbb{R}$ assumed to be continuous. 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

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

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

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

where $a_i^{j_1 \cdots j_n} \in \mathbb{R}$ and the $f_i^j : C_i \to \mathbb{R}$ are continuous. In the special case when the C_i are subsets of $\mathbb{R}$ and the f_i^j are polynomials in s_i, the game is called a **polynomial game**.

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 1: 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

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

This is a polynomial game, and we will return to it periodically to apply the results presented.

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. Two natural topoligies on V_i which will play a role are the weak* and the strong (or norm) topology. The former is the weakest topology such that whenever $f : C_i \to \mathbb{R}$ is continuous, $\sigma \mapsto \int f d\sigma$ is a continuous linear functional on V_i. The latter is the topology arising from the total variation norm.

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 σ. In the case of polynomial games when $f_i^j(s_i) = s_i^j$, 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 $\mathrm{conv}\, C_i$ and $\mathrm{span}\, C_i$ 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 Ch. 2 of Parthasarathy [15] for the proofs).

Proposition 2.3:
- The sets C_i and Δ_i are weak* compact.
- The weak* closures of $\mathrm{conv}\, C_i$ and $\mathrm{span}\, C_i$ are Δ_i and V_i, respectively.
- The norm topology makes V_i a Banach space.

The simplifications which occur in separable games as opposed to general continuous games can be expressed in terms of the following three notions of equivalence between two measures:

Definition 2.4: 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 2.5: Let
- $W_i = \{\text{measures moment equivalent to } 0\}$
- $X_i = \{\text{measures payoff equivalent to } 0\}$
- $Y_i = \{\text{measures almost payoff equivalent to } 0\}$

where 0 denotes the zero measure in V_i.

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) weak* continuous linear functionals, hence they are weak* closed, and therefore closed under the stronger norm topology as well.

We will analyze separable games by considering the quotients of V_i by these subspaces, i.e. V_i 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. The following theorem presents the most fundamental result about separable games.

Theorem 2.6: In a separable game every mixed strategy σ_i is moment equivalent to a finitely-supported mixed strategy τ_i with $|supp(\tau_i)| \leq m_i + 1$. Moreover, if σ_i is countably-supported τ_i can be chosen with $\mathrm{supp}(\tau_i) \subset \mathrm{supp}(\sigma_i)$.

Proof: Note that the map

$$f_i : \sigma_i \mapsto \left(f_i^1(\sigma_i), \ldots, f_i^{m_i}(\sigma_i)\right)$$

is linear and weak* continuous, with kernel W_i. Thus V_i/W_i has dimension at most m_i. Then

$$\begin{aligned} f_i(\Delta_i) = f_i\left(\overline{\mathrm{conv}\, C_i}\right) &\subseteq \overline{f_i(\mathrm{conv}\, C_i)} \\ &= \overline{\mathrm{conv}\, f_i(C_i)} = \mathrm{conv}\, f_i(C_i). \end{aligned} \quad (3)$$

The first three steps follow from Proposition 2.3, continuity of f_i, and linearity of f_i, respectively. The final equality follows from the fact that $\mathrm{conv}\, f_i(C_i)$ is compact, being the convex hull of a compact subset of a finite-dimensional space. The reverse inclusion is obvious, so we have $f_i(\Delta_i) = $

Fig. 1. The space $f_i(\Delta_i) \cong \Delta_i/W_i$ of possible moments for either player 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.

conv $f_i(C_i) = f_i(\text{conv }C_i)$. Thus any mixed strategy is moment equivalent to a finitely-supported mixed strategy, and applying Carathéodory's theorem to the set 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. ∎

This theorem can also be proven by a separating hyperplane argument as applied to zero-sum games by Karlin [9]. Combining this theorem with Glicksberg's result [6] that every continuous game has an equilibrium yields:

Corollary 2.7: Any separable game has a Nash equilibrium in which player i mixes among at most $m_i + 1$ pure strategies.

Example 1 (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 2.6 is shown in Figure 1. In this case it is the same for both players.

For each player the range of the indices in (1) is $0 \leq j_i \leq 3$, so by Corollary 2.7, 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, there is extra structure here to be exploited. For example, u_2 depends on the expected value $E_{\sigma_1}[x^2]$, but not on $E_{\sigma_1}[x]$ or $E_{\sigma_1}[x^3]$. In particular, player 2 is indifferent between the two strategies $\pm x$ of player 1 for all x, insofar as this choice does not affect his payoff (though it does affect what strategy profiles are equilbria). In the following section, we will show how to give improved bounds on the number of strategies played in equilibrium which take these simplifications into account in a systematic manner.

III. THE RANK OF A CONTINUOUS GAME

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

rank precisely, use it to give a bound on the cardinality of the support of equilibrium strategies, and show how to compute the rank of finite and polynomial games. This will generalize the results by Lipton et al. [11] on low-rank two-player finite games to multiplayer finite and polynomial games, simultaneously improving these results slightly. The primary tool will be the notion of almost payoff equivalence from Definition 2.4. In what follows, the dimension of a set will refer to the dimension of its affine hull.

Definition 3.1: The **rank** of a continuous game is defined to be $\rho = (\rho_1, \ldots, \rho_n)$ 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, it is clear that separability implies finite rank. In Section V we show that the converse is also true. Using the rank of a game, Corollary 2.7 can now be improved as follows:

Theorem 3.2: 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. Furthermore, all strategy profiles in the box conv$\{\sigma_1, \tau_1\} \times \cdots \times$ conv$\{\sigma_n, \tau_n\}$ formed by σ and τ are equilibria with payoff $u_i(\sigma)$ to player i.

Proof: By Theorem 2.6, we can assume without loss of generality that each player's mixed strategy σ_i is finitely supported. Fix i, let $\psi_i : V_i \to V_i/Y_i$ 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 but perhaps 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. ∎

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. The same is true of continuous games.

Proposition 3.3: 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 $\tilde{u}_i = u_i|_{\tilde{C}}$. Then the game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$ satisfies $\tilde{\rho}_i \leq \rho_i$ for all i.

Proof: Since $\tilde{\Delta}_i \subseteq \Delta_i$ and $Y_i \cap \tilde{V}_i \subseteq \tilde{Y}_i$,

$$\tilde{\rho}_i = \dim \tilde{\Delta}_i / \tilde{Y}_i \leq \dim \tilde{\Delta}_i / (Y_i \cap \tilde{V}_i)$$
$$= \dim \tilde{\Delta}_i / Y_i \leq \dim \Delta_i / Y_i = \rho_i.$$

∎

Definition 3.4: The game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$ in Proposition 3.3 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 Proposition 3.3 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 IV below. Note that while the proof of Proposition 3.3 is trivial, there exist other kinds of bounds on the cardinality of the support of equilibria (e.g. for special classes of polynomial games as studied by Karlin [9]) which do not share this sampling property.

With the significance of rank made clear by Theorem 3.2 and Proposition 3.3, we will now present a bound on the rank for arbitrary separable games which is tight in the case of finite and polynomial games. For consistency while deriving these results we will use the convention that the summations in (1) begin at zero, even for non-polynomial games. For finite games we will assume each player's strategy set is a set of the form $C_i = \{0, 1, \ldots, m_i\}$ and let $f_i^{j_i}(s_i) = \delta(j_i - s_i)$ where δ is the Kronecker delta function. We will assume that polynomial games are written so that $f_i^j(s_i) = s_i^j$. This way the coefficients in (1) are just the payoffs for finite games and the coefficients of the polynomials for polynomial games.

Despite the fact that V_i may be infinite-dimensional, the problem of computing the rank of a separable game can be reduced to computing the dimension of certain finite-dimensional convex sets. To show this, we first require several preliminary definitions and a lemma. Let $f_i : V_i \to M_i$, where $M_i = \mathbb{R}^{m_i+1} \supseteq V_i / W_i$, be the generalized moment function defined by

$$f_i(\sigma_i) = (f_i^0(\sigma_i), \ldots, f_i^{m_i}(\sigma_i)).$$

We write $x_i = (x_i^0, \ldots, x_i^{m_i})$ for a typical element of M_i and we define multilinear functionals $v_i : M \to \mathbb{R}$ by

$$v_i(x) = \sum_{j_1=0}^{m_1} \cdots \sum_{j_n=0}^{m_n} a_i^{j_1 \cdots j_n} x_1^{j_1} \cdots x_n^{j_n}, \qquad (4)$$

so $u_i(\sigma) = v_i(f_1(\sigma_1), \ldots, f_n(\sigma_n))$. Let $e_i^0, \ldots, e_i^{m_i}$ denote the standard unit vectors in M_i.

Lemma 3.5: Consider the linear functionals on V_i of the form $v_l(f_i(\cdot), x_{-i})$ for $1 \leq l, i \leq n$ where $l \neq i$ and $x_{-i} \in$

M_{-i}. Then

$$Y_i \supseteq \bigcap_{\substack{x_k \in \{e_k^0, \ldots, e_k^{m_k}\} \\ k,l \neq i}} \ker v_l(f_i(\cdot), x_{-i}) \qquad (5)$$

with equality holding for all i if span $f_i(C_i) = M_i$ for all i.

Proof: The definition of Y_i can be written as

$$Y_i = \bigcap_{\substack{s_k \in C_k \\ k,l \neq i}} \ker u_l(\cdot, s_{-i}).$$

Letting $x_k = f_k(s_k)$ and applying (4), we obtain

$$Y_i = \bigcap_{\substack{x_k \in f_k(C_k) \\ k,l \neq i}} \ker v_l(f_i(\cdot), x_{-i}). \qquad (6)$$

Without changing the set on the right hand side, we can take the intersection over the kernels of a larger class of linear functionals, those which are linear combinations of the functionals considered in (6). Using the multilinearity of the v_l yields

$$Y_i = \bigcap_{\substack{x_k \in \text{span } f_k(C_k) \\ k,l \neq i}} \ker v_l(f_i(\cdot), x_{-i}). \qquad (7)$$

If we replace span $f_k(C_k)$ by M_k we will be taking the intersection over a larger collection of sets, so

$$Y_i \supseteq \bigcap_{\substack{x_k \in M_k \\ k,l \neq i}} \ker v_l(f_i(\cdot), x_{-i})$$

with equality if span $f_k(C_k) = M_k$ for all $k \neq i$. Reversing the procedure used to pass from (6) to (7), we may replace the collection of linear functionals on the right hand side by a spanning set. Letting x_k range over the standard unit vectors in M_k for all $k \neq i$ yields such a spanning set (by the multilinearity of the v_l), which proves the lemma. ∎

The linear functional $v_l(f_i(\cdot), x_{-i})$ appearing in (5) can be written as the composition $v_l(\cdot, x_{-i}) \circ f_i$. Since x_{-i} is of the form

$$x_{-i} = \left(e_1^{j_1}, \ldots, e_{i-1}^{j_{i-1}}, e_{i+1}^{j_{i+1}}, \ldots, e_n^{j_n}\right),$$

the matrix for the linear functional $v_l(\cdot, x_{-i})$ is the row vector consisting of the $m_i + 1$ coefficients $a_l^{j_1 \cdots j_n}$ as j_i ranges from 0 to m_i.

Definition 3.6: Let S_i denote the matrix composed of all such row vectors as j_k ranges from 0 to m_k and $k, l \neq i$; the order of these vectors is irrelevant. The matrix S_i has $(n-1)\Pi_{k \neq i}(m_k+1)$ rows and m_i+1 columns. For a set $X \subseteq M_i$ we will write $S_i \cdot X$ to denote the image set $\{S_i x | x \in X\}$.

Theorem 3.7: The rank is bounded by

$$\rho_i \leq \dim [S_i \cdot f_i(C_i)]$$

with equality for all i if span $f_i(C_i) = M_i$ for all i. In particular, equality holds for arbitrary finite games and for polynomial games which satisfy $|C_i| > m_i$.

Proof: Given S_i, we can rephrase Lemma 3.5 as: $Y_i \supseteq \ker[S_i \cdot f_i]$ with equality for all i when $\operatorname{span} f_i(C_i) = M_i$ for all i. Using the definition of rank we have

$$\begin{aligned}
\rho_i &= \dim \Delta_i / Y_i \leq \dim \Delta_i / (\ker[S_i \cdot f_i]) \\
&= \dim[S_i \cdot f_i(\Delta_i)] = \dim[S_i \cdot \operatorname{conv} f_i(C_i)] \\
&= \dim[S_i \cdot f_i(C_i)]
\end{aligned}$$

with equality for all i if $\operatorname{span} f_i(C_i) = M_i$ for all i. The third equality follows from (3) and the final equality follows from the definition of the dimension of a set as the dimension of its affine hull.

In the case of a finite game, we have

$$\begin{aligned}
\operatorname{aff} f_i(C_i) &= \operatorname{aff} f_i(\{0, \ldots, m_i\}) = \operatorname{aff}\{e_i^0, \ldots, e_i^{m_i}\} \\
&= \left\{ (x_i^0, \ldots, x_i^{m_i}) \in \mathbb{R}^{m_i+1} \Big| \sum_{k=0}^{m_i} x_i^k = 1 \right\}.
\end{aligned} \quad (8)$$

In the case of a polynomial game

$$\begin{aligned}
\operatorname{aff} f_i(C_i) &= \operatorname{aff}\{(1, x, x^2, \ldots, x^{m_i}) | x \in C_i\} \\
&= \operatorname{aff}\{e_i^0, e_i^0 + e_i^1, e_i^0 + e_i^2, \ldots, e_i^0 + e_i^{m_i}\} \quad (9) \\
&= \{(x_i^0, \ldots, x_i^{m_i}) \in \mathbb{R}^{m_i+1} | x_i^0 = 1\},
\end{aligned}$$

where we assume that $|C_i| > m_i$ to get the penultimate equality, which is a standard result. Note that in both the finite and polynomial cases we have shown that $\operatorname{aff} f_i(C_i)$ is a hyperplane which does not pass through the origin, so $\operatorname{span} f_i(C_i) = M_i$. ∎

Let F_i be the matrix with one fewer column than S_i whose j^{th} column is the j^{th} column of S_i minus the 0^{th} column of S_i for $1 \leq j \leq m_i$. Let P_i be S_i with the 0^{th} column removed. In view of (8) and (9) we have proven the following.

Proposition 3.8: The rank can be computed or bounded as follows. For any:

- Separable game, $\rho_i \leq \operatorname{rank} S_i$.
- Finite game, $\rho_i = \operatorname{rank} F_i \geq \operatorname{rank} S_i - 1$.
- Polynomial game with $|C_i| > m_i$ for all i, $\rho_i = \operatorname{rank} P_i \geq \operatorname{rank} S_i - 1$.

Now consider a two-player finite game having payoff matrices U_1 and U_2 where player 1 chooses rows and player 2 chooses columns. Then $S_1 = U_2'$ and $S_2 = U_1$, so the general separable game rank bound in Proposition 3.8 applied to Theorem 3.2 shows that the game has an equilibrium in which player 1 mixes among at most $\operatorname{rank} U_2 + 1$ pure strategies and player 2 mixes among at most $\operatorname{rank} U_1 + 1$ pure strategies, as proven by Lipton et al. [11]. The finite game bound proven in Proposition 3.8 yields strictly better results in some cases (by at most one), and applies also to the multiplayer case which was left open in [11].

Example 1 (cont'd): We can apply Proposition 3.8 and Theorem 3.2 to the example with payoffs given by (2). Using the procedure described above, we produce the matrices

$$S_1 = \begin{bmatrix} 0 & 0 & -1 & 0 \\ 4 & 0 & 1 & 0 \\ 0 & 0 & 2 & 0 \\ -4 & 0 & 0 & 0 \end{bmatrix}, S_2 = \begin{bmatrix} 0 & 0 & 0 & 3 \\ -1 & 2 & 0 & 0 \\ 0 & 0 & -3 & 0 \\ -2 & 0 & 0 & 0 \end{bmatrix}$$

and

$$P_1 = \begin{bmatrix} 0 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}, P_2 = \begin{bmatrix} 0 & 0 & 3 \\ 2 & 0 & 0 \\ 0 & -3 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Thus the rank of the example game is

$$\rho = (\operatorname{rank} P_1, \operatorname{rank} P_2) = (1, 3),$$

so in fact there exists an equilibrium in which player 1 mixes among at most 2 pure strategies and player 2 mixes among at most 4 pure strategies.

The manipulations done in this section essentially show how to compute the space Δ_i / Y_i as a linear image of Δ_i / W_i. The space Δ_i / X_i could be computed similarly. Applying these results to the example shows that Δ_1 / X_1, Δ_2 / X_2, and Δ_2 / Y_2 are all isomorphic to Δ_i / W_i, the set shown in Figure 1. On the other hand, since $\operatorname{rank} P_1 = 1$, the set Δ_1 / Y_1 is a 1-dimensional image of Δ_i / W_i, i.e. it is an interval.

IV. COMPUTATION OF NASH EQUILIBRIA AND APPROXIMATE EQUILIBRIA

In this section, we present an algorithm for computing approximate equilibria of separable games with infinite strategy sets which follows directly from the results on the rank of games given in Section III. First, we consider briefly the problem of computing exact equilibria.

The moments of an equilibrium can in principle be computed using the following generalization of the equilibrium formulation given by Başar and Olsder [1]:

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

where x_i are the moments, f_i is the moment function, and v_i is the payoff function on the moment spaces as defined in Section III. Each player also has an auxilliary variable p_i. The optimum objective value of this problem is zero and is attained exactly at the Nash equilibria. 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 computing the payoff 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 given in terms of linear matrix inequalities as in Parrilo's treatment of the zero-sum case [14]. This yields a problem with multiaffine objective and linear matrix inequality constraints.

Example 1 (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.

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

Definition 4.1: 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 $i = 1, \ldots, n$. If $\epsilon = 0$ then σ is called a **Nash equilibrium**.

Throughout the rest of this section, we will consider a separable game for which C_i is a compact interval for each i, and for which the utility funcitons satisfy a Lipschitz condition

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

for some $L_i \geq 0$, all $s_{-i} \in C_{-i}$ and all i. Clearly this is equivalent to requiring the same inequality for all $\sigma_{-i} \in \Delta_{-i}$. We must also be able to compute the values of u_i efficiently, so for example u_i could be a polynomial with rational coefficients. These assumptions could be relaxed, say by making C_i a box in $\mathbb{R}^n$ or by requiring a Hölder condition in place of a Lipschitz condition, but for the sake of clarity we do not do so here.

Fix $\epsilon > 0$. Divide the interval C_i into equal subintervals of length no more than $2\frac{\epsilon}{L_i}$; at most $\lceil \frac{1}{2\epsilon} l(C_i) L_i \rceil$ such intervals are required, where $l(C_i)$ denotes the length of the interval C_i. Let $\tilde{C}_i$ be the set of center points of these intervals and let $\tilde{u}_i$ be the corresponding sampled payoffs. 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 as in Proposition 3.8. By Theorem 3.2 and Proposition 3.3, 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 number of possible choices of at most $\rho_i + 1$ pure strategies from $\tilde{C}_i$ is

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

which is a polynomial in $|\tilde{C}_i|$ since ρ_i is fixed.

This leaves the step of checking whether there exists an equilibrium σ for a given choice of $\text{supp}(\sigma_i) \subseteq \tilde{C}_i$ with $|\text{supp}(\sigma_i)| \leq \rho_i + 1$ for each i, and if so, computing such an equilibrium. If the game has two players, the set of such equilibria for given supports is described by $O(|\tilde{C}_1| + |\tilde{C}_2|)$ linear equalities and inequalities (with more than two players these would become nonlinear), and hence an equilibrium or certificate of nonexistence of an equilibrium for a given support can be found in time polynomial in $|\tilde{C}_1| + |\tilde{C}_2| \propto \frac{1}{\epsilon}$. Since the number of possible supports to be checked is also polynomial in $\frac{1}{\epsilon}$, we have the following result.

Proposition 4.2: For $\epsilon > 0$, an ϵ-equilibrium of a two-player separable game satisfying the conditions of this section can be found in time polynomial in $\frac{1}{\epsilon}$ for fixed ρ and time polynomial in the components of ρ for fixed ϵ.

V. CHARACTERIZATIONS OF SEPARABLE GAMES

The following theorem presents two alternative characterizations of separable games, including the finite rank characterization mentioned above.

Theorem 5.1: 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 $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$.

Proof: $(1 \Rightarrow 3)$ This follows from Theorem 2.6.

$(2 \Rightarrow 1)$ Assume the game has finite rank. Since $V_i = \text{aff}(\Delta_i \cup \{0\})$ we have $\dim V_i/Y_i \leq 1 + \dim \Delta_i/Y_i$, so that V_i/Y_i is finite-dimensional and thus isomorphic to $\mathbb{R}^{r_i}$ for some r_i. Then from the (possibly infinite) collection of kernels of linear functionals whose intersection gives Y_i, r_i functionals can be chosen such that the intersection of their kernels is also Y_i. Note that each of these functionals is the integral of some continuous function $e_i^k(s_i) = \int_{C_{-i}} u_{j_k}(s_i, s_{-i}) d\sigma_{-i}^k$ over C_i. Define $\psi_i(\sigma_i) = (\int e_i^1 d\sigma_i, \ldots, \int e_i^{r_i} d\sigma_i)$ and note that $\ker \psi_i = Y_i$ so we can identify $\psi_i(V_i) = \mathbb{R}^{r_i}$ with V_i/Y_i. In particular, the continuity of the e_i^k shows that the restriction of the map ψ_i to the pure strategy space C_i is continuous.

Define a multilinear functional w_1 on $V_1 \times \Pi_{j=2}^n V_j/Y_j$ as follows. For $\sigma \in V$, let

$$w_1(\sigma_1, \psi_2(\sigma_2), \ldots, \psi_n(\sigma_n)) = u_1(\sigma).$$

By definition of ψ_i this function is well-defined on all of $V_1 \times \Pi_{j=2}^n V_j/Y_j$, and it is straightforward to verify that it is multilinear.

Let $\{b_i^1, \ldots, b_i^{r_i}\}$ be a basis for V_i/Y_i and $\{b_i^{*1}, \ldots, b_i^{*r_i}\}$ be the corresponding algebraic dual basis. Being linear functionals on a finite-dimensional vector space, these dual functionals are continuous. Letting $x_i = \psi_i(s_i)$ for notational

simplicity, we have

$$u_1(s) = w_1(s_1, x_2, \ldots, x_n)$$

$$= w_1\left(s_1, \sum_{j_2=1}^{r_2} b_2^{*j_2}(x_2)b_2^{j_2}, \ldots, \sum_{j_n=1}^{r_n} b_n^{*j_n}(x_n)b_n^{j_n}\right)$$

$$= \sum_{j_2=1}^{r_2} \cdots \sum_{j_n=1}^{r_n} b_2^{*j_2}(x_2)\cdots b_n^{*j_n}(x_n)w_1\left(s_1, b_2^{j_2}, \ldots, b_n^{j_n}\right).$$

Clearly $b_k^{*j_k}(x_j) = b_k^{*j_k}(\psi_k(s_k))$ is a continuous function of s_k for all k. Since the utility functions are continuous, $w_1\left(s_1, e_2^{j_2}, \ldots, e_n^{j_n}\right)$ is a continuous function of s_1 for all $j_2, \ldots, j_n$. Thus u_1 has the separable form given in (1), and the same argument holds for all u_i.

(3 $\Rightarrow$ 2) Assume condition 3. Since Y_i is closed under the norm topology, V_i/Y_i is a Banach space under the quotient norm and the projection transformation $\psi_i : V_i \to V_i/Y_i$ is a bounded linear operator. For the first half of this proof the topologies on V_i and V_i/Y_i will be those induced by these norms.

Since C_i is a bounded subset of V_i, $\psi_i(C_i)$ is a bounded subset of V_i/Y_i. Assume for a contradiction that $\operatorname{span}\psi_i(C_i)$ is infinite-dimensional, so $\psi_i(C_i)$ contains a countable linearly independent subset $B_i = \{\psi_i(s_i^1), \psi_i(s_i^2), \ldots\}$. By a theorem of Erdős and Straus (see [4]), we can assume (replacing B_i by a countable subset if necessary), that if $\sum_{k=1}^{\infty} c^k \psi_i(s_i^k) = 0$ then $c^k = 0$ for all k, a property which is in general stronger than algebraic linear independence.

Let $\{p^k\}$ be a sequence of positive reals summing to one. Define the mixed strategy

$$\sigma_i = \sum_{k=1}^{\infty} p^k s_i^k. \tag{10}$$

Clearly this sum converges in V_i. Since ψ_i is a bounded linear operator,

$$v = \psi_i(\sigma_i) = \sum_{k=1}^{\infty} p^k \psi_i(s_i^k) \tag{11}$$

converges as well. By assumption 3 we can choose an integer N and nonnegative coefficients q^k such that

$$v = \sum_{k=1}^{N} q^k \psi_i(s_i^k). \tag{12}$$

Define $q^k = 0$ for $k > N$. Then

$$\sum_{k=1}^{\infty}(p^k - q^k)\psi_i(s_i^k) = 0, \tag{13}$$

so $p^k = q^k$ for all k. This implies $p^{N+1} = q^{N+1} = 0$, contrary to assumption. Thus $\operatorname{span}\psi_i(C_i)$ is a finite dimensional vector space.

From now on we consider the weak topology on V_i, which makes it a topological vector space. Since Y_i is a closed subspace of V_i, V_i/Y_i is a Hausdorff topological vector space under the quotient topology and ψ_i is continuous with respect to this topology (see 5.7 and 7.3 of Ch.2 in [10]). Since V_i/Y_i

is Hausdorff, the finite-dimensional subspace $\operatorname{span}\psi_i(C_i)$ is closed, and hence its inverse image $\psi_i^{-1}(\operatorname{span}\psi_i(C_i))$ is a closed subspace of V_i containing the set of all finitely supported measures, which is a dense set. Thus $V_i = \psi_i^{-i}(\operatorname{span}\psi_i(C_i))$, so $V_i/Y_i \cong \psi_i(V_i) = \operatorname{span}\psi_i(C_i)$ is finite-dimensional and $\dim \Delta_i/Y_i \leq \dim V_i/Y_i < \infty$. ∎

The following counterexample shows that the containment $\operatorname{supp}\tau_i \subset \operatorname{supp}\sigma_i$ is a necessary part of condition 3 in Theorem 5.1 by showing that there exists a nonseparable continuous game in which every mixed strategy is payoff equivalent to a pure strategy.

Example 2: 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

$$d(x,x') = \sup_i \frac{|x_i - x_i'|}{i}. \tag{14}$$

Define the utilities

$$u_1(x,y) = u_2(x,y) = \sum_{i=1}^{\infty} 2^{-i} x_i y_i. \tag{15}$$

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

$$\begin{aligned}
|u_1(x,y) - u_1(x',y')| &= \left|\sum_{i=1}^{\infty} 2^{-i}(x_i y_i - x_i' y_i')\right| \\
&= \left|\sum_{i=1}^{\infty} 2^{-i}(x_i y_i - x_i' y_i + x_i' y_i - x_i' y_i')\right| \\
&\leq \sum_{i=1}^{\infty} 2^{-i}\left(y_i|x_i - x_i'| + x_i'|y_i - y_i'|\right) \\
&\leq \sum_{i=1}^{\infty} 2^{-i}(2\delta i) = \left(2\sum_{i=1}^{\infty} 2^{-i} i\right)\delta.
\end{aligned} \tag{16}$$

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.

$$u_1(\sigma,\tau) = \int u_1 d(\sigma \times \tau) = \sum_{i=1}^{\infty} 2^{-i}\left(\int x_i d\sigma\right)\left(\int y_i d\tau\right) \tag{17}$$

Thus σ is payoff equivalent to the pure strategy $\left(\int x_1 d\sigma, \int x_2 d\sigma, \ldots\right) \in C_1$ 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 $\sigma = \sum_{i=1}^{\infty} p_i e^i \in V_1$. Suppose σ were 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 i plays σ is

$$u_1(\sigma, e_{N+1}) = \int 2^{-(N+1)} x_{N+1} d\sigma = 2^{-(N+1)} p_{N+1}.$$

(18)

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 payoff equivalent. Thus condition 3 of Theorem 5.1 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, and the containment $\operatorname{supp} \tau_i \subset \operatorname{supp} \sigma_i$ cannot be removed from condition 3 of Theorem 5.1.

VI. CONCLUSIONS AND FUTURE WORK

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. [11] to be extended to the multi-player and polynomial cases. 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 [8]. 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 [14]. 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 extensible to this setting; McKelvey and McLennan have presented a survey of such algorithms [12].

REFERENCES

[1] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999.
[2] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA, 1997.
[3] X. Chen and X. Deng. Settling the complexity of 2-player Nash equilibrium. *Electronic Colloquium on Computational Complexity*, TR05-140, 2005.
[4] P. Erdös and E. G. Straus. On linear independence of sequences in a Banach space. *Pacific Journal of Mathematics*, 3(4):689 – 694, 1953.
[5] F. Germano and G. Lugosi. Existence of sparsely supported correlated equilibria. http://www.econ.upf.es/ lugosi/ssce.pdf, March 2006.
[6] 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.
[7] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1), February 1989.
[8] R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. arXiv:cs.GT/0511021, 2005.
[9] S. Karlin. *Mathematical Methods and Theory in Games, Programming, and Economics*, volume 2: Theory of Infinite Games. Addison-Wesley, Reading, MA, 1959.
[10] J. L. Kelley and I. Namioka. *Linear Topological Spaces*. Springer-Verlag, New York, 1963.
[11] 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.
[12] 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.
[13] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing (STOC 2005)*, New York, NY, 2005. ACM Press.
[14] P. A. Parrilo. Polynomial games and sum of squares optimization. *Submitted to CDC* 2006.
[15] K. R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, New York, NY, 1967.