

Correlated Equilibria in Continuous Games: Characterization and Computation

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

November 18, 2008

Abstract

We present several new characterizations of correlated equilibria in games with continuous utility functions. These have the advantage of being more computationally and analytically tractable than the standard definition in terms of departure functions. We use these characterizations to construct effective algorithms for approximating a single correlated equilibrium or the entire set of correlated equilibria of a game with polynomial utility functions. We then exhibit the rich structure of the set of correlated equilibria by analyzing the simplest of polynomial games, the mixed extension of matching pennies. We show that while the correlated equilibrium set is convex, the structure of its extreme points can be quite complicated. In finite games there can be a superexponential separation between the number of extreme Nash and extreme correlated equilibria. In polynomial games there can exist extreme correlated equilibria which are not finitely supported; we construct a large family of examples using techniques from ergodic theory. These examples show that in general the set of correlated equilibrium distributions of a polynomial game cannot be described by conditions on finitely many joint moments, in marked contrast to the set of Nash equilibria which is always expressible in terms of finitely many moments.

1 Introduction

In finite games correlated equilibria are simpler than Nash equilibria in several senses – mathematically at least, if not conceptually. The set of correlated equilibria is a convex polytope, described by finitely many explicit linear inequalities, while the set of (mixed) Nash equilibria can be essentially any real algebraic variety (set described by polynomial equations on real variables) [7]. The existence of correlated equilibria can be proven by

^{*}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.

| | Nash equilibria (non-zero sum) | Nash equilibria (zero sum) | correlated equilibria |
|------------------|-----------------------------------|-------------------------------|-----------------------|
| Finite games | Semialgebraic set [17] | LP | LP [2] |
| Polynomial games | Semialgebraic set [27] | SDP [21] | ? |

Table 1: Comparison of the simplest known description of different classes of equilibrium sets in finite and polynomial games.

elementary means (linear programming or game theoretic duality [12]), whereas the existence of Nash equilibria seems to require nonconstructive methods (fixed point theorems) or the analysis of complicated algorithms [18, 16]. Computing a sample correlated equilibrium or a correlated equilibrium optimizing some quantity such as social welfare can be done efficiently [11, 19]; strong evidence in complexity theory suggests that the corresponding problems for Nash equilibria are hard [6, 4, 11].

There are several exceptional classes of games for which the above problems about Nash equilibria become easy. The most important here are the zero-sum games. Broadly speaking, Nash equilibria of these games have complexity similar to correlated equilibria of general games. In particular, the set of Nash equilibria is an easily described convex polytope, existence can be proven by duality, and a sample equilibrium can be computed efficiently.

The situation in games with infinite strategy sets is not nearly so clear. For now we restrict attention to the simplest such class of games, those with finitely many players, strategy sets equal to $[-1, 1]$, and polynomial utility functions. Little is known about correlated equilibria of these games, but much is known about Nash equilibria. Most importantly, the set of mixed Nash equilibria is nonempty and admits a finite-dimensional description in terms of the moments of the players' mixed strategies [27].

This set of moments can be described explicitly in terms of polynomial equations and inequalities [27]. The Nash equilibrium conditions are expressible via first order statements, so the set of all moments of Nash equilibria is a real algebraic variety and can be computed in theory, albeit not efficiently in general. In the two-player zero-sum case, the set of Nash equilibria can be described by a semidefinite program (an SDP is a generalization of a linear program which can be efficiently solved, see the appendix), hence we can compute a sample Nash equilibrium or one which optimizes some linear functional in polynomial time [21]. A summary of the results described so far is shown in Table 1.

Contributions The impetus for this paper was to address the bottom right cell of Table 1, the one with the question mark. The table seems to suggest that the set of correlated equilibria of a polynomial game should be describable by a semidefinite program. We will see that this is approximately true, but not exactly. The contribution of this paper is threefold.

- First, we present several new characterizations of correlated equilibria in games with continuous utility functions (polynomiality is not needed here). In particular we show that the standard definition of correlated equilibria in terms of measurable departure

functions is equivalent to other definitions in which the utilities are integrated against all test functions in some class (Theorem 2.9). This characterization does not have any obvious game theoretic significance, but it is extremely useful analytically and it forms the base for our other contributions.

- Second, we present several algorithms for approximating correlated equilibria within an arbitrary degree of accuracy. We present one inefficient linear programming based method as a benchmark, followed by two semidefinite programming based algorithms which perform much better in practice. The first SDP algorithm, called adaptive discretization, iteratively computes a sequence of approximate correlated equilibria supported on finite sets (Section 3.2). We enlarge the support sets at each iteration using a heuristic which guarantees convergence in general and yields fast convergence in practice. The second SDP algorithm, called moment relaxation, does not discretize the strategy spaces but instead works in terms of joint moments. It produces a nested sequence of outer approximations to the set of joint moments of correlated equilibrium distributions, and these approximate equilibrium sets are described by semidefinite programs (Section 3.3). These relaxations depend crucially on one of the correlated equilibrium characterizations we have developed.
- Third, we use our new characterizations to show that the set of correlated equilibria can be surprisingly complex, even in simple games. To this end we analyze the mixed extension of matching pennies, a polynomial game with bilinear payoffs, in detail. We show that the set of correlated equilibria has extreme points with arbitrarily large finite support and also with infinite support (Section 4.2). This cannot happen if a set of measures can be written as those measures satisfying certain conditions on finitely many joint moments, so it provides a counterexample showing that in general the set of correlated equilibria of a polynomial cannot be described in this way (Proposition 4.5).

Related Literature The questions we address and the techniques we use are inspired by existing literature in two main areas. First, our work is related to a number of papers in the game theory literature.

- Aumann defined correlated equilibria in his famous paper [1], focusing on finite games to establish the basic properties and important examples. He obtained existence as a consequence of Nash’s theorem on the existence of Nash equilibria in finite games [18].
- Hart and Schmeidler showed that existence of correlated equilibria in finite games could be proven directly by a duality argument [12]. They then use a careful limiting argument to prove existence of correlated equilibria in continuous games with compact Hausdorff strategy spaces (Theorem 3 of that paper). The germs of ideas in this limiting argument are developed further in Section 2 of the present paper to yield various characterizations of correlated equilibria. It is worth noting that in [12] the authors also consider part (1) of Corollary 2.13 as a candidate definition of correlated

equilibria. They discard it as not obviously capturing the game theoretic idea of correlated equilibrium, but we prove that it is nonetheless an equivalent definition.

- Stoltz and Lugosi have studied learning algorithms which converge to correlated equilibria in continuous games [28]. These procedures do not seem to lead to efficient techniques for computing correlated equilibria, but are interesting in their own right. The authors consider replacing the class of all measurable departure functions with a smaller class, such as simple or continuous departure functions, and study when this yields an equivalent equilibrium notion (see Lemma 2.10 below).
- Separately from the literature on correlated equilibria, the structure of Nash equilibria in zero-sum games with polynomial or separable (polynomial-like) utility functions have been studied in detail by Dresher, Karlin, and Shapley. They show how to cast separable games as finite-dimensional “convex games” by replacing the infinite-dimensional mixed strategy spaces with finite-dimensional spaces of moments [9] and prove existence of equilibria via fixed point arguments [8]. There always exist finitely supported equilibria in separable games as can be shown using the finite-dimensionality of the moment spaces. The rich geometry of these spaces is studied in [14]. Most of these results as well as ad hoc methods for computing equilibria in simple cases are summarized in Karlin’s book [13]. The authors of the present paper have studied generalizations and extensions of these results in nonzero-sum separable games [27].

Second, our work is related to results from the optimization and computer science literature.

- Aumann showed that the set of correlated equilibria of a finite game is defined by polynomially many (in the size of the payoff tables) linear inequalities [2]. However, it was not clear whether this meant they could be computed in polynomial time. This question was settled in the affirmative when Khachian proved that linear programs could be solved in polynomial time; for an overview of this and other more efficient algorithms, see [3]. Papadimitriou extended this result, showing that correlated equilibria can be computed efficiently in many classes of games for which the payoffs can be written succinctly, even if the explicit payoff tables would be exponential in size [19].
- The breakthrough in optimization most directly related to the work in this paper is the development of semidefinite programming, a far-reaching generalization of linear programming which is still polynomial-time solvable (for an overview, see the appendix and [30]). More specifically, the development of sum of squares methods has allowed many optimization problems involving polynomials or moments of measures to be solved efficiently [20]. Parrilo has applied these techniques to efficiently compute Nash equilibria of two-player zero-sum polynomial games [21].

The remainder of this paper is organized as follows. In Section 2 we define the classes of games we study and correlated equilibria thereof, then prove several characterization

theorems. We present algorithms for approximating sample correlated equilibria and the set of correlated equilibria of polynomial games in Section 3. Then we analyze the set of correlated equilibria of the mixed extension of matching pennies in Section 4, with particular emphasis on computing extreme points of this set. Finally, we close with conclusions and directions for future work.

2 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 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 the context when they represent 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. The set of regular Borel probability measures π over a compact Hausdorff space S is denoted by $\Delta(S)$. 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.

2.1 Finite Games

We start with the definition of a finite game.

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

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 realization R_i will be called the recommendation to player i . Given such a recommendation, player i can use 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

$$\mathbb{E} [u_i(t_i, R_{-i}) - u_i(R) | R_i = s_i] \equiv \sum_{s_{-i} \in C_{-i}} \text{Prob}(R = s | R_i = s_i) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0$$

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*

$$\sum_{s_{-i} \in C_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0 \quad (1)$$

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

This proposition shows that the set of correlated equilibria is defined by a finite number of linear equations and inequalities (those in (1) 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 [12]. 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. if no player has an incentive to unilaterally deviate from using 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*

$$\sum_{s \in C} \pi(s) [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] \leq 0 \quad (2)$$

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

2.2 Continuous Games

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

Definition 2.5. A **continuous game** consists of an arbitrary (possibly infinite) set I of players i , each of whom has a pure strategy set C_i which is a compact Hausdorff space and a utility function $u_i : C \rightarrow \mathbb{R}$ which is continuous.

Note that any finite set forms a compact Hausdorff space under the discrete topology 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 are our primary focus when we study computation of correlated equilibria in the sections which

follow. The theorems and proofs below can safely be read with polynomial games in mind, ignoring such topological subtleties as regularity of measures. However the extra generality of arbitrary continuous games requires little additional work in the proofs of the characterization theorems, so we will not formally restrict our attention to polynomial games here.

Definition 2.6. A **polynomial game** is a continuous game with $n < \infty$ players 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 [12] 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

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

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

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 [27]. 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.

To see this, suppose that player 1 chooses his action $x \in [-1, 1]$ according to a mixed strategy σ (a probability distribution over $[-1, 1]$). Each player's utility function is a multivariate polynomial which only contains terms whose degree in x is at most some constant integer d . Then regardless of how everyone chooses their strategies, their expected utility will only depend on σ through the moments $\int x d\sigma(x), \int x^2 d\sigma(x), \dots, \int x^d d\sigma(x)$. Therefore player 1 can switch from σ to any other mixed strategy with the same first d moments without affecting game play, and we can think of the Nash equilibrium problem as one in which each player seeks to choose moments which correspond to an actual probability distribution and form a Nash equilibrium.

On the other hand there is no exact finite-dimensional characterization of the set of correlated equilibria in polynomial games; for a counterexample see Section 4. 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 $\int s_1^{k_1} \dots s_n^{k_n} d\pi$ 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 implicitly imposes constraints on the conditional distributions of the equilibrium

measure. A finite set of moments does not contain enough information about these conditional distributions to check the required constraints 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 $\zeta_i : C_i \rightarrow C_i$. This definition reduces to that of a correlated equilibrium when $\epsilon = 0$.

Compare this definition to the main characterization theorem for ϵ -correlated equilibria immediately below. This theorem shows that ϵ -correlated equilibria can be defined by integrating the utilities against any sufficiently rich class of test functions, instead of by using measurable departure functions. While this characterization does not have an obvious game theoretic interpretation, it allows us to compute correlated equilibria both algorithmically (Section 3) and analytically (Section 4).

Theorem 2.9. A probability measure $\pi \in \Delta(C)$ is an ϵ -correlated equilibrium of a continuous game if and only if for all players i , positive integers k , strategies $t_i^1, \dots, t_i^k \in C_i$, and functions $f_i^1, \dots, f_i^k : C_i \rightarrow [0, 1]$ in one of the classes

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

such that $\sum_{j=1}^k f_i^j(s_i) \leq 1$ for all $s_i \in C_i$, the inequality

$$\sum_{j=1}^k \int f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \leq \epsilon \tag{3}$$

holds.

To prove this, we need several approximation lemmas.

Lemma 2.10 (A special case of Lemma 20 in [28]). *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 players i and all Borel measurable simple functions $\xi_i : C_i \rightarrow C_i$.

Proof. The forward direction is trivial. To prove the reverse, first fix i . Then choose any measurable departure function ζ_i and let $\delta > 0$ be arbitrary. By the continuity of u_i and compactness of the strategy spaces there exists a finite open cover U^1, \dots, U^k of C_i such that $s_i, s'_i \in U^j$ implies $|u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})| < \delta$ for all $s_{-i} \in C_{-i}$ and $j = 1, \dots, k$. Fix any $s_i^j \in U^k$ for all j . Define a simple measurable departure function ξ_i by $\xi_i(s_i) = s_i^j$ where $j = \min\{l : \zeta_i(s_i) \in U^l\}$. Then $|u_i(\zeta_i(s_i), s_{-i}) - u_i(\xi_i(s_i), s_{-i})| < \delta$ for all $s \in C$, so

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

Letting δ go to zero completes the proof. \square

Lemma 2.11. *If C is a compact Hausdorff space, μ is a finite regular Borel measure on C , $f^1, \dots, f^k : C \rightarrow [0, 1]$ are measurable functions such that $\sum_{j=1}^k f^j \leq 1$, and $\delta > 0$, then there exist continuous functions $g^1, \dots, g^k : C \rightarrow [0, 1]$ such that $\mu(\{x \in C : f^j(x) \neq g^j(x)\}) < \delta$ for all j and $\sum_{j=1}^k g^j \leq 1$.*

Proof. We can apply Lusin's theorem which states exactly this result in the case $k = 1$ [24]. If $k > 1$, then we can apply the $k = 1$ case with $\frac{\delta}{k}$ in place of δ to each of the f^j . Call the resulting continuous functions \tilde{g}^j . Then $\mu(\{x \in C : f^j(x) \neq \tilde{g}^j(x) \text{ for some } j\}) < \delta$. But $\sum_{j=1}^k f^j \leq 1$, so $\mu(\{x \in C : \sum_{j=1}^k \tilde{g}^j(x) > 1\}) < \delta$. Let $h(x) = \max\{1, \sum_{j=1}^k \tilde{g}^j(x)\}$ so $h : C \rightarrow [1, \infty)$ is a continuous map. Define $g^j(x) = \frac{\tilde{g}^j(x)}{h(x)}$. Then the g^j are continuous, sum to at most unity, and are equal to the f^j wherever all of the \tilde{g}^j equal the f^j , i.e. except on a set of measure at most δ . \square

Lemma 2.12. *If $C \subset \mathbb{R}^d$ is compact, $f^1, \dots, f^k : C \rightarrow [0, 1]$ are continuous functions such that $\sum_{j=1}^k f^j \leq 1$, and $\delta > 0$, then there exist polynomials $p^1, \dots, p^k : C \rightarrow [0, 1]$ which are squares such that $|f^j(x) - p^j(x)| \leq \delta$ for all $x \in C$ and $\sum_{j=1}^k p^j \leq 1$.*

Proof. By the Stone-Weierstrass theorem, any continuous function on a compact subset of \mathbb{R}^d can be approximated by a polynomial arbitrarily well with respect to the sup norm. Approximating the square root of a nonnegative function f using this theorem and squaring the resulting polynomial shows that a nonnegative continuous function on a compact subset of \mathbb{R}^d can be approximated arbitrarily well by a square of a polynomial with respect to the sup norm.

Let \tilde{p}^j be a square of a polynomial which approximates f^j within $\frac{\delta}{2k}$ in the sup norm. Let $p^j = \frac{\tilde{p}^j}{1 + \frac{\delta}{2}}$. Then p^j is always within $\frac{\delta}{2}$ of \tilde{p}^j , hence p^j approximates f^j within δ in the sup norm. Furthermore for all $x \in C$ we have

$$\sum_{j=1}^k p^j(x) = \frac{1}{1 + \frac{\delta}{2}} \sum_{j=1}^k \tilde{p}^j(x) \leq \frac{1}{1 + \frac{\delta}{2}} \sum_{j=1}^k \left(f^j(x) + \frac{\delta}{2k} \right) \leq \frac{1}{1 + \frac{\delta}{2}} \left(1 + \frac{\delta}{2} \right) = 1. \quad \square$$

Proof of Theorem 2.9. First we prove that if π is an ϵ -correlated equilibrium then (3) holds in the case where the f_i^j are simple. We can choose a partition B_i^1, \dots, B_i^l of C_i into disjoint

measurable sets such that $f_i^j = \sum_{m=1}^l c_{jm} \chi_{B_i^m}$ where $c_{jm} \in [0, 1]$ and $\chi_{B_i^m}$ denotes the indicator function which is unity on B_i^m and zero elsewhere. Define a departure function $\zeta_i : C_i \rightarrow C_i$ piecewise on the B_i^m as follows. If

$$\int_{B_i^m \times C_{-i}} [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi$$

is nonnegative for some j define $\zeta_i(s_i) = t_i^j$ for all $s_i \in B_i^m$ where j is chosen to maximize the above integral. If the integral is negative for all j define $\zeta_i(s_i) = s_i$ for all $s_i \in B_i^m$. Then we have

$$\sum_{j=1}^k c_{jm} \int_{B_i^m \times C_{-i}} [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \leq \int_{B_i^m \times C_{-i}} [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)]$$

for all m . Summing over m and using the definition of an ϵ -correlated equilibrium yields (3) in the case where the f_i^j are simple.

Conversely suppose that (3) holds for all measurable simple functions. Let $\zeta_i : C_i \rightarrow C_i$ be any simple departure function. Let t_i^1, \dots, t_i^k be the range of ζ_i and $B_i^j = \zeta_i^{-1}(\{t_i^j\})$. Defining $f_i^j = \chi_{B_i^j}$, (3) says exactly that π satisfies the ϵ -correlated equilibrium condition for the departure function ζ_i . By Lemma 2.10, π is an ϵ -correlated equilibrium.

Any simple function can be written as a sum of weighted characteristic functions, so by making several of the t_i^j the same, we see that (3) for weighted characteristic functions is the same as (3) for simple measurable functions. If the inequality (3) holds for all simple measurable functions, a standard limiting argument proves that it holds for all measurable f_i^j , hence for all continuous f_i^j .

Suppose conversely that (3) holds for all continuous f_i^j . Fix any measurable f_i^j satisfying the assumptions of the theorem. Define a signed measure π_i^j on C_i by $\pi_i^j(B_i) = \int_{B_i \times C_{-i}} [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi$. Let $\mu_i = \sum_{j=1}^k |\pi_i^j|$ and fix any $\delta > 0$. Then by the Lemma 2.11 there exist continuous functions $g_i^j : C_i \rightarrow [0, 1]$ which sum to at most unity and equal the f_i^j except on a set of μ_i measure at most δ . Therefore

$$\left| \sum_{j=1}^k \int f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi - \sum_{j=1}^k \int g_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \right| \\ \leq \sum_{j=1}^k \int |f_i^j(s_i) - g_i^j(s_i)| d\pi_i^j \leq 2k\delta,$$

so

$$\sum_{j=1}^k \int f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \leq \epsilon + 2k\delta.$$

But δ was arbitrary, so (3) holds for all measurable f_i^j .

Finally assume $C_i \subset \mathbb{R}^{k_i}$ for some k_i . If (3) holds for all continuous f_i^j , then it holds for all squares of polynomials. Suppose conversely that it holds for all squares of polynomials.

Let f_i^j be any continuous functions satisfying the assumptions of the theorem and $\delta > 0$. Let p_i^j be polynomials squares which approximate the f_i^j within δ in the sup norm and satisfy the assumptions of the theorem, as provided by Lemma 2.12. Then

$$\begin{aligned} & \left| \sum_{j=1}^k \int f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi - \sum_{j=1}^k \int p_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \right| \\ & \leq \sum_{j=1}^k \int |f_i^j(s_i) - p_i^j(s_i)| d\pi_i^j \leq \delta \sum_{j=1}^k \int d\pi_i^j, \end{aligned}$$

so

$$\sum_{j=1}^k \int f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi \leq \epsilon + \delta \sum_{j=1}^k \int d\pi_i^j.$$

But δ was arbitrary and the integrals on the right are finite, so (3) holds for all continuous f_i^j . \square

Several simplifications occur when specializing Theorem 2.9 to the $\epsilon = 0$ case, yielding the following characterization. We will use the polynomial condition of this corollary in Section 3 to develop algorithms for computing (approximate) correlated equilibria. The characteristic function condition will allow us to compute extreme correlated equilibria of an example game in Section 4.

Corollary 2.13. *A joint measure π is a correlated equilibrium of a continuous game if and only if*

$$\int f_i(s_i) [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) \leq 0 \quad (4)$$

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. When $\epsilon = 0$ the $k = 1$ case of equation (3) implies the $k > 1$ cases. Furthermore $\epsilon = 0$ makes (3) homogeneous, so it is unaffected by positive scaling of the f_i^j , which allows us to drop the assumption $f_i \leq 1$. \square

Theorem 2.9 also has important topological implications for the structure of ϵ -correlated equilibria. Recall that the weak* topology on the set of probability distributions $\Delta(C)$ over a compact Hausdorff space is the weakest topology which makes $\pi \mapsto \int f d\pi$ a continuous functional whenever $f : C \rightarrow \mathbb{R}$ is a continuous function.

Corollary 2.14. *The set of ϵ -correlated equilibria of a continuous game is weak* compact.*

Proof. By the continuous test function condition in Theorem 2.9, the set of ϵ -correlated equilibria is defined by conditions of the form $\int f d\pi \leq \epsilon$ where f ranges over continuous functions of the form $\sum_{j=1}^k f_i^j(s_i) [u_i(t_i^j, s_{-i}) - u_i(s)]$. By definition this presents the set of ϵ -correlated equilibria as the intersection of a family of weak* closed sets. Hence the set of ϵ -correlated equilibria is a closed subset of $\Delta(C)$. But $\Delta(C)$ is compact by the Banach-Alaoglu theorem [25], so the set of ϵ -correlated equilibria is compact. \square

Corollary 2.15. *If π^k is a sequence of ϵ^k -correlated equilibria and $\epsilon^k \rightarrow 0$, then some π^k is a correlated equilibrium or the sequence π^k has a limit point which is a correlated equilibrium.*

Proof. Assume no π^k is a correlated equilibrium. Then since $\epsilon^k \rightarrow 0$, the sequence π^k contains infinitely many points. The space $\Delta(C)$ with the weak* topology is compact by the Banach-Alaoglu theorem [25], hence any infinite set has a limit point. Let $\pi \in \Delta(C)$ be a limit point of the sequence π^k . For any $\epsilon > 0$ there exists k_0 such that for all $k \geq k_0$, π^k is an ϵ -correlated equilibrium. The set $\Delta(C)$ is Hausdorff [25], so π is also a limit point of $\{\pi^k\}_{k \geq k_0}$. Since the set of ϵ -correlated equilibria is compact by Corollary 2.14, the limit point π must be an ϵ -correlated equilibrium for all $\epsilon > 0$, i.e. a correlated equilibrium. \square

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

Theorem 2.16. *A probability measure $\pi \in \Delta(\tilde{C})$, where $\tilde{C} = \prod_{j \in I} \tilde{C}_j$ 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*

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

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

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

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

3 Computing correlated equilibria

We focus in this section on developing algorithms that can compute approximate correlated equilibria with arbitrary accuracy. We consider three types of algorithms, which we will illustrate in turn using the example below.

Example 3.1. Consider the polynomial game with two players, x and y , each choosing their strategies from the interval $C_x = C_y = [-1, 1]$. Their utilities are given by

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

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). 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 $(\int u_x d\pi, \int u_y d\pi)$.

3.1 Static Discretization Methods

The techniques in this subsection are general enough to apply to arbitrary continuous games with finitely many players, 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 $\tilde{C} = \prod_{i=1}^n \tilde{C}_i$ 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} . 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. *Consider a continuous game with finitely many players, strategy sets C_i , and payoffs u_i . For any $\epsilon > 0$, there exists a finite open cover $U_i^1, \dots, U_i^{l_i}$ of C_i such that if*

$\tilde{C}_i \subseteq C_i$ is a finite set chosen to contain at least one point from each U_i^l , then all correlated equilibria of the finite 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, so for any $\epsilon > 0$ we can choose a finite open cover $U_i^1, \dots, U_i^{l_i}$ such that if s_i varies within one of the U_i^l and $s_{-i} \in C_{-i}$ is held fixed, the value of 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,

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

for all i and all $s_i, t_i \in \tilde{C}_i$. Any $t_i \in C_i$ belongs to the same U_i^l as some $\tilde{t}_i \in \tilde{C}_i$, so

$$\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(\tilde{t}_i, s_{-i}) - u_i(s) + \epsilon] \leq \epsilon \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) = \epsilon.$$

Therefore the assumptions of Theorem 2.16 are satisfied with $\epsilon_{i,s_i} = \epsilon \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s)$. \square

Combined with Corollary 2.15 and the fact that all finite games have correlated equilibria [12], this proposition shows that any continuous game with finitely many players has a correlated equilibrium (the case of an arbitrary set of players has appeared in [12]) which can be computed as a limit point of a sequence of correlated equilibria of sampled games when the discretization gets finer and finer. The proof also shows that if the utilities are Lipschitz functions, such as polynomials, then the U_i^l can in fact be chosen to be balls with radius proportional to ϵ . 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 $[-1, 1]$. In this case $\epsilon = O\left(\frac{1}{d}\right)$ where $d = \max_i |\tilde{C}_i|$.

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.

In fact we can improve this convergence rate to $\epsilon = O\left(\frac{1}{d^2}\right)$ if we include the endpoints ± 1 in \tilde{C}_i as well and assume that the utilities have bounded second derivatives. This fact is based on the following technical lemma.

Lemma 3.3. *Let $x_0 > 0$. If $f(\pm x_0) \leq 0$ and $\frac{d^2 f}{dx^2} \geq -M$ on $[-x_0, x_0]$, then $f(x) \leq \frac{Mx_0^2}{2}$ for all $x \in [-x_0, x_0]$. Furthermore this bound is achieved by $f(x) = \frac{M}{2}(x_0 - x)(x_0 + x)$ at $x = 0$.*

Proof. Fix $x_1 \in \arg \max_{[-x_0, x_0]} f(x)$. Replacing $f(x)$ by $f(-x)$ if necessary we can assume that $x_1 \geq 0$. Then $\frac{df}{dx}(x_1) = 0$, so the function $g(x) = f(x_1 + \frac{x_0 - x_1}{x_0}|x|)$ satisfies all the




Figure 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 $\Theta\left(\frac{1}{d}\right)$.

assumptions of the theorem (in particular it is twice differentiable despite the $|x|$ term) and achieves its maximum value of $f(x_1)$ at $x = 0$.

Therefore to prove the lemma it suffices to prove that $f(0) \leq \frac{x_0^2}{2}$ under the additional assumption that $\frac{df}{dx}(0) = 0$. By assumption $\frac{d^2f}{dx^2} \geq -M$. Integrating from 0 to x , we have $\frac{df}{dx}(x) \geq -Mx$. Integrating again from 0 to x_0 yields $f(0) - \frac{Mx_0^2}{2} \leq f(x_0) \leq 0$. \square

Corollary 3.4. *If $f : [-1, 1] \rightarrow \mathbb{R}$ is a function whose second derivatives is bounded below by $-M$ such that f is nonpositive at $d + 1$ evenly spaced points from -1 to 1 inclusive, then $\max_{-1 \leq x \leq 1} f(x) \leq \frac{M}{2d^2}$.*

Proof. Apply Lemma 3.3 between each pair of adjacent points where f is nonpositive. \square

The extremal example in Lemma 3.3 makes it clear that we cannot achieve a tighter bound on the values of f by constraining higher derivatives. Furthermore, if we do not constrain f to be nonpositive at the endpoints ± 1 , then we can produce a linear f with $\max_{-1 \leq x \leq 1} f(x) = \Theta\left(\frac{1}{d}\right)$ even if we constrain the first derivative of f . This weaker convergence bound leads to $\epsilon = \Theta\left(\frac{1}{d}\right)$ in the example above.

Corollary 3.5. *Consider a continuous game with strategy spaces $C_i = [-1, 1]$ for all players and utilities such that $\frac{\partial^2 u_i}{\partial s_i^2} \geq -M$ for some M . Let \tilde{C}_i be the set which consists of $d + 1$ points equally spaced between -1 and 1 inclusive for all i . Then a correlated equilibrium of the finite game with strategy spaces \tilde{C}_i and utilities $u_i|_{\tilde{C}_i}$ is an $\frac{M}{2d^2}$ -correlated equilibrium of the original game.*

Proof. Let π be such a correlated equilibrium. Then by assumption the function:

$$f_{i,s_i}(t_i) = \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)]$$

is zero for $t_i \in \tilde{C}_i$ and has second derivative bounded below by $M \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s)$. By Corollary 3.4, $\epsilon_{i,s_i} = \max_{-1 \leq t_i \leq 1} f_{i,s_i}(t_i) \leq \frac{M}{2d^2} \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s)$. Therefore $\epsilon = \max_i \sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} \leq \frac{M}{2d^2}$ and Theorem 2.16 completes the proof. \square

3.2 Adaptive Discretization Methods

3.2.1 A family of convergent adaptive discretization algorithms

In this section we consider continuous games with finitely many players and provide two algorithms (the second is in fact a parametrized family of algorithms which generalizes the first) to compute a sequence of ϵ^k -correlated equilibria such that $\lim_{k \rightarrow \infty} \epsilon^k = 0$. By Corollary 2.15 any limit point of this sequence is a correlated equilibrium. We will show that for polynomial games these algorithms can be implemented efficiently using semidefinite programming.

Informally, these algorithms work as follows. Each iteration k begins with a finite set $\tilde{C}_i^k \subseteq C_i$ of strategies which each player i is allowed to play with positive probability in that iteration; the initial choice of this set at iteration $k = 0$ is arbitrary. We then compute the “best” ϵ -correlated equilibrium in which players are restricted to use only these strategies, i.e., the one which minimizes ϵ (subject to some extra technical conditions needed to ensure convergence).

Given the optimal objective value ϵ^k and optimal probability distribution π^k , there is some player i who can improve his payoff by ϵ^k if he switches from his recommended strategies to certain other strategies. We interpret these other strategies as good choices for that player to use to help make ϵ^k smaller in later iterations k . Therefore we add these strategies to \tilde{C}_i^k to get \tilde{C}_i^{k+1} and repeat this process for iteration $k + 1$.

Algorithm 3.6. Fix a continuous game with finitely many players. Let $k = 0$ and for each player fix a finite subset $\tilde{C}_i^0 \subseteq C_i$.

- Let π^k be an ϵ^k -correlated equilibrium of the game having minimal ϵ^k subject to two extra conditions. First, π^k must be supported on \tilde{C}^k . Second, we require that π^k be an exact correlated equilibrium of the finite game induced when deviations from the recommended strategies are restricted to the set \tilde{C}^k , i.e. when we replace the condition $t_i \in C_i$ in Definition 2.8 with $t_i \in \tilde{C}_i^k$.

That is to say, let ϵ^k be the optimal value of the following optimization problem, and

π^k be an optimal assignment to the decision variables.

$$\begin{aligned}
& \text{minimize } \epsilon \\
& \text{subject to} \\
& \sum_{s_{-i} \in \tilde{C}_{-i}^k} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0 \quad \text{for all } i \text{ and } s_i, t_i \in \tilde{C}_i^k \\
& \sum_{s_{-i} \in \tilde{C}_{-i}^k} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq \epsilon_{i,s_i} \quad \text{for all } i, s_i \in \tilde{C}_i^k \text{ and } t_i \in C_i \\
& \sum_{s_i \in \tilde{C}_i^k} \epsilon_{i,s_i} \leq \epsilon \quad \text{for all } i \\
& \pi(s) \geq 0 \quad \text{for all } s \in \tilde{C} \\
& \sum_{s \in \tilde{C}^k} \pi(s) = 1
\end{aligned}$$

- If $\epsilon^k = 0$, terminate.
- For each player i for whom $\sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} = \epsilon$, form \tilde{C}_i^{k+1} from \tilde{C}_i^k by adding in at least one strategy t_i which makes $\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] = \epsilon_{i,s_i}$ for each $s_i \in \tilde{C}_i^k$ such that $\epsilon_{i,s_i} > 0$.
- For all other players i , let $\tilde{C}_i^{k+1} = \tilde{C}_i^k$.
- Let $k = k + 1$ and repeat.

Note that all steps of this algorithm are well-defined. First, the optimization problem is feasible. To see this let π^k be any exact correlated equilibrium of the finite game with strategy spaces \tilde{C}_i^k and utilities u_i restricted to \tilde{C}^k ; such an equilibrium exists because all finite games have correlated equilibria [12]. The u_i are bounded on C (being continuous functions on a compact set), so by making ϵ and the ϵ_{i,s_i} large, we see that π^k is a feasible solution of the problem. Second, the optimal objective value is achieved by some π^k because the space of probability measures on \tilde{C}^k is compact and ϵ is bounded below by zero. Third, the set of $t_i \in C_i$ making the ϵ -correlated equilibrium constraints tight at the optimum is nonempty by optimality of π^k and continuity of u_i . This set consists only of strategies which are not in \tilde{C}_i^k because we have the constraint that the deviations in utility be nonpositive for $t_i \in \tilde{C}_i^k$.

To show that Algorithm 3.6 converges, we will view it as a member of the following family of algorithms with the parameters set to $\alpha = 0$ and $\beta = 1$. Varying these parameters corresponds to adding some slack in the exact correlated equilibrium constraints and allowing some degree of suboptimality in the choice of strategies added to \tilde{C}_i^k to form \tilde{C}_i^{k+1} . Such changes make little conceptual difference, but could be helpful in practice by making the optimization problem strictly feasible and allowing it to be solved to within a known fraction of the optimal objective value rather than all the way to optimality. We will prove that all algorithms in this family converge, that is, with these algorithms ϵ^k converges to zero in the limit.

Algorithm 3.7. Fix a continuous game with finitely many players and parameters $0 \leq \alpha < \beta \leq 1$. Let $k = 0$ and for each player fix a finite subset $\tilde{C}_i^0 \subseteq C_i$.

- Choose ϵ^k to be the smallest number for which there exists π^k such that:
 - π^k is a probability distribution supported on \tilde{C}^k ,
 - π^k is an ϵ^k -correlated equilibrium of the game,
 - π^k is not an ϵ -correlated equilibrium for any $\epsilon < \epsilon^k$,
 - π^k is an $\alpha\epsilon^k$ -correlated equilibrium of the game when strategy deviations are restricted to \tilde{C}^k (i.e., when the condition $t_i \in C_i$ is changed to $t_i \in \tilde{C}_i^k$ in Definition 2.8).
- If $\epsilon^k = 0$, terminate.
- For at least one value of i , form \tilde{C}_i^{k+1} from \tilde{C}_i^k by adding strategies $t_{i,s_i} \in C_i$ such that

$$\sum_{s \in \tilde{C}^k} \pi^k(s) [u_i(t_{i,s_i}, s_{-i}) - u_i(s)] \geq \beta\epsilon^k.$$

- For all other values of i , let $\tilde{C}_i^{k+1} = \tilde{C}_i^k$.
- Let $k = k + 1$ and repeat.

In this case it is not immediately obvious that the first step of the algorithm is well-defined, i.e. that a minimal ϵ^k exists. To see this note that we could choose π^k to be an exact correlated equilibrium when strategy deviations are restricted to \tilde{C}^k , so there exists a pair $(\pi^{k,1}, \epsilon^{k,1})$ satisfying the four conditions under the first bullet above. Choose some sequence $(\pi^{k,l}, \epsilon^{k,l})$, $l = 1, 2, \dots$, of pairs satisfying these conditions such that $\epsilon^k = \lim_{l \rightarrow \infty} \epsilon^{k,l}$ is the infimum over all ϵ^k values of pairs satisfying these conditions. Passing to a subsequence if necessary we can assume without loss of generality that the $\pi^{k,l}$ converge to some π^k . It is clear from the proof of Corollary 2.15 that π^k is an ϵ^k -correlated equilibrium supported on \tilde{C}^k which is an $\alpha\epsilon^k$ -correlated equilibrium when deviations are restricted to \tilde{C}^k .

From Theorem 2.16 we see that a small change in the probabilities $\pi(s)$ of a distribution supported on a finite set results in a correspondingly small change of the minimal ϵ for which that distribution is an ϵ -correlated equilibrium. Therefore π^k is not an ϵ -correlated equilibrium for any $\epsilon < \epsilon^k$. Note that this final step depends crucially on the fact that \tilde{C}^k is finite and fixed while l varies. Also note that this subtlety disappears if $\alpha = 0$ because in that case it wouldn't matter if the limiting distribution had a smaller ϵ value. It is clear that the remaining steps of the algorithm are well-defined.

Theorem 3.8. Fix a continuous game with finitely many players. Algorithms 3.6 and 3.7 converge in the sense that $\epsilon^k \rightarrow 0$.

Proof. Suppose not, so there exists $\epsilon > 0$ and infinitely many values of k such that $\epsilon^k \geq \epsilon$. For each i let B_i^1, \dots, B_i^l be a finite open cover of C_i such that $u_i(s_i, s_{-i}) - u_i(t_i, s_{-i}) \leq \frac{1}{2}(\beta - \alpha)\epsilon$ when s_i and t_i belong to the same set B_i^l and $s_{-i} \in C_{-i}$. Such a cover exists by the compactness of the C_i and the continuity of the u_i . There are finitely many sets B_i^l so there is some iteration k , which we can take to satisfy $\epsilon^k \geq \epsilon$, such that for all i all of the sets B_i^l which will ever contain an element of \tilde{C}_i^k at some iteration k already do.

Note that π^k is an $\alpha\epsilon^k$ -correlated equilibrium when strategy choices are restricted to \tilde{C}_i^k , and $\epsilon^k > 0$ so we have $\beta\epsilon^k > \alpha\epsilon^k$. By the minimality of ϵ^k , the set $\tilde{C}_i^{k+1} \setminus \tilde{C}_i^k$ is nonempty for some player i (that is to say, it is always possible to perform the third step of the algorithm). Choose such an i and let $t_{i,s_i} \in \tilde{C}_i^{k+1}$ satisfy

$$\sum_{s \in \tilde{C}^k} \pi^k(s) [u_i(t_{i,s_i}, s_{-i}) - u_i(s)] \geq \beta\epsilon^k$$

for all $s_i \in \tilde{C}_i^k$. By assumption, for any choice of $r_{i,s_i} \in \tilde{C}_i^k$ we have

$$\sum_{s \in \tilde{C}^k} \pi^k(s) [u_i(r_{i,s_i}, s_{-i}) - u_i(s)] \leq \alpha\epsilon^k,$$

so

$$\sum_{s \in \tilde{C}^k} \pi^k(s) [u_i(t_{i,s_i}, s_{-i}) - u_i(r_{i,s_i}, s_{-i})] \geq (\beta - \alpha)\epsilon^k.$$

By construction of k , we can choose $r_{i,s_i} \in \tilde{C}_i^k$ to lie in the same set B_i^l as t_{i,s_i} for each $s_i \in \tilde{C}_i^k$. Thus

$$\begin{aligned} (\beta - \alpha)\epsilon &\leq (\beta - \alpha)\epsilon^k \\ &\leq \left| \sum_{s \in \tilde{C}^k} \pi^k(s) [u_i(t_{i,s_i}, s_{-i}) - u_i(r_{i,s_i}, s_{-i})] \right| \\ &\leq \sum_{s \in \tilde{C}^k} \pi^k(s) |u_i(t_{i,s_i}, s_{-i}) - u_i(r_{i,s_i}, s_{-i})| \\ &\leq \sum_{s \in \tilde{C}^k} \pi^k(s) \frac{(\beta - \alpha)\epsilon}{2} = \frac{(\beta - \alpha)\epsilon}{2}, \end{aligned}$$

a contradiction. □

Now we will illustrate Algorithm 3.6 on two examples.

Example 3.1 (continued). In Figure 2 we illustrate Algorithm 3.6 initialized with $\tilde{C}_x^0 = \tilde{C}_y^0 = \{0\}$. In this case convergence is obtained in three iterations, significantly faster than the static discretization method. The resulting strategy sets were $\tilde{C}_x^2 = \tilde{C}_y^2 = \{0, 1\}$.




Figure 2: Convergence of Algorithm 3.6 (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.

| k | ϵ^k | $\tilde{C}_x^k \setminus \tilde{C}_x^{k-1}$ | $\tilde{C}_y^k \setminus \tilde{C}_y^{k-1}$ | $\tilde{C}_z^k \setminus \tilde{C}_z^{k-1}$ |
|-----|--------------|---|---|---|
| 0 | 0.99 | {0} | {0} | {0} |
| 1 | 4.16 | | | {0.89} |
| 2 | 5.76 | {-1} | | |
| 3 | 0.57 | | {1} | |
| 4 | 0.28 | {0.53} | | {0.50, 0.63} |
| 5 | 0.16 | | {0.49, 0.70} | |
| 6 | 10^{-7} | | {-1, 0.60} | {-0.60, 0.47} |

Table 2: Output of Algorithm 3.6 on a three player polynomial game with utilities of degree 4 and randomly chosen coefficients.

Example 3.9. For a more complex illustration, we consider a polynomial game with three players, choosing strategies $x, y,$ and $z \in [-1, 1]$. The utilities were chosen to be polynomials with terms up to degree 4 in all the variables and the coefficients were chosen independently according to a normal distribution with zero mean and unit variance (their actual values are omitted). Algorithm 3.6 proceeds as in Table 2, which shows the value of ϵ^k and the new strategies added on each iteration. The terminal probability distribution π^6 does not display any obvious structure; in particular it is not a Nash equilibrium (product distribution).

3.2.2 Implementing these algorithms with semidefinite programs

To implement these algorithms for polynomial games, we must be able to do two things. First, we need to solve optimization problems with finitely many decision variables, linear objective functions and two types of constraints: nonnegativity constraints on linear functionals of the decision variables, and nonnegativity constraints on univariate polynomials whose coefficients are linear functionals of the decision variables. That is to say, we must be able to handle constraints of the form $p(t) \geq 0$ for all $t \in [-1, 1]$, where the coefficients of the polynomial p are linear in the decision variables. Second, we need to extract values of t for which such polynomial inequalities are tight at the optimum.

Both of these tasks can be done simultaneously by casting the problem as a **semidefinite program (SDP)**. For an overview of semidefinite programs and a summary of the necessary results (both of which are classical), see the appendix.

In the optimization problem in Algorithm 3.6 we have a finite number of univariate polynomials in t_i whose coefficients are linear in the decision variables $\pi(s)$ and ϵ_{i,s_i} . We wish to constrain these coefficients to allow only polynomials which are nonnegative for all $t_i \in [-1, 1]$. By Propositions A.3 and A.4 in the appendix this is the same as asking that these coefficients equal certain linear functions of matrices (i.e., sums along antidiagonals) which are constrained to be symmetric and positive semidefinite. Therefore we can write this optimization problem as a semidefinite program.

As a special case of convex programs, semidefinite programs have a rich duality theory which is useful for theoretical and computational purposes. In particular, SDP solvers keep

| | | | |
|-----|-----|-----|-----|
| | a | b | c |
| a | 0 | 1 | 0 |
| b | 1 | 5 | 7 |
| c | 0 | 7 | 0 |

Table 3: A finite symmetric game with identical utilities for which Algorithm 3.7 with $\alpha = \beta = 1$ does not converge when started with strategy sets $\tilde{C}_1^0 = \tilde{C}_2^0 = \{a\}$.

track of feasible primal and dual solutions in order to determine when optimality is reached. It can be shown that the dual data obtained by an SDP solver run on this optimization problem will encode the values of t_i making the polynomial inequalities tight at the optimum [20].

The process of generating an SDP from the optimization problem in the algorithms above, solving it, and extracting an optimal solution along with t_i values from the dual can all be automated. We have done so using the SOSTOOLS MATLAB toolbox for the pre- and post-processing and SeDuMi for solving the semidefinite programs efficiently [22, 29].

3.2.3 A nonconvergent limiting case

Note that in the algorithms above the convergence of the sequence ϵ^k is not necessarily monotone. If we were to let $\alpha = \beta$ (a case we did not allow above), the sequence would become monotone nonincreasing. If we were to furthermore fix $\alpha = \beta = 1$, then the condition that π be an exact (or $\alpha\epsilon^k$ -) correlated equilibrium when deviations are restricted to \tilde{C}_i^k would become redundant and could be removed.

These changes would simplify the behavior of Algorithm 3.7 conceptually as well as reducing the size of the SDP solved at each iteration, so we would like to adopt them if possible. However, the resulting algorithm may not converge, in the sense that ϵ^k may remain bounded away from zero.

Example 3.10. Consider the game shown in Table 3, which is symmetric and has identical utilities for both players. Let $\tilde{C}_1^0 = \tilde{C}_2^0 = \{a\}$ and apply Algorithm 3.6, but remove the condition that π^k be an exact correlated equilibrium when deviations are restricted to \tilde{C}_i^k . The only probability distribution supported on \tilde{C}^0 is $\delta_{(a,a)}$ which has an objective value of $\epsilon^0 = 1$. It is easy to see that \tilde{C}_i^1 is formed by simply adding each player's best response to a , so that $\tilde{C}_1^1 = \tilde{C}_2^1 = \{a, b\}$. We will argue that the unique solution to the optimization problem in iteration $k = 1$ is also $\delta_{(a,a)}$, hence $\tilde{C}_i^2 = \tilde{C}_i^1$ and the algorithm gets “stuck”, so that $\epsilon^k = \epsilon^0 = 1$ for all k .

For a probability distribution π , let π^T denote π with the players interchanged. By symmetry and convexity, if π is an optimal solution then so is $\frac{\pi + \pi^T}{2}$, which is a symmetric probability distribution with respect to the two players. Hence an optimal solution which is symmetric always exists. We will parametrize such distributions by $\pi = p\delta_{(a,a)} + q\delta_{(a,b)} + q\delta_{(b,a)} + r\delta_{(b,b)}$, where $p, q, r \geq 0$ and $p + 2q + r = 1$. Define a departure function $\zeta : C_1 \rightarrow C_1$ by $\zeta(a) = b$, $\zeta(b) = \zeta(c) = c$. Then for π to be an ϵ -correlated equilibrium it must satisfy

| k | ϵ^k | $\tilde{C}_x^k \setminus \tilde{C}_x^{k-1}$ | $\tilde{C}_y^k \setminus \tilde{C}_y^{k-1}$ |
|-----|--------------|---|---|
| 0 | 2 | $\{-1\}$ | $\{-1\}$ |
| 1 | 4 | $\{0\}$ | $\{0\}$ |
| 2 | 0 | $\{1\}$ | $\{1\}$ |

Table 4: Output of Algorithm 3.6 for a polynomial game on which Algorithm 3.7 with $\alpha = \beta = 1$ does not converge to a correlated equilibrium.

the following condition:

$$\begin{aligned} \epsilon &\geq \sum_{s_1 \in \tilde{C}_1^1} \epsilon_{1,s_1} \geq \sum_{s \in \tilde{C}^1} \pi(s) [u_1(\zeta(s_1), s_2) - u_1(s_1, s_2)] \\ &= p + 4q - q + 2r = 1 + q + r. \end{aligned}$$

We know we can achieve $\epsilon = 1$ with $p = 1$ (i.e. $\pi = \pi^0 = \delta_{(a,a)}$), and this inequality shows that if $p < 1$ then $\epsilon > 1$. Therefore the minimal ϵ value in iteration $k = 1$ is unity and is achieved by $\pi = \delta_{(a,a)}$. Furthermore we have shown that this is the unique symmetric probability distribution which achieves the minimal value of ϵ . Hence any other (not necessarily symmetric) optimal solution $\hat{\pi}$ satisfies $\frac{\hat{\pi} + \hat{\pi}^T}{2} = \delta_{(a,a)}$. But $\delta_{(a,a)}$ is an extreme point of the convex set of probability distributions on \tilde{C}^1 , so we must in fact have $\hat{\pi} = \delta_{(a,a)}$. Therefore $\pi^1 = \pi^0 = \delta_{(a,a)}$ is the unique optimal solution on iteration $k = 1$, so the procedure must get stuck as claimed. That is, $\tilde{C}_i^k = \{a, b\}$ and $\epsilon^k = 1$ for all $k \geq 1$.

The same behavior can occur in polynomial games, as can be shown by “embedding” the above finite game in a polynomial game. For example, we can take $C_x = C_y = [-1, 1]$ and

$$\begin{aligned} u_x(x, y) = u_y(x, y) &= (1 - x^2)(3y^2 + 6y + 5) \\ &\quad + (1 - y^2)(3x^2 + 6x + 5). \end{aligned}$$

Then if $\tilde{C}_x^0 = \tilde{C}_y^0 = \{-1\}$ the same analysis as above shows that $\tilde{C}_x^k = \tilde{C}_y^k = \{-1, 0\}$ and $\epsilon^k = 2$ for all $k \geq 1$.

Example 3.11. If we run Algorithm 3.6 on this polynomial game, the iterations proceed as in Table 4. The correlated equilibrium obtained in iteration 2 is

$$\begin{aligned} \pi^2 &= 0.4922\delta(x = 0, y = 1) + 0.4922\delta(x = 1, y = 0) \\ &\quad + 0.0156\delta(x = 1, y = 1), \end{aligned}$$

i.e., a probability of 0.4922 is assigned to each of the outcomes $(x, y) = (0, 1)$ and $(x, y) = (1, 0)$ and a probability of 0.0156 is assigned to $(x, y) = (1, 1)$.

3.3 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 Corollary 2.13. In particular, a measure π on C is a correlated equilibrium if and only if

$$\int p^2(s_i) [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) \leq 0 \quad (5)$$

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

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

then (5) 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 $c \in \mathbb{R}^{d+1}$ and $t_i \in C_i$, i.e. if and only if $M_i^d(t_i)$ is negative semidefinite for all $t_i \in C_i$.

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 semidefinite for all $t_i \in [-1, 1]$ for a given d we can use a semidefinite program (Proposition A.5 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 semidefinite conditions 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.

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 2nd 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 before being projected into utility space), so this proves that this example game has a unique correlated equilibrium.

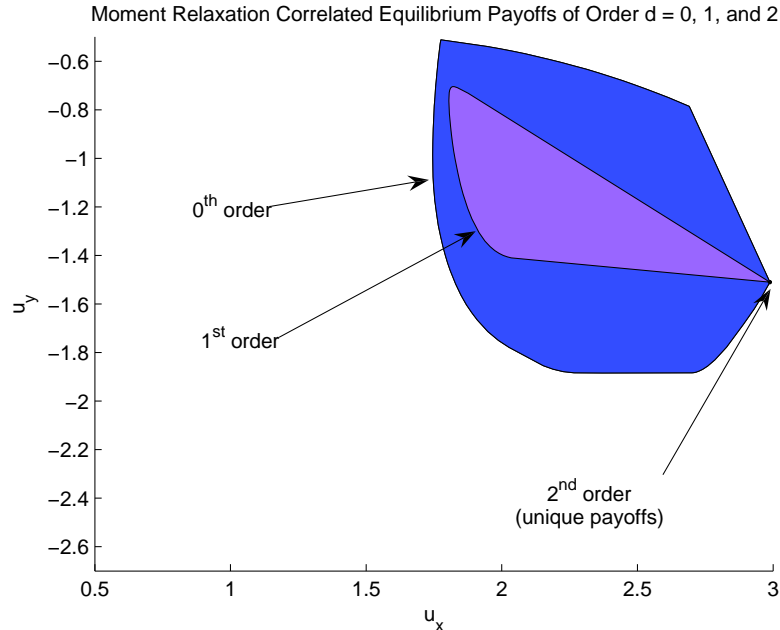


Figure 3: Semidefinite relaxations approximating the set of correlated equilibrium payoffs. The second order relaxation is a singleton, so this game has a unique correlated equilibrium payoff (and in fact a unique correlated equilibrium).

4 Structure of the set of correlated equilibria

We will demonstrate the complexity of the extreme points of the set of correlated equilibria through several example games. In particular we will show that there can be a superexponential separation between the number of extreme Nash and extreme correlated equilibria of a finite game, and that polynomial games can have extreme correlated equilibria which are not finitely supported. This implies that in general the condition that a measure μ be a correlated equilibrium of a given polynomial game cannot be defined in terms of conditions on finitely many joint moments of μ . In contrast, the condition that an n -tuple of mixed strategies be a Nash equilibrium of a polynomial game can always be expressed in terms of finitely many moments [27].

First we fix notation. Given a space S the set $\Delta(S)$ will denote the set of Borel probability measures on S (as above) and the set $\Delta^*(S)$ will denote the set of finite positive measures on S . In particular $\Delta(S)$ is the set of measures in $\Delta^*(S)$ with unit mass. If S is finite then $\Delta(S)$ is a simplex and $\Delta^*(S)$ is an orthant in $\mathbb{R}^{|S|}$. For any $p \in S$, define $\delta_p \in \Delta(S)$ to be the measure which assigns unit mass to the point p . Let $I = [-1, 1] \subset \mathbb{R}$.

We will focus on two related examples, one with finite strategy sets and one with infinite strategy sets. We will develop them in parallel by analyzing arbitrary games satisfying the following condition.

Assumption 4.1. The game is a zero-sum continuous game with two players, called x and

Table 5: Utilities for matching pennies

| (u_x, u_y) | $x = -1$ | $x = 1$ |
|--------------|-----------|-----------|
| $y = -1$ | $(1, -1)$ | $(-1, 1)$ |
| $y = 1$ | $(-1, 1)$ | $(1, -1)$ |

y . The strategy sets C_x and C_y are compact subsets of $I = [-1, 1]$, each of which contains a positive element and a negative element. Player x chooses a strategy $x \in C_x$ and player y chooses $y \in C_y$. The utility functions are $u_x(x, y) = xy = -u_y(x, y)$.

Example 4.2. Fix an integer $n > 0$. Let C_x and C_y each have $2n$ elements, n of which are positive and n of which are negative. If we take $n = 1$ and $C_x = C_y = \{-1, 1\}$ then we recover the matching pennies game¹, as shown in Table 5.

Example 4.3. Let $C_x = C_y = [-1, 1]$. Then the game is essentially the mixed extension of matching pennies. That is to say, suppose two players play matching pennies and choose their strategies independently, playing 1 with probabilities $p \in [0, 1]$ and $q \in [0, 1]$. Define the utilities for the mixed extension to be the expected utilities under this random choice of strategies. Letting $x = 2p - 1$ and $y = 2q - 1$, the utility to the first player is xy and the utility to the second player is $-xy$. Therefore this example is the mixed extension of matching pennies, up to an affine scaling of the strategies.

4.1 Extreme Nash equilibria

We will now show that the extreme points of the sets of Nash equilibria in games satisfying Assumption 4.1 are well-behaved. Since these games are zero-sum, the set of Nash equilibria can be viewed as a Cartesian product of two (weak*) compact convex sets, the sets of maximin and minimax strategies. By the Krein-Milman theorem, such sets can be completely characterized in terms of their extreme points [25].

We define Nash equilibria in two-player games, which will be sufficient for our purposes, as well as the standard notions of extreme point and extreme ray from convex analysis.

Definition 4.4. A **Nash equilibrium** is a pair $(\sigma, \tau) \in \Delta(C_x) \times \Delta(C_y)$ such that $u_x(x, \tau) \leq u_x(\sigma, \tau)$ for all $x \in C_x$ and $u_y(\sigma, y) \leq u_y(\sigma, \tau)$ for all $y \in C_y$.

Definition 4.5. A point x in a convex set K is an **extreme point** if $x = \lambda y + (1 - \lambda)z$ for $y, z \in K$ and $\lambda \in [0, 1]$ implies $x = y = z$.

Definition 4.6. A convex set K such that $x \in K$ and $\lambda \geq 0$ implies $\lambda x \in K$ is called a **convex cone**. A point $x \neq 0$ is an **extreme ray** of the convex cone K if $x = y + z$ and $y, z \in K$ implies that y or z is a scalar multiple of x .

¹By inspection of the utilities we can see that for any C_x and C_y with at least two points, the rank of this game in the sense of [27] is $(1, 1)$ (and in fact also in the stronger sense of Theorem 3.3 of that paper). The notion of the rank of a game is related to the rank of the payoff matrices and will not play a significant role in this paper; we merely wish to note that under this definition of complexity of payoffs the games we consider are extremely simple.

The Nash equilibria of games satisfying Assumption 4.1 take the following particularly simple form.

Proposition 4.7. *A pair $(\sigma, \tau) \in \Delta(C_x) \times \Delta(C_y)$ is a Nash equilibrium of a game satisfying Assumption 4.1 if and only if $\int x d\sigma(x) = \int y d\tau(y) = 0$.*

Proof. If $\int x d\sigma(x) = 0$ then $u_y(\sigma, y) = 0$ for all $y \in C_y$, so any $\tau \in \Delta(C_y)$ is a best response to σ . If $\int y d\tau(y) = 0$ as well then σ is also a best response to τ , so (σ, τ) is a Nash equilibrium.

Suppose for a contradiction that there exists a Nash equilibrium (σ, τ) such that $\int x d\sigma(x) > 0$; the other cases are similar. Player y must play a best response, so $\int y d\tau(y) < 0$, which is possible by assumption. Player x plays a best response to that, so $\int x d\sigma(x) < 0$, a contradiction. \square

We introduce the notion of extreme Nash equilibrium in the context of zero-sum games. For an extension of this definition to two-player non-zero sum finite games and a proof that extreme Nash equilibria are always extreme points of the set of correlated equilibria in this setting, see [5] or [10].

Definition 4.8. A Nash equilibrium $(\sigma, \tau) \in \Delta(C_x) \times \Delta(C_y)$ is **extreme** if σ and τ are extreme points of the maximin ($\{\sigma \in \Delta(C_x) \mid \int x d\sigma(x) = 0\}$) and minimax ($\{\tau \in \Delta(C_y) \mid \int y d\tau(y) = 0\}$) sets, respectively.

Proposition 4.9. *Consider a game satisfying Assumption 4.1. A pair $(\sigma, \tau) \in \Delta(C_x) \times \Delta(C_y)$ is an extreme Nash equilibrium if and only if σ and τ are each either δ_0 or of the form $\alpha\delta_u + \beta\delta_v$ where $u < 0$, $v, \alpha, \beta > 0$, $\alpha + \beta = 1$, and $\alpha u + \beta v = 0$.*

Proof. By Proposition 4.7 we must show that these distributions are the extreme points of the set of probability distributions having zero mean. Since δ_0 is an extreme point of the set of probability distributions, it must be an extreme point of the subset which has zero mean. To see that $\alpha\delta_u + \beta\delta_v$ is also an extreme point, suppose we could write it as a convex combination of two other probability distributions with zero mean. The condition that both be positive measures implies that both must be of the form $\alpha'\delta_u + \beta'\delta_v$. But α and β as specified above are the unique coefficients which make this be a probability measure with zero mean. Therefore $\alpha' = \alpha$ and $\beta' = \beta$, so $\alpha\delta_u + \beta\delta_v$ cannot be written as a nontrivial convex combination of probability distributions with zero mean, so it is an extreme point.

Suppose σ were an extreme point which was not of one of these types. Then σ could not be supported on one or two points, so either $[0, 1]$ or $[-1, 0)$ could be partitioned into two sets of positive measure. We will only treat the first case; the second is similar. Let $[0, 1] = A \cup B$ where $A \cap B = \emptyset$ and $\sigma(A), \sigma(B) > 0$. Since σ has zero mean we must have $\sigma([-1, 0)) > 0$ as well.

For a set D we define the restriction measure $\sigma|_D$ by $\sigma|_D(C) = \sigma(D \cap C)$ for all C . Then $\sigma = \sigma|_A + \sigma|_B + \sigma|_{[-1, 0)}$. Let $a = \int_A x d\sigma(x)$, $b = \int_B x d\sigma(x)$, and $c = \int_{[-1, 0)} x d\sigma(x)$. Since $\sigma([-1, 0)) > 0$ and x is less than zero everywhere on $[-1, 0)$, we must have $c < 0$. By assumption $a + b + c = 0$. Therefore we can write:

$$\sigma = \left(\sigma|_A + \frac{a}{|c|} \sigma|_{[-1, 0)} \right) + \left(\sigma|_B + \frac{b}{|c|} \sigma|_{[-1, 0)} \right)$$

Being an extreme point of the set of probability measures with zero mean, σ must be an extreme ray of the set of positive measures with first moment equal to zero. But this means that we cannot write $\sigma = \sigma_1 + \sigma_2$ where the σ_i are positive measures with zero first moment unless σ_i is a multiple of σ . Neither of the measures in parentheses above is a multiple of σ , so we have a contradiction. \square

We illustrate this proposition on both examples introduced at the beginning of this section.

Example 4.2 (cont'd). In this case neither C_x nor C_y contains zero, so the only extreme Nash equilibria are those in which σ and τ are of the form $\alpha\delta_u + \beta\delta_v$ for $u < 0$ and $v > 0$. For any choice of u and v there are unique α and β satisfying the conditions of Proposition 4.9. There are n possible choices for each of u and v for each of the two players, so there are n^4 extreme Nash equilibria.

Example 4.3 (cont'd). Since $C_x = C_y = [-1, 1]$, there are infinitely many extreme Nash equilibria in this case. However, they are all finitely supported and the size of their support is always either one or two. Furthermore the condition that (σ, τ) be a Nash equilibrium is equivalent to both having zero mean. This illustrates the general facts that in games with polynomial utility functions the Nash equilibrium conditions only involve finitely many moments of σ and τ and the extreme Nash equilibria (when defined, i.e., for zero-sum games) have uniformly bounded support.

4.2 Extreme correlated equilibria

We will show that even in finite games, the number of extreme correlated equilibria can be many more than the number of extreme Nash equilibria. It makes sense to compare these because all extreme Nash equilibria of a two-player finite game are automatically extreme correlated equilibria [10]. Also, for rank $(1, 1)$ finite games we will show that the size of the support of extreme correlated equilibria can be arbitrarily large.

In the case of polynomial games we will show that even in the rank $(1, 1)$ case, there can be extreme correlated equilibria with arbitrarily large finite support and without finite support. This implies that the set of correlated equilibria cannot be characterized in terms of finitely many joint moments.

Note that all of the characterizations of exact correlated equilibria in Section 2 are in terms of homogeneous (that is, invariant under positive scaling) linear constraints on μ . The only requirement on μ that is not homogeneous is the probability measure condition $\mu(I \times I) = 1$. It will be convenient to consider measures $\mu \in \Delta^*(C_x \times C_y)$ satisfying all the correlated equilibrium conditions except for this normalization.

Definition 4.10. We refer to a measure $\mu \in \Delta(C_x \times C_y)$ satisfying the conditions of Definition 2.7 as a **proper correlated equilibrium** and a measure $\mu \in \Delta^*(C_x \times C_y)$ satisfying all the conditions except normalization as a **homogeneous correlated equilibrium**. When the distinction is irrelevant or clear from context we simply use the term correlated equilibrium to refer to either. In the context of homogeneous correlated equilibria the term

extreme will refer to extreme rays; for proper correlated equilibria it will refer to extreme points.

When $\mu \neq 0$ is a homogenous correlated equilibrium, $\frac{1}{\mu(I \times I)}\mu$ is a proper correlated equilibrium. The set of homogenous correlated equilibria is a convex cone. The extreme rays of this cone are exactly those measures which are positive multiples of the extreme points of the set of proper correlated equilibria.

The following proposition characterizes correlated equilibria of games satisfying Assumption 4.1 and is analogous to Proposition 4.7 for Nash equilibria. Note how the Nash equilibrium measures were characterized in terms of their moments but the correlated equilibria are not.

Proposition 4.11. *For a game satisfying Assumption 4.1 and a measure $\mu \in \Delta^*(C_x \times C_y)$, the following are equivalent:*

1. μ is a correlated equilibrium;
2. If $d\kappa = xyd\mu$ then the marginals

$$\kappa_x(A) := \int_{A \times I} xyd\mu(x, y) \quad \text{and} \quad \kappa_y(A) := \int_{I \times A} xyd\mu(x, y)$$

are both the zero measure, i.e., equal zero for all measurable $A \subseteq I$;

- 3.

$$\lambda_x(A) := \int_{A \times I} yd\mu(x, y) \quad \text{and} \quad \lambda_y(A) := \int_{I \times A} xd\mu(x, y)$$

are both the zero measure.

Proof. (1 \Rightarrow 2) We will consider only κ_x ; κ_y is similar. The characteristic function version of Corollary 2.13 with $f_i = \chi_I$ implies that

$$x' \int_{I \times I} yd\mu(x, y) \leq \int_{I \times I} xyd\mu(x, y) \leq y' \int_{I \times I} xd\mu(x, y)$$

for all $x' \in C_x, y' \in C_y$. By assumption it is possible to choose x' and y' either positive or negative, so $\int_{I \times I} xyd\mu(x, y) = 0$. Furthermore the same argument with any A implies that $\int_{A \times I} xyd\mu(x, y) \geq 0$. Therefore we have

$$0 = \int_{I \times I} xyd\mu(x, y) = \int_{A \times I} xyd\mu(x, y) + \int_{(I \setminus A) \times I} xyd\mu(x, y) \geq 0 + 0 = 0$$

for all A , so the inequality must be tight and we get $\int_{A \times I} xyd\mu(x, y) = 0$ for all A .

(2 \Rightarrow 3) Substituting this equation into the characteristic function version of Corollary 2.13 yields $x' \int_{A \times I} yd\mu(x, y) \leq 0$ for all $x' \in C_x$ and all measurable A . But we can choose x' to be positive or negative by assumption, so we must have $\int_{A \times I} yd\mu(x, y) = 0$ for all measurable $A \subseteq I$.

(3 \Rightarrow 1) Suppose we know that λ_x is the zero measure. We can approximate xy on $I \times I$ by functions of the form $\sum_{k=1}^r x_i y \chi_{A_i}(x)$, where the measurable sets A_i have small diameter and partition I , $x_i \in A_i$, and χ_{A_i} is the indicator function for A_i . By definition of the Lebesgue integral and the dominated convergence theorem [24] we have

$$\int_{A \times I} xy d\mu(x, y) = \int_A x d\lambda_x(x) = \int_A x d0(x) = 0.$$

Thus $\int_{A \times I} (x - x') y d\mu(x, y) \geq 0$ for all $x' \in C_x$ and all A . □

Proposition 4.12. *Fix a game satisfying Assumption 4.1. Let $k > 0$ be even and x_1, \dots, x_{2k} and y_1, \dots, y_{2k} be such that:*

1. $x_i \in C_x$ and $y_i \in C_y$ are all nonzero;
2. the sequences $x_1, x_3, \dots, x_{2k-1}$ and $y_1, y_3, \dots, y_{2k-1}$ alternate in sign;
3. $x_{2i} = x_{2i-1}$ and $y_{2i} = y_{2i+1}$ for all i when subscripts are interpreted mod $2k$.

Then $\mu = \sum_{i=1}^{2k} \frac{1}{|x_i y_i|} \delta_{(x_i, y_i)}$ is an extreme correlated equilibrium. Furthermore, all finitely supported extreme correlated equilibria whose support does not contain points with $x = 0$ or $y = 0$ are of this form.

Proof. To show that μ is a correlated equilibrium define $d\kappa(x, y) = xy d\mu(x, y)$. Then $\kappa = \sum_{i=1}^{2k} \text{sign}(x_i) \text{sign}(y_i) \delta_{(x_i, y_i)}$. Defining the projection κ_x as in Proposition 4.11, we have

$$\begin{aligned} \kappa_x &= \sum_{i=1}^{2k} \text{sign}(x_i) \text{sign}(y_i) \delta_{x_i} = \sum_{i=1}^k \text{sign}(x_{2i}) (\text{sign}(y_{2i}) + \text{sign}(y_{2i-1})) \delta_{x_{2i}} \\ &= \sum_{i=1}^k \text{sign}(x_{2i}) (0) \delta_{x_{2i}} = 0, \end{aligned}$$

because $x_{2i} = x_{2i-1}$ and y_{2i} differs in sign from y_{2i-1} by assumption. The same argument shows that $\kappa_y = 0$, so μ is a correlated equilibrium.

To see that μ is extreme, suppose $\mu = \mu' + \mu''$ where μ' and μ'' are correlated equilibria. Clearly $\mu' = \sum_{i=1}^{2k} \alpha_i \delta_{(x_i, y_i)}$ for some $\alpha_i \geq 0$. Define $d\kappa' = xy d\mu'(x, y)$, so $\kappa' = \sum_{i=1}^{2k} \alpha_i x_i y_i \delta_{(x_i, y_i)}$. By assumption

$$\kappa'_x = \sum_{i=1}^{2k} x_{2i} (\alpha_{2i-1} y_{2i-1} + \alpha_{2i} y_{2i}) \delta_{x_{2i}}$$

is the zero measure. Since the x_{2i} are distinct and nonzero we must have $\alpha_{2i-1} y_{2i-1} + \alpha_{2i} y_{2i} = 0$ for all i . Similarly since $\kappa'_y = 0$ we have $\alpha_{2i+1} x_{2i+1} + \alpha_{2i} x_{2i} = 0$ for all i (with subscripts interpreted mod $2k$).

The x_i and y_i are all nonzero, so fixing one α_i fixes all the others by these equations. That is to say, these equations have a unique solution up to multiplication by a scalar, so




Figure 4: The support of an extreme correlated equilibrium. In the notation of Proposition 4.12, $k = 1$, $x_1 = 0.4$, $x_3 = -0.6$, $y_1 = 0.2$, and $y_3 = -0.8$.

μ' is a positive scalar multiple of μ . But the splitting $\mu = \mu' + \mu''$ was arbitrary, so μ is extreme.

A similar argument shows that any finitely supported correlated equilibrium μ whose support does not contain any points with $x = 0$ or $y = 0$ can be written as $\mu = \mu' + \mu''$ where μ' is a correlated equilibrium and μ'' is a correlated equilibrium of the above form. Therefore μ cannot be extreme unless it is of this form. \square

Example 4.2 (cont'd). For some examples of extreme correlated equilibria of games of this type, see Figures 4 and 5.

To count the number of extreme correlated equilibria of this game for general n we must count the number of essentially different sequences of x_i and y_i of the type mentioned in Proposition 4.12. Fix k and let $k = 2r$ where $1 \leq r \leq n$. Note that cyclically shifting the sequences of x_i 's and y_i 's by two does not change μ , nor does reversing the sequence. Therefore we can assume without loss of generality that $x_1, y_1 > 0$. We then have n possible choices for x_1, y_1, x_3 , and y_3 , $n - 1$ possible choices for x_5, x_7, y_5 , and y_7 , etc., for a total of $\left(\frac{n!}{(n-r)!}\right)^4$ possible choices of the x_i and y_i . These will always be essentially different (i.e., give rise to different μ) unless we cyclically permute the sequences of x_i and y_i by some multiple of four, in which case the resulting sequence is essentially the same. The number of such cyclic permutations is r . Therefore the total number of extreme correlated equilibria is

$$e(n) = \sum_{r=1}^n \frac{1}{r} \left(\frac{n!}{(n-r)!} \right)^4.$$

We will see that $e(n) = \Theta\left(\frac{1}{n}(n!)^4\right)$. That is to say, $e(n)$ is asymptotically upper and lower bounded by a constant times $\frac{1}{n}(n!)^4$. The expression $\frac{1}{n}(n!)^4$ is just the final term in




Figure 5: The support of another extreme correlated equilibrium. In the notation of Proposition 4.12, $k = 2$, $x_1 = 0.4$, $x_3 = -0.4$, $x_5 = 0.6$, $x_7 = -0.6$, $y_1 = 0.6$, $y_3 = -0.4$, $y_5 = 0.4$, and $y_7 = -0.6$.

the summation for $e(n)$, so the lower bound is clear. Define

$$f(n) = \frac{e(n)}{\frac{1}{n}(n!)^4} = \sum_{s=0}^{n-1} \frac{n}{n-s} \cdot \frac{1}{(s!)^4}.$$

Then $f(n) \geq 1$ for all n . We will now show that $f(n)$ is also bounded above. Intuitively this is not surprising as the terms in the summation for $f(n)$ die off extremely fast as s grows.

For all $1 \leq s < n-1$ we have that the ratio of term $s+1$ in the summation to term s is:

$$\frac{\frac{n}{n-s-1} \cdot \frac{1}{((s+1)!)^4}}{\frac{n}{n-s} \cdot \frac{1}{(s!)^4}} = \frac{n-s}{n-s-1} \cdot \frac{1}{(s+1)^4} \leq \frac{1}{8},$$

so for $n > 1$ we can bound the sum by a geometric series:

$$f(n) - 1 = \sum_{s=1}^{n-1} \frac{n}{n-s} \cdot \frac{1}{(s!)^4} \leq \frac{n}{n-1} \sum_{t=0}^{\infty} \frac{1}{8^t} = \frac{8n}{7(n-1)} \leq \frac{16}{7}.$$

Therefore $1 \leq f(n) \leq \frac{23}{7}$ for all n , so $e(n) = \Theta\left(\frac{1}{n}(n!)^4\right)$ as claimed. Comparing this to the results of the previous section in which we saw that the number of extreme Nash equilibria of this game is n^4 , we see that in this case there is a super-exponential separation between the number of extreme Nash and the number of extreme correlated equilibria. This implies, for example, that computing all extreme correlated equilibria is not an efficient method for computing all extreme Nash equilibria, even though all extreme Nash equilibria are extreme correlated equilibria and recognizing whether an extreme correlated equilibrium is an extreme Nash equilibrium is easy. There are simply too many extreme correlated equilibria.

Next we will prove a more abstract version of Proposition 4.12 which includes certain extreme points which are not finitely supported. First we need a brief digression to ergodic theory. The first definition is the standard definition of compatibility between a measure and a transformation on a space. The second definition expresses one notion of what it means for a transformation to “mix up” a space – in this case that the space cannot be partitioned into two sets of positive measure which do not interact under the transformation. Then we state the main ergodic theorem and a corollary which we will apply to exhibit extreme correlated equilibria of games satisfying Assumption 4.1.

Definition 4.13. Given a measure $\mu \in \Delta^*(S)$ on a space S , a measurable function $g : S \rightarrow S$ is called **(μ -)measure preserving** if $\mu(g^{-1}(A)) = \mu(A)$ for all measurable $A \subseteq S$. Note that if g is invertible (in the measure theoretic sense that an almost everywhere inverse exists), then this is equivalent to the condition that $\mu(g(A)) = \mu(A)$ for all A .

Definition 4.14. Given a measure $\mu \in \Delta^*(S)$, a μ -measure preserving transformation g is called **ergodic** if $\mu(A \Delta g^{-1}(A)) = 0$ implies $\mu(A) = 0$ or $\mu(A) = \mu(S)$, where $A \Delta B$ denotes the symmetric difference $(A \setminus B) \cup (B \setminus A)$.

Example 4.15. Fix a finite set S and a permutation $g : S \rightarrow S$. Let μ be counting measure on S . Then g is measure preserving. Furthermore, a set T satisfies $\mu(g^{-1}(T) \Delta T) = 0$ if and only if $g^{-1}(T) = T$ if and only if T is a union of cycles of g . Therefore g is ergodic if and only if it consists of a single cycle.

Example 4.16. Fix $\alpha \in \mathbb{R}$. Let $S = [0, 1)$ and let μ be Lebesgue measure on S . Define $g : S \rightarrow S$ by $g(x) = (x + \alpha) \bmod 1 = (x + \alpha) - \lfloor x + \alpha \rfloor$. Then g is μ -measure preserving because Lebesgue measure is translation invariant. It can be shown that g is ergodic if and only if α is irrational. For a proof and more examples, see [26].

The following is one of the core theorems of ergodic theory. We will only use it to prove the corollary which follows, so it need not be read in detail. The proof can be found in any text on ergodic theory, e.g. [26].

Theorem 4.17 (Birkhoff’s ergodic theorem). *Fix a probability measure μ and a μ -measure preserving transformation g . Then for any $f \in \mathcal{L}^1(\mu)$:*

- $\tilde{f}(x) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(g^k(x))$ exists μ -almost everywhere,
- $\tilde{f} \in \mathcal{L}^1(\mu)$,
- $\int \tilde{f} d\mu = \int f d\mu$,
- $\tilde{f}(g(x)) = \tilde{f}(x)$ μ -almost everywhere, and
- if g is ergodic then $\tilde{f}(x) = \int f d\mu$ μ -almost everywhere.

Corollary 4.18. *Suppose μ and ν are probability measures such that ν is absolutely continuous with respect to μ . If a transformation g preserves both μ and ν and g is ergodic with respect to μ , then $\nu = \mu$.*

Proof. Fix any measurable set A . Let f be the indicator function for A , i.e. the function equal to unity on A and zero elsewhere. Applying Birkhoff's ergodic theorem to f and μ yields $\tilde{f}(x) = \mu(A)$ μ -almost everywhere. Since ν is absolutely continuous with respect to μ , $\tilde{f}(x) = \mu(A)$ ν -almost everywhere also. If we now apply Birkhoff's ergodic theorem to ν we get:

$$\nu(A) = \int f d\nu = \int \tilde{f} d\nu = \int \mu(A) d\nu = \mu(A). \quad \square$$

Proposition 4.19. *Fix measures ν_1, ν_2, ν_3 , and $\nu_4 \in \Delta^*((0, 1])$ and maps $f_i : (0, 1] \rightarrow (0, 1]$ such that $\nu_{i+1} = \nu_i \circ f_i^{-1}$ (interpreting subscripts mod 4). The portion of the measure μ in the i^{th} quadrant of $I \times I$ will be constructed in terms of f_i and ν_i . Define $j_i : (0, 1] \rightarrow I \times I$ by $j_1(x) = (x, f_1(x))$, $j_2(x) = (-f_2(x), x)$, $j_3(x) = (-x, -f_3(x))$, and $j_4(x) = (f_4(x), -x)$. Let $|\kappa| = \sum_{i=1}^4 \nu_i \circ j_i^{-1}$. If Assumption 4.1 is satisfied, $\text{supp}|\kappa| \subseteq C_x \times C_y$, and $\frac{1}{|xy|} \in \mathcal{L}^1(|\kappa|)$ then $d\mu = \frac{1}{|xy|} d|\kappa|$ is a correlated equilibrium.*

If in addition $f_4 \circ f_3 \circ f_2 \circ f_1 : (0, 1] \rightarrow (0, 1]$ is ergodic with respect to ν_1 , then μ is extreme.

Proof. First we must show that μ is a correlated equilibrium. It is a finite measure by the assumption $\frac{1}{|xy|} \in \mathcal{L}^1(|\kappa|)$. Define $g : I \times I \rightarrow I \times I$ as follows.

$$g(x, y) = \begin{cases} j_1(x) & \text{if } x > 0, y < 0 \\ j_2(y) & \text{if } x > 0, y > 0 \\ j_3(-x) & \text{if } x < 0, y > 0 \\ j_4(-y) & \text{if } x < 0, y < 0 \\ \text{arbitrary} & \text{otherwise} \end{cases}$$

The function g is $|\kappa|$ -measure preserving. To see this fix any measurable set $B \subseteq (0, 1] \times (0, 1]$. Let $A = j_1^{-1}(B)$. Then $|\kappa|(B) = |\kappa|(A \times (0, 1]) = \nu_1(A)$ by definition of $|\kappa|$. But $g^{-1}(B) = g^{-1}(A \times (0, 1]) = A \times [-1, 0)$, so $|\kappa|(g^{-1}(B)) = |\kappa|(A \times [-1, 0)) = \nu_4(j_4^{-1}(A \times [-1, 0))) = \nu_4(f_4^{-1}(A)) = \nu_1(A)$ by assumption. Therefore g is measure preserving for subsets of $(0, 1] \times (0, 1]$. The arguments for the other quadrants are similar and since g maps each quadrant into a different quadrant, g is measure preserving on its entire domain.

Using the terminology of Proposition 4.11, $d\kappa = \text{sign}(x) \text{sign}(y) d|\kappa|$. We have seen that $|\kappa|(A \times (0, 1]) = |\kappa|(A \times [-1, 0))$, so $\kappa(A \times (0, 1]) = -\kappa(A \times [-1, 0))$. Since $\kappa(A \times \{0\}) = 0$, we have $\kappa(A \times I) = 0$, or equivalently, $\kappa_x(A) = 0$. A similar argument implies $\kappa_x(A) = 0$ if $A \subseteq [-1, 0)$. Clearly $\kappa_x(\{0\}) = 0$ by definition of κ_x , so κ_x is the zero measure. A similar argument shows that κ_y is the zero measure, so μ is a correlated equilibrium by Proposition 4.11.

Now we will show via several steps that μ is extreme. Write $\mu = \mu_1 + \mu_2$ where the μ_i are nonzero correlated equilibria. Since these are all positive measures, the μ_i are absolutely continuous with respect to μ . Define $d|\kappa_i| = |xy| d\mu_i$.

Next we show that g is $|\kappa_i|$ -measure preserving. We will demonstrate this fact for $B \subseteq (0, 1] \times (0, 1]$. As above, we define $A = j_1^{-1}(B)$. Then $|\kappa_i|(B) = |\kappa_i|(A \times (0, 1])$ since

$(A \times (0, 1]) \triangle B$ has $|\kappa|$ measure zero and $|\kappa_i|$ is absolutely continuous with respect to $|\kappa|$. Furthermore, $|\kappa_i|(g^{-1}(B)) = |\kappa_i|(A \times [-1, 0))$. But μ_i is a correlated equilibrium so $\kappa_i(A \times (0, 1]) = -\kappa_i(A \times [-1, 0))$. Hence $|\kappa_i|(g^{-1}(B)) = |\kappa_i|(A \times [-1, 0)) = |\kappa_i|(A \times (0, 1]) = |\kappa_i|(B)$. Again, the proof is the same for B contained in other quadrants, so g is $|\kappa_i|$ -measure preserving.

For the second-to-last step we prove that g is ergodic with respect to $|\kappa|$. Suppose $B \subseteq I \times I$ is such that $|\kappa|(g^{-1}(B) \triangle B) = 0$. Let Q_i be the intersection of B with the i^{th} quadrant. Then $|\kappa|(g^{-1}(Q_{i+1}) \triangle Q_i) = 0$, so $|\kappa|(g^{-4}(Q_1) \triangle Q_1) = 0$. Let $A = j_1^{-1}(Q_1)$. Then $|\kappa|(g^{-4}(Q_1) \triangle Q_1) = \nu_1((f_4 \circ f_3 \circ f_2 \circ f_1)^{-1}(A) \triangle A) = 0$. By assumption the map $f_4 \circ f_3 \circ f_2 \circ f_1$ is ergodic, so $\nu_1(A) = 0$ or $\nu_1(A) = \nu_1((0, 1]) = |\kappa|((0, 1] \times (0, 1])$. Therefore $|\kappa|(Q_1) = \nu_1(A) = 0$ or $|\kappa|(Q_1) = |\kappa|((0, 1] \times (0, 1])$. In either case since g is $|\kappa|$ -measure preserving we get $|\kappa|(Q_i) = |\kappa|(Q_1)$ for all i . Therefore $|\kappa|(B) = 0$ or $|\kappa|(B) = |\kappa|(I \times I)$, so g is ergodic with respect to $|\kappa|$.

Normalizing $|\kappa|$ and $|\kappa_i|$ to be probability measures, we can apply Corollary 4.18 to obtain $|\mu_i| = \frac{|\kappa_i|(I \times I)}{|\kappa|(I \times I)} |\kappa|$. By assumption the set on which $|xy|$ is zero has μ measure zero. Therefore

$$d\mu_i = \frac{1}{|xy|} d|\kappa_i| = \frac{|\kappa_i|(I \times I)}{|\kappa|(I \times I)} \frac{1}{|xy|} d|\kappa| = \frac{|\kappa_i|(I \times I)}{|\kappa|(I \times I)} d\mu,$$

so $\mu_i = \frac{|\kappa_i|(I \times I)}{|\kappa|(I \times I)} \mu$ and μ is extreme. □

Above we have constructed μ and g so that g maps the quadrants counter-clockwise – quadrant 1 to quadrant 2, etc. However, the same argument would go through if g mapped the quadrants clockwise.

To view Proposition 4.12 as a special case of Proposition 4.19, let each ν_i be a uniform probability measure over a finite subset of $(0, 1]$. The function g is defined by $g(x_i, y_i) = (x_{i+1}, y_{i+1})$ and the f_i are defined to be compatible with this. The map $f_4 \circ f_3 \circ f_2 \circ f_1$ is a permutation on the support of ν_1 , which is precisely the positive values of x_i . By construction this permutation consists of a single cycle, hence it is ergodic.

Example 4.3 (cont'd). We can combine Example 4.16 and Proposition 4.19 to exhibit extreme points of the set of correlated equilibria for this game which are not finitely supported. Let $0 < a < b < 1$. Let ν_i be Lebesgue measure on $[a, b]$ for all i . Fix α such that $\frac{\alpha}{b-a}$ is irrational. Define $f_1 : [a, b] \rightarrow [a, b]$ by $f_1(x) = (x - a + \alpha \bmod (b - a)) + a$. This is just an affinely scaled version of Example 4.16 so f_1 is ν_i -measure preserving and ergodic. Define f_1 on $(0, 1] \setminus [a, b]$ arbitrarily, because that is a set of measure zero. Let $f_2, f_3, f_4 : (0, 1] \rightarrow (0, 1]$ be the identity. These data satisfy all the assumptions of Proposition 4.19. In particular, since $0 < a < b < 1$, xy is bounded away from zero on the support of $|\kappa|$. Therefore $\frac{1}{|xy|} \in \mathcal{L}^1(|\kappa|)$. Since ν_i is not finitely supported, μ is an extreme correlated equilibrium which is not finitely supported. The support of μ is shown in Figure 6 with parameters $a = 0.2$, $b = 0.8$, and $\alpha = \frac{1}{\sqrt{5}}$.

Definition 4.4. Given a compact Hausdorff space K we say that a set of measures $\mathcal{M} \subseteq \Delta(K)$




Figure 6: The support set of an extreme correlated equilibrium which is not finitely supported. Extremality of this equilibrium depends sensitively on the choices of endpoints for the line segments. In the case pictured there are segments connecting: $(0.2, -0.2)$ to $(0.8, -0.8)$; $(-0.2, -0.2)$ to $(-0.8, -0.8)$; $(-0.2, 0.2)$ to $(-0.8, 0.8)$; $(0.2, 0.2 + \frac{1}{\sqrt{5}})$ to $(0.8 - \frac{1}{\sqrt{5}}, 0.8)$; and $(0.8 - \frac{1}{\sqrt{5}}, 0.2)$ to $(0.8, 0.2 + \frac{1}{\sqrt{5}})$.

is **describable by moments** if there exists an integer d , continuous maps $g_1, \dots, g_d : K \rightarrow \mathbb{R}$, and a set $M \subseteq \mathbb{R}^d$ such that a measure μ is in \mathcal{M} if and only if $(\int g_1 d\mu, \dots, \int g_d d\mu) \in M$.

The results of [21] show that the maximin and minimax strategy sets of a two-player zero-sum polynomial game can always be described by moments. Introducing a similar notion for n -tuples of moments, the set of Nash equilibria can always be described by moments in any polynomial game [27]. However, combining this example with the following proposition we see that the set of correlated equilibria of a polynomial game cannot in general be described by moments.

Proposition 4.5. *Let $\mathcal{M} \subseteq \Delta^*(K)$ be a set of measures describable by moments. Then all extreme points of \mathcal{M} have finite support and this support is uniformly bounded by $d + 1$, where d is the integer associated with the description of \mathcal{M} by moments.*

Proof. Suppose there exists a measure $\mu \in \mathcal{M}$ which is extreme and supported on more than $d + 1$ points, so we can partition the domain of μ into $d + 2$ sets B_1, \dots, B_{d+2} of positive measure. Then μ is a convex combination of the measures $(d + 2)\mu|_{B_i}$, but only if all the weights in the convex combination are equal to $\frac{1}{d+2} > 0$. The moments of μ are a convex combination of the moments of these measures. Since there are d moments, there exist convex combinations of at most $d + 1$ of the measures $(d + 2)\mu|_{B_i}$ with the same moments as μ by Carathéodory's theorem. We can write μ as a strict convex combination of such measures, contradicting extremality of μ . \square

5 Conclusions

We have proven several new characterizations of correlated equilibria in continuous games, and applied them to compute correlated equilibria, both algorithmically and analytically. This has allowed us to prove that in general the set of correlated equilibria of a polynomial game cannot be described exactly in terms of finitely many moments. Nonetheless, a sample correlated equilibrium or the entire set of correlated equilibria can be approximated within an arbitrary degree of accuracy by a sequence of semidefinite programs.

These results leave several open questions. If we define a moment map to be any map of the form $\pi \mapsto (\int f_1 d\pi, \int f_2 d\pi, \dots, \int f_k d\pi)$, then we have shown that the set of correlated equilibria is not the inverse image of any set under any moment map. On the other hand, since moment maps are linear and weak* continuous, we know that the image of the set of correlated equilibria under any moment map is convex and compact. Supposing the utilities and the f_i are polynomials, is there anything more we can say about this image? In particular, is it semialgebraic (i.e., describable in terms of finitely many polynomial inequalities)?

For any continuous game, the set of correlated equilibria is nonempty, and this can be proven constructively as in [12]. Under the same assumptions we can prove the existence of a Nash equilibrium, but the proof is nonconstructive, or at least does not seem to give an efficient algorithm for constructing an equilibrium [27]. In the case of polynomial games, existence of a Nash equilibrium immediately gives existence of a finitely supported Nash equilibrium by Carathéodory's theorem, which is constructive [27]. Therefore there exists a finitely supported correlated equilibrium of any polynomial game. Is there a constructive way to prove this fact directly, without going through Nash equilibria? Such a proof could potentially lead to a provably efficient algorithm for computing a sample correlated equilibrium of a polynomial game.

Finally, we note that while the adaptive discretization and moment relaxation algorithms converge in general and work well in practice, we do not know of any results regarding rate of convergence. If we regard the probability distributions produced by these algorithms at the k^{th} iteration as ϵ^k -correlated equilibria, how fast does ϵ^k converge to zero?

Acknowledgements

The authors would like to thank Professor Muhamet Yildiz for a productive discussion which led to an early formulation of the characterization theorems in Section 2 as well as the moment relaxation methods presented in Subsection 3.3. The first author would also like to thank Prof. Cesar E. Silva for many discussions about ergodic theory, and in particular for the statement and proof of Corollary 4.18 using Birkhoff's ergodic theorem. Figures were produced using the SeDuMi package for MATLAB [29].

A Semidefinite programming, sums of squares, and moments of measures

Definition A.1. A semidefinite program is an optimization problem of the form:

$$\begin{aligned} & \text{minimize} && L(S) \\ & \text{subject to} && T(S) = v \\ & && S \text{ is a symmetric matrix} \\ & && S \succeq 0 \text{ (positive semidefinite),} \end{aligned}$$

where L is a given linear functional, T is a given linear transformation, v is a given vector, and S is a square matrix of decision variables.

Semidefinite programs are convex and generalize linear programs (T and v can be designed to make S diagonal, in which case the condition $S \succeq 0$ is the same as the condition that $S \geq 0$ elementwise). Many problems can be expressed exactly or approximately as semidefinite programs, and this is important because semidefinite programs can be solved efficiently by interior point methods. For details and a variety of examples see [30] and [20].

The square of a real-valued function is nonnegative on its entire domain, as is 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 [23].

Proposition A.2. *A univariate polynomial p is nonnegative on \mathbb{R} if and only if it is a sum of squares.*

Proof. A simpler version of the proof of the following proposition. □

Proposition A.3 (Markov-Lukács [15]). *A univariate polynomial $p(x)$ is nonnegative on the interval $[-1, 1]$ if and only if $p(x) = s(x) + (1 - x^2)t(x)$ where s and t are both sums of squares of polynomials.*

Proof. Direct algebraic manipulations show that the set of polynomials of the form $s(x) + (1 - x^2)t(x)$ where s and t are sums of squares of polynomials in x is closed under multiplication and contains all polynomials of the following forms: a for $a \geq 0$, $(x - a)^2 + b^2$ for $a, b \in \mathbb{R}$, $x - a$ for $a \leq -1$, and $a - x$ for $a \geq 1$. By assumption $p(x)$ factors as a product of terms of these types, because any real root of p in the interval $(-1, 1)$ must have even multiplicity. □

These sum of squares conditions are easy to express using linear equations and semidefinite constraints.

Proposition A.4. *A univariate polynomial $p(x) = \sum_{k=0}^{2d} p_k x^k$ of degree at most $2d$ is a sum of squares of polynomials if and only if there exists a symmetric positive semidefinite matrix $Q \in \mathbb{R}^{(d+1) \times (d+1)}$ such that $p_k = \sum_{i+j=k} Q_{ij}$ (numbering the rows and columns of Q from 0 to d).*

Proof. Relating the coefficients of $p(x)$ to the entries of Q in this way is the same as writing $p(x) = \mathbf{x}^T Q \mathbf{x}$ where $\mathbf{x} = [1 \ x \ x^2 \ \cdots \ x^d]^T$. Thought of in this way, saying that $p(x)$ is a sum of squares is the same as saying that $Q = \sum_i q_i q_i^T$ for some column vectors q_i and in this case Q is clearly positive semidefinite. Conversely, if Q is positive semidefinite then there exists a matrix F such that $Q = F^T F$, so $p(x) = \mathbf{x}^T Q \mathbf{x} = \sum_i [F \mathbf{x}]_i^2$. \square

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 non-negative 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 [23]. Furthermore, for some special classes of multivariate polynomials, the sum of squares condition is exact.

Proposition A.5. *A matrix $M(t)$ whose entries are univariate polynomials in t is positive semidefinite on $[-1, 1]$ if and only if $x^T M(t) x = S(x, t) + (1 - t^2)T(x, t)$ where S and T are polynomials which are sums of squares.*

Now suppose we wish to answer the question of whether a finite sequence (μ^0, \dots, μ^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 $[-1, 1]$. This necessary condition for moments to correspond to a measure turns out to be sufficient [14] and can be written in terms of semidefinite constraints.

Proposition A.6. *The condition that a finite sequence of numbers (μ^0, \dots, μ^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 is obtained. 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] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA, 1997.

- [4] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.
- [5] M. Cripps. Extreme correlated and Nash equilibria in two-person games. <http://www.olin.wustl.edu/faculty/cripps/CES2.DVI>, November 1995.
- [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 (STOC)*, pages 71 – 78, New York, NY, 2006. ACM Press.
- [7] Ruchira Datta. Universality of Nash equilibrium. *Mathematics of Operations Research*, 28(3):424 – 432, August 2003.
- [8] M. Dresher 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. Dresher, 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] Fe. S. Evangelista and T. E. S. Raghavan. A note on correlated equilibrium. *International Journal of Game Theory*, 25(1):35 – 41, March 1996.
- [11] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1:80 – 93, 1989.
- [12] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1), February 1989.
- [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] M. G. Kreĭn and A. A. Nudel'man. *The Markov Moment Problem and Extremal Problems*, volume 50 of *Translations of Mathematical Monographs*. American Mathematical Society, 1977.
- [16] C. E. Lemke and J. T. Howson, Jr. Equilibrium points in bimatrix games. *SIAM Journal on Applied Math*, 12:413 – 423, 1964.

- [17] R. J. Lipton and E. Markakis. Nash equilibria via polynomial equations. In *Proceedings of LATIN*, 2004.
- [18] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286 – 295, September 1951.
- [19] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, New York, NY, 2005. ACM Press.
- [20] P. A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 2000.
- [21] P. A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control (CDC)*, 2006.
- [22] S. Prajna, A. Papachristodoulou, P. Seiler, and P. A. Parrilo. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*, 2004.
- [23] 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.
- [24] W. Rudin. *Real & Complex Analysis*. WCB / McGraw-Hill, New York, 1987.
- [25] W. Rudin. *Functional Analysis*. McGraw-Hill, New York, 1991.
- [26] C. E. Silva. *Invitation to Ergodic Theory*. American Mathematical Society, Providence, RI, 2007.
- [27] N. D. Stein, A. Ozdaglar, and P. A. Parrilo. Separable and low-rank continuous games. *International Journal of Game Theory*, To appear.
- [28] G. Stoltz and G. Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, to appear.
- [29] Jos F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.*, 11/12(1-4):625–653, 1999. Interior point methods.
- [30] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49 – 95, 1996.