Characterization and Computation of Correlated Equilibria in Infinite Games

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

Abstract— **Motivated by recent work on computing Nash equilibria in two-player zero-sum games with polynomial payoffs by semidefinite programming and in arbitrary polynomial-like games by discretization techniques, we consider the problems of characterizing and computing correlated equilibria in games with infinite strategy sets. We prove several characterizations of correlated equilibria in continuous games which are more analytically tractable than the standard definition and may be of independent interest. Then we use these to construct algorithms for approximating correlated equilibria of polynomial games with arbitrary accuracy, including a sequence of semidefinite programming relaxation algorithms and discretization algorithms.**

I. Introduction

Aumann's [1] notion of correlated equilibrium has received much attention as a generalization of the Nash equilibrium [5] which is both justifiable in theory [2] and efficiently computable in practice [6]. The idea of a correlated equilibrium is that each player receives a private recommendation of what strategy to play, but these recommendations may be correlated. If all the players know the joint distribution of the recommendations, then they can each compute the joint conditional distribution of their opponents' recommendations given their own recommendation. If each player's recommendation is always a best response to this conditional distribution, then the distribution of recommendations is called a correlated equilibrium. If additionally the recommendations to each player are independent, then the distribution is called a Nash equilibrium.

Hart and Schmeidler [3] have proven the existence of correlated equilibria in a large class of games with an arbitrary set of players and infinite strategy sets. Papadimitriou and Roughgarden [7] have observed that in general finite games with many players require an exponential amount of data to describe. They have proven conditions under which correlated equilibria of games which admit some succinct description can be computed efficiently. Papadimitriou [6] later extended these results to construct polynomial time algorithms for computing correlated equilibria in broader classes of succinctly representable games. While computing correlated equilibria in finite games has been studied extensively, to our knowledge there has not been work on the corresponding problem for games with infinite strategy sets.

This paper has two goals. The first goal is to characterize the correlated equilibria of a large class of games with infinite strategy sets. The definition of a correlated equilibrium in this setting involves a quantifier ranging over the set of all measurable self-maps of a given space. This is a complex object which cannot be easily parameterized. Hence alternative characterizations of correlated equilibria are a prerequisite for designing algorithms to compute correlated equilibria in these games. We give a characterization in which the above quantifier is replaced with a quantifier over all continuous maps from the space into $[0, \infty)$. This allows for a sequence of finitely parameterizable relaxations by using polynomials of fixed degree in place of arbitrary continuous functions.

The second goal is to construct algorithms for computing or approximating correlated equilibria in continuous games, particularly games with polynomial payoff functions. We give three such algorithms and compare them using a running example. First, as a baseline, we discuss the properties of the approximation algorithm obtained by discretizing the players' strategy spaces without regard to the game structure. Second, we suggest a method for choosing a sequence of discretizations adaptively which seems to perform better in practice. Third, we construct a sequence of semidefinite relaxations which converges to a description of the set of correlated equilibria in terms of joint moments.

The rest of the paper is organized as follows. In Section II we define correlated equilibria in finite and continuous games and extend several known characterizations of correlated equilibria in finite games to continuous games. We then use these characterizations to give several approximation algorithms for correlated equilibria of polynomial games in Section III. We close with conclusions and directions for future research in Section IV.

II. Characterizations of Correlated Equilibria

In this section we will define finite and continuous games along with correlated equilibria thereof. We will present several known characterizations of correlated equilibria in finite games, and then show how these naturally extend to continuous games.

Some notational conventions used throughout are that subscripts refer to players, while superscripts are frequently used for other indices (it will be clear from context when they represent 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. The set of Borel probability measures π over a metric space S is denoted $\Delta(S)$. For simplicity we will write $\pi(s)$ in place of $\pi(\{s\})$ for the measure of a singleton $\{s\} \subseteq S$. All polynomials will be assumed to have real coefficients.

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

This research was funded in part by National Science Foundation grant DMI-0545910 and AFOSR MURI subaward 2003-07688-1.

A. Finite Games

We begin with the definition.

Definition 2.1: A **finite game** consists of $n < \infty$ **players**, each of whom has a finite **pure strategy set** C_i and a **utility** or **payoff function** $u_i : C \to \mathbb{R}$, where $C = \Pi_{j=1}^n C_j$.

Each player's objective is to maximize his (expected) utility. We now consider what it would mean for the players to maximize their utility if their strategy choices were correlated. Let R be a random variable taking values in C distributed according to some measure $\pi \in \Delta(C)$. A realization of R is a **pure strategy profile** (a choice of pure strategy for each player) and the i^{th} component of the instantiation R_i will be called the recommendation to player i. Given such a recommendation, player i can use the conditional probability to form a posteriori beliefs about the recommendations given to the other players. A distribution π is defined to be a correlated equilibrium if no player can ever expect to unilaterally gain by deviating from his recommendation, assuming the other players play according to their recommendations.

Definition 2.2: A **correlated equilibrium** of a finite game is a joint probability measure $\pi \in \Delta(C)$ such that if R is a random variable distributed according to π then

$$\sum_{s_{-i} \in C_{-i}} \text{Prob}(R = s | R_i = s_i) \left[u_i(t_i, s_{-i}) - u_i(s) \right] \leq 0 \tag{1}$$

for all players i, all $s_i \in C_i$ such that $\text{Prob}(R_i = s_i) > 0$, and all $t_i \in C_i$.

While this definition captures the idea we have described above, the following characterization is easier to apply and visualize.

Proposition 2.3: A joint probability measure $\pi \in \Delta(C)$ is a correlated equilibrium of a finite game if and only if

$$\sum_{s_{-i} \in C_{-i}} \pi(s) \left[u_i(t_i, s_{-i}) - u_i(s) \right] \leq 0 \tag{2}$$

for all players i and all $s_i, t_i \in C_i$.

Proof: Using the definition of conditional probability we can rewrite the definition of a correlated equilibrium as the condition that

$$\sum_{s_{-i} \in C_{-i}} \frac{\pi(s)}{\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i})} \left[u_i(t_i, s_{-i}) - u_i(s) \right] \leq 0$$

for all i, all $s_i \in C_i$ such that $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in C_i$. The denominator does not depend on the variable of summation so it can be factored out of the sum and cancelled, yielding the simpler condition that (2) holds for all i, all $s_i \in C_i$ such that $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in C_i$. But if $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) = 0$ then the left hand side of (2) is zero regardless of i and t_i, so the equation always holds trivially in this case. ■

This proposition shows that the set of correlated equilibria is defined by a finite number of linear equations and inequalities (those in (2) along with $\pi(s) \geq 0$ for all $s \in C$ and $\sum_{s \in C} \pi(s) = 1$) and is therefore convex and even polyhedral. It can be shown via linear programming duality that this set is nonempty (see [3]); this can be shown alternatively by appealing to the fact that Nash equilibria exist and are the same as correlated equilibria which are product distributions.

We can think of correlated equilibria as joint distributions corresponding to recommendations which will be given to the players as part of an extended game. The players are then free to play any function of their recommendation (this is called a **departure function**) as their strategy in the game. If it is a Nash equilibrium of this extended game for each player to play his recommended strategy (i.e. to use the identity departure function), then the distribution is a correlated equilibrium. This interpretation is justified by the following alternative characterization of correlated equilibria.

Proposition 2.4: A joint probability measure $\pi \in \Delta(C)$ is a correlated equilibrium of a finite game if and only if

$$\sum_{s \in C} \pi(s) \left[u_i(\zeta_i(s_i), s_{-i}) - u_i(s) \right] \leq 0 \tag{3}$$

for all players i and all functions $\zeta_i : C_i \to C_i$.

Proof: By substituting $t_i = \zeta_i(s_i)$ into (2) and summing over all $s_i \in C_i$ we obtain (3) for any i and any $\zeta_i : C_i \to C_i$. For the converse, define ζ_i for any $s_i, t_i \in C_i$ by

$$\zeta_i(r_i) = \begin{cases} t_i & r_i = s_i \\ r_i & \text{else.} \end{cases}$$

Then all the terms in (3) except the s_i terms cancel yielding (2). ■

B. Continuous Games

Again we begin with the definition of this class of games.

Definition 2.5: A **continuous game** consists of $n < \infty$ players, each of whom has a pure strategy set C_i which is a compact metric space and a utility function $u_i : C \to \mathbb{R}$ which is continuous.

Note that any finite set forms a compact metric space under the discrete metric and any function out of such a set is continuous, so the class of continuous games includes the finite games. Another class of continuous games are the polynomial games, which we study more in Section III below.

Definition 2.6: A **polynomial game** is a continuous game in which the pure strategy spaces are $C_i = [-1, 1]$ for all players and the utility functions are polynomials.

Defining correlated equilibria in continuous games requires somewhat more care than in finite games. The standard definition as used in [3] is a straightforward generalization of the characterization of correlated equilibria for finite games in Proposition 2.4. In this case we must add the additional assumption that the departure functions be Borel measurable to ensure that the integrals are defined.

Definition 2.7: A **correlated equilibrium** of a continuous game is a joint probability measure $\pi \in \Delta(C)$ such that

$$\int \left[u_i(\zeta_i(s_i), s_{-i}) - u_i(s) \right] d\pi(s) \leq 0$$

for all i and all Borel measurable functions $\zeta_i : C_i \to C_i$.

Preprint submitted to 46th IEEE Conference on Decision and Control.
Received March 9, 2007.

The problem of computing Nash equilibria of polynomial games can be formulated exactly as a finite-dimensional nonlinear program or as a system of polynomial equations and inequalities [11]. The key feature of the problem which makes this possible is the fact that it has an explicit finite-dimensional formulation in terms of the moments of the players' mixed strategies.

On the other hand to our knowledge no exact finite-dimensional characterization of the set of correlated equilibria in polynomial games is known. Given the characterization of Nash equilibria in terms of moments, a natural attempt would be to try to characterize correlated equilibria in terms of the joint moments, i.e. the values $\int s_1^{k_1} \cdots s_n^{k_n} d\pi$ for nonnegative integers k_i and joint measures π. In fact we will be able to obtain such a characterization below, albeit in terms of infinitely many joint moments. The reason this attempt fails to yield a finite dimensional formulation is that the definition of a correlated equilibrium imposes conditions on the conditional distributions of the equilibrium measure. A finite set of moments does not seem to contain enough information about these conditional distributions to check the required conditions exactly. Therefore we also consider approximate correlated equilibria.

Definition 2.8: An **ϵ-correlated equilibrium** of a continuous game is a joint probability measure $\pi \in \Delta(C)$ such that

$$\int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \leq \epsilon$$

for all i and all Borel measurable functions $\zeta_i : C_i \to C_i$.

This definition reduces to that of a correlated equilibrium when $\epsilon = 0$.

Before considering any algorithms for computing approximate correlated equilibria, we will prove several alternative characterizations of exact and approximate correlated equilibria which are more amenable to analysis than the definition. These characterizations may also be of independent interest.

We begin with a technical lemma which we will use to prove the other characterization theorems. A more general version of this lemma has appeared as Lemma 20 in [12].

Lemma 2.9: Simple departure functions (those with finite range) suffice to define ϵ-correlated equilibria in continuous games. That is to say, a joint measure π is an ϵ-correlated equilibrium if and only if

$$\int [u_i(\xi_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \leq \epsilon$$

for all i and all Borel measurable simple functions $\xi_i : C_i \to C_i$.

Proof: The forward direction is trivial. To prove the reverse, first fix i. Then choose any measurable departure function ζ_i and let ξ_i^k be a sequence of simple measurable departure functions converging to ζ_i pointwise; such a sequence exists because the simple measurable departure functions on the compact metric space C_i are dense in the measurable departure functions in the sense of pointwise convergence. Then $u_i(\xi_i^k(s_i), s_{-i})$ converges to $u_i(\zeta_i(s_i), s_{-i})$ pointwise since u_i is continuous. Thus

$$\begin{aligned} &\int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \\ &\quad = \int \lim_{k\to\infty} \left[u_i(\xi_i^k(s_i), s_{-i}) - u_i(s)\right] d\pi(s) \\ &\quad = \lim_{k\to\infty} \int \left[u_i(\xi_i^k(s_i), s_{-i}) - u_i(s)\right] d\pi(s) \\ &\quad \leq \lim_{k\to\infty} \epsilon = \epsilon, \end{aligned}$$

where the second equality follows from Lebesgue's Dominated Convergence Theorem [10] and the inequality is by assumption. ■

The following characterization is a generalization of the standard formulation of correlated equilibria in finite games in terms of linear constraints presented in Proposition 2.3.

Theorem 2.10: A joint measure π is a correlated equilibrium of a continuous game if and only if

$$\int_{B_i \times C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)]\, d\pi(s) \leq 0 \qquad (4)$$

for all i, $t_i \in C_i$, and measurable subsets $B_i \subseteq C_i$.

Proof: ($\Rightarrow$) Fix i, $t_i \in C_i$ and a measurable set $B_i \subseteq C_i$. Define the function $\xi_i : C_i \to C_i$ by $\xi_i(s_i) = t_i$ if $s_i \in B_i$ and $\xi_i(s_i) = s_i$ otherwise. Then ξ_i is measurable, so by definition of a correlated equilibrium we have

$$\begin{aligned} &\int_{B_i \times C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)]\, d\pi(s) \\ &\qquad = \int [u_i(\xi_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \leq 0. \end{aligned}$$

($\Leftarrow$) By Lemma 2.9 it suffices to show that the condition defining correlated equilibria holds for all simple measurable departure functions ξ_i. Fix such a function and let $\{t_i^1, \ldots, t_i^k\}$ be its range. Let $B_i^k = \xi_i^{-1}(\{t_i^k\})$, which is measurable by assumption. Note that the sets $B_i^1, \ldots, B_i^k$ partition C_i, so we have

$$\begin{aligned} &\int [u_i(\xi_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \\ &\quad = \sum_{j=1}^{k} \int_{B_i^j \times C_{-i}} [u_i(\xi_i(s_i), s_{-i}) - u_i(s)]\, d\pi(s) \\ &\quad = \sum_{j=1}^{k} \int_{B_i^j \times C_{-i}} \left[u_i(t_i^j, s_{-i}) - u_i(s)\right] d\pi(s) \leq \sum_{j=1}^{k} 0 = 0, \end{aligned}$$

where the inequality follows from (4). ■

The preceding lemma and theorem can be extended to an arbitrary set of players. This result is interesting on its own since in [3] it was stated that for an arbitrary set of players and compact Hausdorff strategy spaces, the analog of Theorem 2.10 does not hold.

The next theorem is an alternative characterization of correlated equilibria in continuous games, which we will use in Section III to develop algorithms for computing (approximate) correlated equilibria.

Preprint submitted to 46th IEEE Conference on Decision and Control.
Received March 9, 2007.

Theorem 2.11: A joint measure π is a correlated equilibrium of a continuous game if and only if

$$\int f_i(s_i)\left[u_i(t_i, s_{-i}) - u_i(s)\right] d\pi(s) \leq 0 \qquad (5)$$

for all i and $t_i \in C_i$ as f_i ranges over any of the following sets of functions from C_i to $[0, \infty)$:

1) Characteristic functions of measurable sets,
2) Measurable simple functions,
3) Bounded measurable functions,
4) Continuous functions,
5) Squares of polynomials (if $C_i \subset \mathbb{R}^{k_i}$ for some k_i).

Proof: Condition 1 holds if and only if π is a correlated equilibrium; this is just a restatement of Theorem 2.10. Clearly condition 3 $\Rightarrow$ 1. By the linearity of the integral, 1 $\Rightarrow$ 2. Since every measurable function is the pointwise limit of simple measurable functions, 2 $\Rightarrow$ 3 by Lebesgue's Dominated Convergence Theorem (here we use the boundedness of the chosen measurable function).

Clearly 3 $\Rightarrow$ 4. On the other hand the set of bounded measurable functions is a subset of $\mathcal{L}^1(\pi_i)$, where π_i is the i^{th} marginal distribution of π. By Theorem 3.14 in [10] for example, the set of continuous functions on C_i is thus dense in the set of bounded measurable functions, with respect to the $\mathcal{L}^1(\pi_i)$ norm. Since the left hand side of (5) is a continuous linear functional of $f_i \in \mathcal{L}^1(\pi_i)$, we have 4 $\Rightarrow$ 3.

Finally we assume that $C_i \subseteq \mathbb{R}^{k_i}$ for some k_i. Clearly 4 $\Rightarrow$ 5. Let $f_i : C_i \to [0, \infty)$ be a continuous function. By the Stone-Weierstrass theorem $\sqrt{f_i}$ can be approximated arbitrary well by a polynomial p, with respect to the sup norm. Thus f_i can be approximated arbitrarily well by a square of a polynomial, with respect to the sup norm. Since convergence in the sup norm implies pointwise convergence, another application of Lebesgue's Dominated Convergence Theorem shows that 5 $\Rightarrow$ 4. ∎

Finally, we will consider ϵ-correlated equilibria which are supported on some finite subset. In this case, we obtain another generalization of Proposition 2.3.

Theorem 2.12: A probability measure $\pi \in \Delta(\tilde{C})$, where $\tilde{C} = \Pi_{j=1}^n \tilde{C}_j$ is a finite subset of C, is an ϵ-correlated equilibrium of a continuous game if and only if there exist ϵ_{i,s_i} such that

$$\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s)\left[u_i(t_i, s_{-i}) - u_i(s)\right] \leq \epsilon_{i,s_i}$$

for all players i, all $s_i \in \tilde{C}_i$, and all $t_i \in C_i$, and

$$\sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} \leq \epsilon$$

for all players i.

Proof: If we replace t_i with $\zeta_i(s_i)$ in the first inequality then sum over all $s_i \in \tilde{C}_i$ and combine with the second inequality, we get the equivalent condition that

$$\sum_{s \in \tilde{C}} \pi(s)\left[u_i(\zeta_i(s_i), s_{-i}) - u_i(s)\right] \leq \epsilon$$

holds for all i and any function $\zeta_i : \tilde{C}_i \to C_i$. This is exactly the definition of an ϵ-correlated equilibrium in the case when π is supported on $\tilde{C}$. ∎

III. Computing Correlated Equilibria

When faced with a game to analyze, we may be interested in any of the following problems:

(P1) computing a single correlated equilibrium,
(P2) computing a projection of the entire set of correlated equilibria, or
(P3) optimizing some objective function over the set of correlated equilibria.

Since computing exact correlated equilibria in continuous games is intractable, we focus in this section on developing algorithms that can compute approximate correlated equilibria with arbitrary accuracy. We consider three types of algorithms. In the first, the strategy sets are discretized without regard to the structure of the game. This algorithm applies to any continuous game, but may not use computational resources efficiently and may yield conservative performance estimates. The other two classes of algorithms apply only to polynomial games, but take advantage of the algebraic structure present and function better than the first in practice.

Example 3.1: We will use the following polynomial game to illustrate the algorithms presented below. The game has two players, x and y, who each choose their strategies from the interval $C_x = C_y = [-1, 1]$. Their utilities are given by

$$\begin{aligned} u_x(x,y) &= 0.596x^2 + 2.072xy - 0.394y^2 + 1.360x \\ &\quad - 1.200y + 0.554 \\ u_y(x,y) &= -0.108x^2 + 1.918xy - 1.044y^2 - 1.232x \\ &\quad + 0.842y - 1.886. \end{aligned}$$

where the coefficients have been selected at random. This example is convenient, because as Figure 3 shows, the game has a unique correlated equilibrium (the players choose $x = y = 1$ with probability one), which makes convergence easier to see. For the purposes of visualization and comparison, we will project the computed equilibria and approximations thereof into expected utility space, i.e. we will plot pairs $\left(\int u_x d\pi, \int u_y d\pi\right)$.

A. Static Discretization Methods

The techniques in this subsection are general enough to apply to arbitrary continuous games, so we will not restrict our attention to polynomial games here. The basic idea of **static discretization methods** is to select some finite subset $\tilde{C}_i \subset C_i$ of strategies for each player and limit his strategy choice to that set. Restricting the utility functions to the product set $\tilde{C} = \Pi_{i=1}^n \tilde{C}_i$ produces a finite game, called a **sampled game** or **sampled version** of the original continuous game. The simplest computational approach is then to consider the set of correlated equilibria of this sampled game. This set is defined by the linear inequalities in Proposition 2.3 along with the conditions that π be a probability measure on $\tilde{C}$, so in principle it is possible to solve any of the problems (P1) - (P3) for the discretized

Preprint submitted to 46th IEEE Conference on Decision and Control.
Received March 9, 2007.

game using standard linear programming techniques. The complexity of this approach in practice depends on the number of points in the discretization.

The question is then: what kind of approximation does this technique yield? In general the correlated equilibria of the sampled game may not have any relation to the set of correlated equilibria of the original game. The sampled game could, for example, be constructed by selecting a single point from each strategy set, in which case the unique probability measure over $\tilde{C}$ is automatically a correlated equilibrium of the sampled game but is a correlated equilibrium of the original game if and only if the points chosen form a pure strategy Nash equilibrium. Nonetheless, it seems intuitively plausible that if a large number of points were chosen such that any point of C_i were near a point of $\tilde{C}_i$ then the set of correlated equilibria of the finite game would be "close to" the set of correlated equilibria of the original game in some sense, despite the fact that each set might contain points not contained in the other.

To make this precise, we will show how to choose a discretization so that the correlated equilibria of the finite game are ϵ-correlated equilibria of the original game.

Proposition 3.2: Given a continuous game with strategy sets C_i and payoffs u_i along with any $\epsilon > 0$, there exist $\delta_i > 0$ such that if all points of C_i are within δ_i of a point of the finite set $\tilde{C}_i \subseteq C_i$ (such a $\tilde{C}_i$ exists since C_i is a compact metric space, hence totally bounded) then all correlated equilibria of the sampled game with strategy spaces $\tilde{C}_i$ and utilities $u_i|_{\tilde{C}}$ will be ϵ-correlated equilibria of the original game.

Proof: Note that the utilities are continuous functions on a compact set, hence uniformly continuous. Therefore for any $\epsilon > 0$ we can choose $\delta_i > 0$ such that if we change any of the arguments of u_i by no more than δ_i, then u_i changes by no more than ϵ. Let $\tilde{C}$ satisfy the stated assumption and let π be any correlated equilibrium of the corresponding finite game. Then by Proposition 2.3,

$$\sum_{s_{-i}\in\tilde{C}_{-i}} \pi(s)\left[u_i(t_i, s_{-i}) - u_i(s)\right] \leq 0$$

for all i and all $s_i, t_i \in \tilde{C}_i$. Any $t_i \in C_i$ is within δ of some $\tilde{t}_i \in \tilde{C}_i$, so

$$\begin{aligned}\sum_{s_{-i}\in\tilde{C}_{-i}} &\pi(s)\left[u_i(t_i, s_{-i}) - u_i(s)\right] \\ &\leq \sum_{s_{-i}\in\tilde{C}_{-i}} \pi(s)\left[u_i\left(\tilde{t}_i, s_{-i}\right) - u_i(s) + \epsilon\right] \\ &\leq \epsilon \sum_{s_{-i}\in\tilde{C}_{-i}} \pi(s).\end{aligned}$$

Therefore the assumptions of Theorem 2.12 are satisfied with $\epsilon_{i,s_i} = \epsilon \sum_{s_{-i}\in\tilde{C}_{-i}} \pi(s)$. ■

The proof shows that if the utilities are Lipschitz functions, such as polynomials, then the δ_i can in fact be chosen proportional to ϵ, so the number of points needed in $\tilde{C}_i$ is $O\left(\frac{1}{\epsilon}\right)$. More concretely, if the strategy spaces are $C_i = [-1, 1]$ as in a polynomial game, then $\tilde{C}_i$ can be chosen to be uniformly spaced within $[-1, 1]$, and if this is done ϵ will be $O\left(\frac{1}{d}\right)$ where $d = \max_i \left|\tilde{C}_i\right|$.

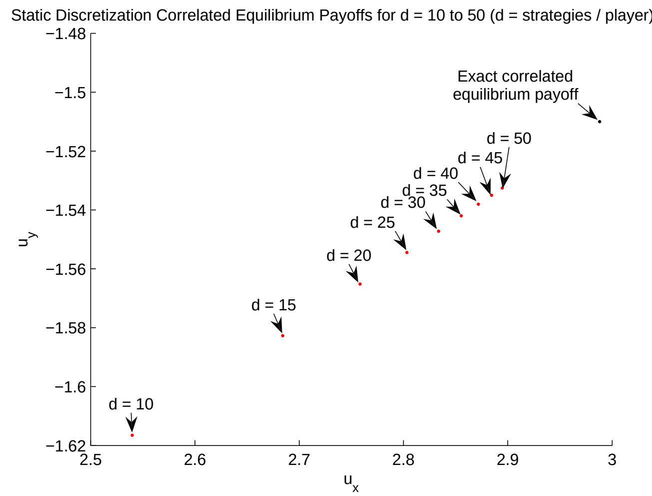

Fig. 1. Convergence of a sequence of ϵ-correlated equilibria of the game in Example 3.1 computed by a sequence of static discretizations, each with some number d of equally spaced strategies chosen for each player. The axes represent the utilities received by players x and y. It can be shown that the convergence in this example happens at a rate $\Theta\left(\frac{1}{d}\right)$.

Example 3.1 (continued): Figure 1 is a sequence of static discretizations for this game for increasing values of d, where d is the number of points in $\tilde{C}_x$ and $\tilde{C}_y$. These points are selected by dividing $[-1, 1]$ into d subintervals of equal length and letting $\tilde{C}_x = \tilde{C}_y$ be the set of midpoints of these subintervals. For this game it is possible to show that the rate of convergence is in fact $\Theta\left(\frac{1}{d}\right)$ so the worst case bound on convergence rate is achieved in this example.

B. Adaptive Discretization Methods

For the next two subsections, we restrict attention to the case of polynomial games. While static discretization methods are straightforward, they do not exploit the algebraic structure of polynomial games. Furthermore, the sampling of points in C_i to produce an ϵ-correlated equilibrium via Proposition 3.2 is conservative, in two senses. First, the ϵ-correlated equilibrium produced by that method may in fact be an ϵ^*-correlated equilibrium for some ϵ^* which is much less than ϵ; below we will show how to compute the minimal value of ϵ^* for a given joint probability measure. Second, it can be shown that any polynomial game has a Nash equilibrium, and hence a correlated equilibrium, which is supported on a finite set (see [11]). Hence, at least in principle there is no need for the number of points in $\tilde{C}_i$ to grow without bound as $\epsilon \to 0$. In this subsection we consider methods in which the points in the discretization are chosen more carefully.

An **adaptive discretization method** is an iterative procedure in which the finite set of strategies $\tilde{C}_i$ available to player i changes in some way on each iteration; we let $\tilde{C}_i^k$ denote the strategies available to player i on the k^{th} iteration.

Preprint submitted to 46th IEEE Conference on Decision and Control.
Received March 9, 2007.

The goal of such a method is to produce a sequence of ϵ^k-correlated equilibria with $\epsilon^k \to 0$.

There are many possible update rules to generate $\tilde{C}_i^{k+1}$ from $\tilde{C}_i^k$. The simplest are the **dense update rules** in which $\tilde{C}_i^1 \subseteq \tilde{C}_i^2 \subseteq \ldots$ and $\bigcup_{k=1}^{\infty} \tilde{C}_i^k$ is dense in C_i for all i. However, if such a method adds points without regard to the problem structure many iterations may be wasted adding points which do not get the algorithm any closer to a correlated equilibrium. Furthermore, the size of the discretized strategy sets $\tilde{C}_i^k$ may become prohibitively large before the algorithm begins to converge. Therefore it seems advantageous to choose the points to add to $\tilde{C}_i^k$ in a structured way, and it may also be worthwhile to delete points which don't seem to be in use after a particular iteration.

To get a handle on the convergence properties of these algorithms, we will use the ϵ-correlated equilibrium characterization in Theorem 2.12 since we are dealing with sampled strategy spaces. By that theorem, we can begin with a product set $\tilde{C} \subseteq C$ and find the joint measures $\pi \in \Delta(\tilde{C})$ which correspond to ϵ-correlated equilibria with minimal ϵ values by solving the following optimization problem:

$$\begin{aligned} &\text{minimize} && \epsilon \\ &\text{s.t.} && \\ \sum_{s_{-i}\in\tilde{C}_{-i}} \pi(s)\left[u_i(t_i,s_{-i}) - u_i(s)\right] &\le \epsilon_{i,s_i} && \text{for all } i,\ s_i\in\tilde{C}_i, \\ & && \text{and } t_i \in C_i \\ \sum_{s_i\in\tilde{C}_i} \epsilon_{i,s_i} &\le \epsilon && \text{for all } i \\ \pi(s) &\ge 0 && \text{for all } s\in\tilde{C} \\ \sum_{s\in\tilde{C}} \pi(s) &= 1 \end{aligned} \qquad (6)$$

For fixed s_{-i}, the functions $u_i(t_i, s_{-i})$ are univariate polynomials in t_i, so this problem can be solved exactly as a semidefinite program (Lemmas A.2 and A.3 in the appendix).

If the sequence of optimal ϵ values tends to zero for any game under a given update rule, we say that rule converges. Dense update rules converge by Proposition 3.2. Given the problem (6), a natural category of update rules are those which select an optimal solution to the problem, remove any strategies which are assigned zero or nearly zero probability in this solution, then add some or all of the values t_i which make the inequalities tight in this optimal solution into $\tilde{C}_i^k$ to obtain $\tilde{C}_i^{k+1}$. This corresponds to selecting constraints in Definition 2.7 which are maximally violated by the chosen optimal solution, so we call these **maximally violated constraint update rules**. These rules seem to perform well in practice, but it is not known whether they converge in general.

Example 3.1 (continued): In Figure 2 we illustrate an adaptive discretization method using a maximally violated constraint update rule. The solver was initialized with $\tilde{C}_x^0 = \tilde{C}_y^0 = \{0\}$. At each iteration the ϵ-correlated equilibrium π of minimal ϵ-value was computed. Then ϵ was reported and one player's sampled strategy set was enlarged, the player for whom the constraint $\sum_{s_i\in\tilde{C}_i^k} \epsilon_{i,s_i} \le \epsilon$ was tight. To choose which points to add to $\tilde{C}_i^k$, the algorithm identified the points $s_i \in \tilde{C}_i^k$ which were assigned positive probabiility under π. For each such s_i the values of $t_i \in C_i$ making the constraints in (6) tight were added to $\tilde{C}_i^k$ to obtain $\tilde{C}_i^{k+1}$. The other player's strategy set was not changed.

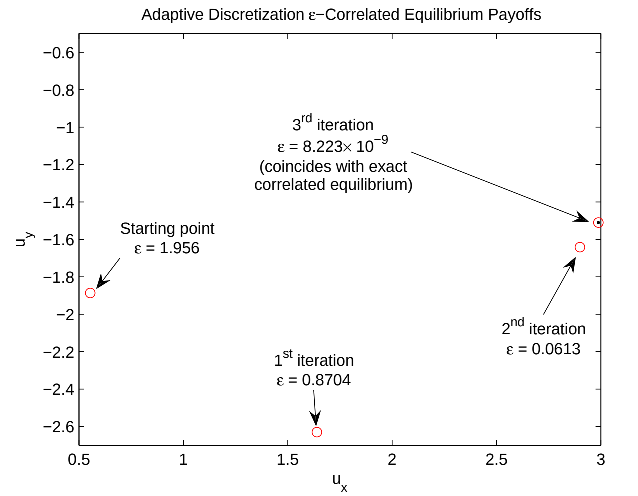

Fig. 2. Convergence of an adaptive discretization method with a maximally violated constraint update rule (note the change in scale from Figure 1). At each iteration, the expected utility pair is plotted along with the computed value of ϵ for which that iterate is an ϵ-correlated equilibrium of the game. In this case convergence to $\epsilon = 0$ (to within numerical error) occurred in three iterations.

In this case convergence happened in three iterations, significantly faster than the static discretization method. The resulting strategy sets were $\tilde{C}_x^3 = \{0, 1\}$ and $\tilde{C}_y^3 = \{0, 0.9131, 1\}$.

C. Moment Relaxation Methods

In this subsection we again consider only polynomial games. The **moment relaxation methods** for computing correlated equilibria have a different flavor from the discretization methods discussed above. Instead of using tractable finite approximations of the correlated equilibrium problem derived via discretizations, we begin with the alternative exact characterization given in condition 5 of Theorem 2.11. In particular, a measure π on C is a correlated equilibrium if and only if

$$\int p^2(s_i)\left[u_i(t_i, s_{-i}) - u_i(s)\right] d\pi(s) \le 0 \qquad (7)$$

for all i, $t_i \in C_i$, and polynomials p. If we wish to check all these conditions for polynomials p of degree less than or equal to d, we can form the matrices

$$S_i^d = \begin{bmatrix} 1 & s_i & s_i^2 & \cdots & s_i^d \\ s_i & s_i^2 & s_i^3 & \cdots & s_i^{d+1} \\ s_i^2 & s_i^3 & s_i^4 & \cdots & s_i^{d+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s_i^d & s_i^{d+1} & s_i^{d+2} & \cdots & s_i^{2d} \end{bmatrix}.$$

Let c be a column vector of length $d+1$ whose entries are the coefficients of p, so $p^2(s_i) = c' S_i^d c$. If we define

$$M_i^d(t_i) = \int S_i^d \left[u_i(t_i, s_{-i}) - u_i(s) \right] d\pi(s).$$

then (7) is satisfied for all p of degree at most d if and only if $c' M_i^d(t_i) c \leq 0$ for all c, i.e. if and only if $M_i^d(t_i)$ is negative semidefinite.

The matrix $M_i^d(t_i)$ has entries which are polynomials in t_i with coefficients which are linear in the joint moments of π. To check the condition that $M_i^d(t_i)$ be negative semidefinte for all $t_i \in [-1,1]$ for a given d we can use a semidefinite program (Lemma A.4 in the appendix), so as d increases we obtain a sequence of semidefinite relaxations of the correlated equilibrium problem and these converge to the exact condition for a correlated equilibrium.

We can also let the measure π vary by replacing the moments of π with variables and constraining these variables to satisfy some necessary conditions for the moments of a joint measure on C (see appendix). These conditions can be expressed in terms of linear matrix inequalities and there is a sequence of these conditions which converges to a description of the exact set of moments of a joint measure π. Thus we obtain a nested sequence of semidefinite relaxations of the set of moments of measures which are correlated equilibria, and this sequence converges to the set of correlated equilibria. In this way we can use moment relaxation methods to solve problems (P2) and (P3) given above.

Example 3.1 (continued): Figure 3 shows moment relaxations of orders $d = 0, 1,$ and 2. Since moment relaxations are outer approximations of the set of a correlated equilibria and the 2$^{\text{nd}}$ order moment relaxation corresponds to a unique point in expected utility space, all correlated equilibria of the example game have exactly this expected utility. In fact, the set of points in this relaxation is a singleton (even without being projected into utility space), so this proves that the example game has a unique correlated equilibrium.

IV. Conclusions and Future Work

We have shown how to generalize several characterizations of correlated equilibria in finite games to the larger class of games with continuous utility functions. In games with infinite strategy sets the definition of a correlated equilibrium involves a quantifier ranging over the set of all measurable functions from the strategy set to itself, which may be quite complicated and is not easily parametrized. The characterizations we present simplify the definition by allowing the quantifiers to range over more manageable sets. This makes it possible to construct algorithms which approximate correlated equilibria arbitrarily well.

Several questions about these algorithms remain open. Empirical evidence suggests that adaptive discretization algorithms using maximally violated constraint update rules converge quickly, but none has yet been proven to converge in general. Perhaps the primary question is whether there exists such a rule under which convergence is guaranteed, or even better, under which a statement about rate of convergence can be proven. Although the moment relaxation methods are known to converge, the problem of finding a bound on the rate of convergence for these algorithms is also still open.

Computing Nash equilibria of two-player zero-sum games and correlated equilibria of arbitrary games are two of the main equilibrium problems in game theory which lead to convex optimization problems. In [8] and this paper it has been shown that both of these can be solved using sum of squares methods in the case of polynomial games. We leave the task of extending these results to other convex equilibrium-type problems in polynomial games for future work.

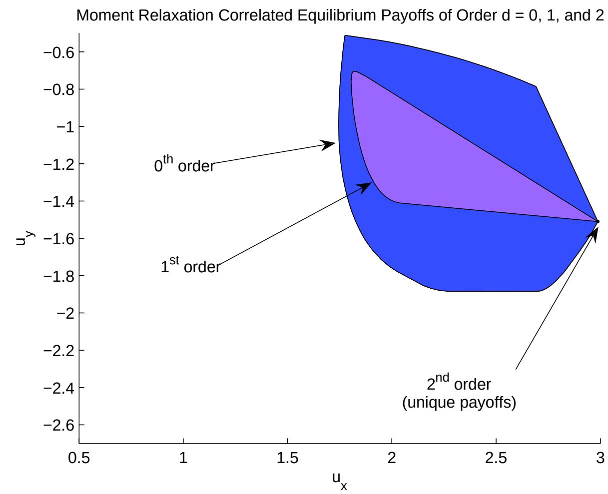

Fig. 3. Moment relaxations approximating the set of correlated equilibrium payoffs. The second order relaxation is a singleton, so this game has a unique correlated equilibrium.

V. Acknowledgements

The authors would like to thank Professor Muhamet Yildiz for a productive discussion which led to the formulation of Theorems 2.10 and 2.11 as well as the moment relaxation methods presented in Subsection 3-C.

Appendix
Sum of squares techniques

Below we will summarize some of the sum of squares results used in the algorithms above. Broadly, sum of squares methods allow nonnegativity conditions on polynomials to be expressed exactly or approximately as small semidefinite programs, and hence to be used in optimization problems which can be solved efficiently by interior point methods for semidefinite programming. The condition that a list of numbers correspond to the moments of a measure is dual to polynomial nonnegativity and can also be represented by similar semidefinite conditions.

The idea of sum of squares techniques is that the square of a real-valued function is nonnegative on its entire domain, and hence the same is true of a sum of squares of real-valued functions. In particular, any polynomial of the form

$p(x) = \sum p_k^2(x)$, where p_k are polynomials, is guaranteed to be nonnegative for all x. This gives a sufficient condition for a polynomial to be nonnegative. It is a classical result that this condition is also necessary if p is univariate [9].

Lemma A.1: A univariate polynomial p is nonnegative on $\mathbb{R}$ if and only if it is a sum of squares.

Frequently we are interested in polynomials which are nonnegative only on some interval such as $[-1, 1]$. These can be characterized almost as simply.

Lemma A.2: A univariate polynomial p is nonnegative on $[-1, 1]$ if and only if $p(x) = s(x) + (1 - x^2)t(x)$ where s and t are sums of squares.

These sum of squares conditions are easy to express using linear equations and semidefinite constraints. The proof of the following claim proceeds by factoring the positive semidefinite matrix P as a product $P = Q'Q$.

Lemma A.3: A univariate polynomial $p(x) = \sum_{k=0}^{d} p_k x^k$ of degree d is a sum of squares if and only if there exists a $(d + 1) \times (d + 1)$ positive semidefinite matrix P which satisfies $p_k = \sum_{i+j=k} P_{i,j}$ when the rows and columns of P are numbered 0 through d.

Similar semidefinite characterizations exist for multivariate polynomials to be sums of squares. While the condition of being a sum of squares does not characterize general nonnegative multivariate polynomials exactly, there exist sequences of sum of squares relaxations which can approximate the set of nonnegative polynomials (on e.g. $\mathbb{R}^k$, $[-1, 1]^k$, or a more general semialgebraic set) arbitrarily tightly [9]. Furthermore, for some special classes of multivariate polynomials, the sum of squares condition is exact. For example, the condition that a square matrix $M(t)$ whose entries are polynomials be positive semidefinite for $-1 \leq t \leq 1$ amounts to checking whether the multivariate polynomial $x'M(t)x$ is nonnegative for $1 \leq t \leq 1$ and all $x \in \mathbb{R}^k$.

Lemma A.4: A matrix $M(t)$ whose entries are univariate polynomials in t is positive semidefinite on $[-1, 1]$ if and only if $x'M(t)x = S(x, t) + (1 - t^2)T(x, t)$ where S and T are polynomials which are sums of squares.

Now suppose we wish to answer the question of whether a finite sequence $(\mu^0, \ldots, \mu^k)$ of reals correspond to the moments of a measure on $[-1, 1]$, i.e. whether there exists a positive measure μ on $[-1, 1]$ such that $\mu^i = \int x^i d\mu(x)$. Clearly if such a measure exists then we must have $\int p(x)d\mu(x) \geq 0$ for any polynomial p of degree at most k which is nonnegative on $[-1, 1]$. But any such integral is a linear combination of the moments $(\mu^0, \ldots, \mu^k)$ by definition and the polynomials p which are nonnegative on $[-1, 1]$ can be characterized with semidefinite constraints using Lemmas A.2 and A.3. Therefore this necessary condition for $(\mu_0, \ldots, \mu_k)$ to be the moments of a measure on $[-1, 1]$ can be written in terms of semidefinite constraints. It turns out that this condition is also sufficient [4], so we have:

Lemma A.5: The condition that a finite sequence of numbers $(\mu^0, \ldots, \mu^k)$ be the moments of a positive measure on $[-1, 1]$ can be written in terms of linear equations and semidefinite matrix constraints.

One can formulate similar questions about whether a finite sequence of numbers corresponds to the joint moments $\int x_1^{i_1} \cdots x_k^{i_k} d\mu(x)$ of a positive measure μ on $[-1, 1]^k$ (or a more general semialgebraic set). Using a sequence of semidefinite relaxations of the set of nonnegative polynomials on $[-1, 1]^k$, a sequence of necessary conditions for joint moments can be obtained in the same way as the conditions for moments of univariate measures. While no single one of these conditions is sufficient for a list of numbers to be joint moments, these conditions approximate the set of joint moments arbitrarily closely.

References

[1] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67 – 96, 1974.

[2] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1 – 18, January 1987.

[3] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1), February 1989.

[4] S. Karlin and L. S. Shapley. *Geometry of Moment Spaces*. American Mathematical Society, Providence, RI, 1953.

[5] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286 – 295, September 1951.

[6] 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.

[7] C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multiplayer games. In *Proceedings of the sixteenth annual ACM-SIAM symposium on discrete algorithms (SODA)*, 2005.

[8] P. A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*, 2006.

[9] B. Reznick. Some concrete aspects of Hilbert's 17th problem. In C. N. Delzell and J. J. Madden, editors, *Real Algebraic Geometry and Ordered Structures*, pages 251 – 272. American Mathematical Society, 2000.

[10] W. Rudin. *Real & Complex Analysis*. WCB / McGraw-Hill, New York, 1987.

[11] N. D. Stein, A. Ozdaglar, and P. A. Parrilo. Separable and low-rank continuous games. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2849 – 2854, 2006.

[12] G. Stoltz and G. Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, to appear.

Preprint submitted to 46th IEEE Conference on Decision and Control.
Received March 9, 2007.