

Polynomial games and sum of squares optimization

Pablo A. Parrilo

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology, Cambridge, MA 02139

Abstract—We study two-person zero-sum games, where the payoff function is a polynomial expression in the actions of the players. This class of games was introduced by Dresher, Karlin, and Shapley in 1950. We show that the value of the game, and the corresponding optimal strategies, can be obtained by solving a single semidefinite programming problem. In addition, we show how the results extend, with suitable modifications, to a general class of semialgebraic games.

I. INTRODUCTION

Game theory ([5], [12], [1]) is a well-established mathematical framework used to model and analyze the decision process of multiple decision makers having possibly conflicting objectives. For games in strategic form, several different notions of “solutions” or “equilibria” have been proposed over the years, with the most celebrated (for the class of games we will discuss) being the Nash equilibrium. This notion is characterized by a “no unilateral deviation” property. For two-player, zero-sum games, this concept coincides with the much better understood minimax value of a game.

Although much of the research efforts have been concentrated in developing computational techniques for equilibria of finite games (i.e., where each player has a finite number of pure strategies), there is much interest, particularly with a view towards applications, in so-called infinite games. In this important class, the players have access to an *infinite* number of non-equivalent pure strategies, and furthermore they are allowed to randomize over their choice.

This paper presents, for the first time, an appealing computational approach to the determination of optimal strategies for zero-sum infinite games, in the case where the payoffs are general polynomial functions, and the strategy sets are semi-algebraic. The solution strongly relies on techniques based on sum of squares (SOS) decompositions of polynomials, and the associated semidefinite programming characterization of the cone of SOS polynomials and its dual cone of moments.

Since its beginnings, game theory has been inextricably linked with optimization, and this relationship is nowhere stronger than in the case of zero-sum games, where convex duality-based ideas permeate both the theoretical and computational aspects of the subject. As we shall see, the solution presented here shares many of these features, and can be understood as a natural and complete generalization of the well-known linear programming (LP) characterization for finite games.

This research was funded in part by AFOSR MURI subaward 2003-07688-1.

An outline of the paper follows: in Section II we formally define the class of polynomial games considered. We discuss next several previously known characterizations of their solutions, as well as our main contributions. Section III contains several examples illustrating the application of the methods. In Section IV, we briefly describe several appealing extensions to general semialgebraic sets and the conic setting. Finally, we present some conclusions and future research directions.

II. POLYNOMIAL GAMES

Polynomial games were originally introduced and studied by Dresher, Karlin, and Shapley in 1950 [4]. In this paper, we concentrate on the computational aspects. In the basic setup of these games, there are two players (which we will denote as Player 1 and Player 2), which simultaneously and independently choose actions parametrized by real numbers x , y , respectively, in the interval $[-1, 1]$. The payoff associated with these choices is a polynomial function in x and y (a formal definition appears below in Section II-A).

The characterization and existence of optimal strategies for polynomial games has been worked out in [4]. However, no computational procedure was previously available to efficiently obtain the value of this class of games (existing approaches are reviewed in Section II-B). In fact, in the preface of [10], Kuhn and Tucker present a list of several important open questions regarding games with an infinite number of pure strategies, and in particular they write:

(7) To find a computational technique of general applicability for zero-sum two-person games with infinite sets of strategies. A constructive method for obtaining the optimal strategies for polynomial-like games or some large class of non-trivial continuous games would constitute a considerable contribution to this problem.

The main objective in this paper is precisely to illustrate how this goal can be achieved by using recent techniques from sum of squares optimization and semidefinite programming. In particular, we will show how to characterize and compute the optimal solution of this class of games by solving a single semidefinite programming problem.

A. Problem setup

For simplicity of presentation, we first consider games on the square $\Omega = [-1, 1] \times [-1, 1]$, as opposed to the more common $[0, 1] \times [0, 1]$. This is mainly for aesthetic and notational reasons, because the semidefinite programming

conditions take a particularly compact form in this case. It is straightforward to convert one type of game into the other by a linear transformation of the strategy space. Furthermore, the results can be transparently extended to finite unions of arbitrary intervals. Most of the results presented in this subsection are from [4].

In the class of games we consider the pure actions of each player are given by real numbers x and y , that belong to the closed interval $[-1, 1]$. The payoff is given by a polynomial expression

$$P(x, y) = \sum_{i=0}^n \sum_{j=0}^m p_{ij} x^i y^j, \quad (1)$$

that assigns payments from Player 2 to Player 1. Thus, Player 1 wants to choose his strategy x to maximize $P(x, y)$, while Player 2 tries to make this expression as small as possible. Although we do not expand on this here, many of the convenient properties of polynomial games can be extended to *separable* games, where the payoff is a finite sum of products of functions in the strategies of each player; see Section IV-A and [19]. When these assumptions are not satisfied, the resulting games can be quite complicated from a computational viewpoint. In particular, Glicksberg and Gross have given an example of a zero-sum game with rational payoffs, whose uniquely defined value is a transcendental number [7].

Since we are interested in zero-sum games, the appropriate equilibrium concept is that of a minimax (or Nash) equilibrium. For the game described by (1), we can consider as usual lower and upper bounds on the value of the game. These can be derived by computing

$$\max_x \min_y P(x, y) \quad \text{and} \quad \min_y \max_x P(x, y).$$

Since we have made no assumptions on the payoff $P(x, y)$, in general the maxmin will be different from the minmax (in other words, the game does not necessarily admit a solution in pure strategies). As is well-known from the finite case, an equality between these expressions can be recovered by allowing randomization over pure strategies (i.e., allowing mixed strategies).

The *mixed strategies* of each player correspond to probability measures ν, μ over the set of pure strategies, i.e., over the interval $[-1, 1]$. As we shall see, by enlarging the allowable set of strategies in this way, we will recover the notion of minimax equilibrium for these games. We illustrate the concept first with an example.

Example 2.1: Consider the game on $[-1, 1] \times [-1, 1]$, with payoff function given by $P(x, y) = (x - y)^2$. Since Player 2 wants to minimize her payoffs, she should try to “guess” the number chosen by Player 1. Conversely, the first player should try to make his number as difficult to guess as possible (in the sense defined by $P(x, y)$). It is easy to see in this case that the optimal strategy for Player 1 is to randomize between $x = -1$ or $x = 1$ with equal probability, and the optimal strategy of Player 2 is to always choose $y = 0$. No player has incentive to deviate from these strategies, and thus the value of the game is equal to 1.

When considering mixed strategies, and similar to the finite case, we need to consider the expressions

$$\max_{\nu} \min_{\mu} E_{\nu \times \mu}[P(x, y)] \quad \text{and} \quad \min_{\mu} \max_{\nu} E_{\nu \times \mu}[P(x, y)],$$

where $E_{\nu \times \mu}[\cdot]$ denotes the expectation under the product measure. We can rewrite these as bilinear expressions

$$\max_{\nu_i} \min_{\mu_j} \sum_{i=0}^n \sum_{j=0}^m p_{ij} \nu_i \mu_j \quad \min_{\mu_j} \max_{\nu_i} \sum_{i=0}^n \sum_{j=0}^m p_{ij} \nu_i \mu_j, \quad (2)$$

where ν_i, μ_j are the *moments* of the measures ν, μ , i.e.,

$$\nu_i := \int_{-1}^1 x^i d\nu, \quad \mu_j := \int_{-1}^1 y^j d\mu.$$

It is well-known that the moment spaces (i.e., the image of the probability measures under the moment map given above) are compact convex sets in \mathbb{R}^{n+1} and \mathbb{R}^{m+1} [9]. Since the objective function in the problems (2) is bilinear, and the feasible sets are convex and compact, a generalized version of the standard minimax theorem can be used to show that these two quantities are exactly equal [4]. Furthermore there exist measures ν^*, μ^* that satisfy the saddle-point condition:

$$\sum_{i=0}^n \sum_{j=0}^m p_{ij} \nu_i \mu_j^* \leq \sum_{i=0}^n \sum_{j=0}^m p_{ij} \nu_i^* \mu_j^* \leq \sum_{i=0}^n \sum_{j=0}^m p_{ij} \nu_i^* \mu_j. \quad (3)$$

The key fact here is that, due to the separable structure of the payoffs, the optimal strategies can be characterized only in terms of their first m (or n) moments. Higher moments are irrelevant, at least in terms of the payoffs of the players.

From the previous discussion, we have the following result, essentially contained in [4]:

Theorem 2.2: Consider the two-player zero-sum game on $[-1, 1] \times [-1, 1]$, with payoff given by (1). Then, the value of the game is well-defined, and there exist optimal mixed strategies ν^*, μ^* satisfying a saddle-point condition. Furthermore, without loss of generality, the optimal measures can be taken to be discrete, with at most $\min(n, m) + 1$ atoms.

This result has been extended in several directions, for instance general two-player zero-sum continuous-kernel games (e.g. [1, Corollary 4.3]), Nash equilibria of separable games [19], etc.

B. Earlier results and approaches

Polynomial games generated considerable excitement shortly after they were introduced, as a possible bridge between finite and infinite games. However, the lack of efficient computational methods soon caused somewhat of a disappointment. In [8, p. 78], we find:

”It was soon realized that polynomial games of large degree possess solutions of great complexity which are impossible to so much as describe in qualitative terms, let alone calculate.”

As we will see, the notable advances in optimization theory in the last decade can be put to good use to update in a significant manner this assessment. Indeed, there were good

reasons for the lack of satisfactory solutions to this problem. Until the development of semidefinite programming, and the connections with sums of squares techniques, even the simple minimization of a univariate polynomial was not amenable to convex methods. Since solving a game cannot be easier than optimizing (consider, e.g., a game that does not depend on the actions of one of the players), this fact was a significant roadblock towards an efficient solution method.

Among several other interesting results for polynomial games worth mentioning, Gale and Gross [6] have shown how to explicitly produce a polynomial game that has two *a priori* specified measures as the unique minimax equilibria.

In terms of computation, besides the relatively direct schemes that approximate the infinite game by a (large) matrix game obtained by discretizing the strategy space, we have found in the related literature only two earlier references on effective computational methods for this kind of games. The first one is an approach presented by Karlin in [8], and is a technique based on the explicit mapping of convex sets. While relatively easy to implement for low dimension, its complexity quickly becomes unmanageable for larger problems. The other approach, by Tiskin, is based on a generic cutting plane method, and appears in [20].

C. The value of the game

The derivation of the value of the game can be done in a very similar way as in the finite actions case. We characterize first “security strategies” that provide a minimum guaranteed payoff. We can then invoke convex duality to prove that this actually yields the unique value of the game.

Proceeding along these lines, by analogy to the finite case, a security strategy of Player 2 can be computed by solving

$$\min_{\gamma, \mu} \gamma \quad \text{s.t.} \quad \begin{cases} E_{\mu}[P(x, y)] \leq \gamma \quad \forall x \in [-1, 1] \\ \int_{-1}^1 d\mu(y) = 1. \end{cases} \quad (4)$$

Indeed, if Player 2 plays the mixed strategy μ obtained from the solution of this problem, the best that Player 1 can do is to choose a value of x that maximizes $E_{\mu}[P(x, y)]$, thus limiting his gain (and Player 2’s loss) to γ .

Since $P(x, y)$ is polynomial, this expectation can be equivalently written in terms of the first n moments of the measure μ , i.e.,

$$E_{\mu}[P(x, y)] = \int_{-1}^1 P(x, y) d\mu(y) = \sum_{i=0}^n \sum_{j=0}^m p_{ij} \mu_j x^i.$$

Notice that this is a univariate polynomial in the action x of Player 1, with coefficients that depend *affinely* on the moments μ_j of the mixed strategy of Player 2. This property will be crucial in our developments.

Consider now the problem (4), but instead of writing it in terms of the decision variable μ (which is a probability measure), let us use instead the moments $\{\mu_j\}_{j=0}^m$. The problem is then reduced to the minimization of the safety level γ , subject to the conditions:

- The univariate polynomial $\gamma - \sum_{i=0}^n \sum_{j=0}^m p_{ij} x^i \mu_j$ is nonnegative on $[-1, 1]$.

- The sequence $\{\mu_j\}_{j=0}^m$ is a valid moment sequence for a probability measure supported in $[-1, 1]$.

We can rewrite this in a more compact form, as the optimization problem

$$\min \gamma \quad \text{s.t.} \quad \begin{cases} \gamma - \sum_{i=0}^n \sum_{j=0}^m p_{ij} x^i \mu_j \in \mathcal{P}_n \\ \mu \in \mathcal{M}_m \\ \mu_0 = 1 \end{cases} \quad (5)$$

where \mathcal{P}_n is the set of univariate polynomials of degree n nonnegative in $[-1, 1]$, and \mathcal{M}_m is the set of $m + 1$ first moments of a nonnegative measure with support on the same interval.

In order to convert the abstract formulation (5) into a concrete optimization problem that we can solve, we need a computationally convenient representation of these sets. This is provided in the next section.

D. SDP characterization of nonnegativity and moments

We show in this section that the sets $\mathcal{P}_n, \mathcal{M}_m$ introduced above are exactly representable in terms of semidefinite programming conditions. These results are relatively well-known by now, and follow directly from the fact that nonnegative univariate polynomials are sums of squares, and classical results on the moment problem.

We characterize below a few basic properties of the sets of nonnegative polynomials and measures.

Lemma 2.3: The sets $\mathcal{P}_n, \mathcal{M}_n$ are $(n + 1)$ -dimensional proper cones (i.e., closed, convex, pointed and solid). Furthermore, they are convex duals, i.e., $\mathcal{P}_n = \mathcal{M}_n^*$. Remarkably, these cones admit nice representations in terms of semidefinite programming conditions. We introduce below some simple notation that will allow for a concise and explicit presentation of the results.

Let \mathcal{S}^d denote the set of $d \times d$ real symmetric matrices, and define the linear operator $\mathcal{H} : \mathbb{R}^{2d-1} \rightarrow \mathcal{S}^d$ as:

$$\mathcal{H} : \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_{2d-1} \end{bmatrix} \mapsto \begin{bmatrix} a_1 & a_2 & \dots & a_d \\ a_2 & a_3 & \dots & a_{d+1} \\ \vdots & \vdots & \ddots & \vdots \\ a_d & a_{d+1} & \dots & a_{2d-1} \end{bmatrix}.$$

By its definition, \mathcal{H} simply takes a vector, and constructs the associated Hankel matrix. Its corresponding adjoint is given by a linear map $\mathcal{H}^* : \mathcal{S}^d \rightarrow \mathbb{R}^{2d-1}$,

$$\mathcal{H}^* : \begin{bmatrix} m_{11} & m_{12} & \dots & m_{1d} \\ m_{12} & m_{22} & \dots & m_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ m_{1d} & m_{2d} & \dots & m_{dd} \end{bmatrix} \mapsto \begin{bmatrix} m_{11} \\ 2m_{12} \\ m_{22} + 2m_{13} \\ \vdots \\ m_{d-1,d-1} + 2m_{d-2,d} \\ m_{dd} \end{bmatrix},$$

that “flattens” a matrix into a vector by adding all the entries along antidiagonals.

A polynomial $f(x)$ is a sum of squares (SOS) if it can be written as $f(x) = \sum_i g_i(x)^2$ for some polynomials g_i . Using the notation above, we can write the following simple SDP condition for polynomial nonnegativity:

Lemma 2.4: The polynomial $p(x) = \sum_{k=0}^{2d} p_k x^k$ is non-negative (or SOS) if and only if there exists $S \in \mathcal{S}^{d+1}$, $S \succeq 0$ such that

$$\begin{bmatrix} p_0 \\ p_1 \\ \vdots \\ p_{2d} \end{bmatrix} = \mathcal{H}^*(S).$$

Proof: For univariate polynomials, nonnegativity is equivalent to SOS (see, e.g., [16], [14]). Furthermore, letting $[x]_d := [1, x, \dots, x^d]^T$, for every $S \in \mathcal{S}^{d+1}$ we have

$$p(x) = \langle \mathcal{H}^*(S), [x]_{2d} \rangle = [x]_d^T S [x]_d,$$

and factorizing $S \succeq 0$, we obtain a sum of squares decomposition of $p(x)$. ■

We can give a similar characterization of polynomials non-negative in an interval. For this, define

$$L_1 = \begin{bmatrix} I_{d \times d} \\ 0_{1 \times d} \end{bmatrix}, \quad L_2 = \begin{bmatrix} 0_{1 \times d} \\ I_{d \times d} \end{bmatrix}.$$

Lemma 2.5: The polynomial $p(x) = \sum_{k=0}^{2d} p_k x^k$ is non-negative in $[-1, 1]$ if and only if there exist $Z \in \mathcal{S}^{d+1}$, $W \in \mathcal{S}^d$, $Z \succeq 0$, $W \succeq 0$ such that

$$\begin{bmatrix} p_0 \\ p_1 \\ \vdots \\ p_n \end{bmatrix} = \mathcal{H}^*(Z + L_1 W L_1^T - L_2 W L_2^T).$$

Proof: Follows directly from the characterization of univariate polynomials nonnegative in an interval, and Lemma 2.4. Indeed, it is well known that

$$p(x) \geq 0 \quad \forall x \in [-1, 1] \quad \Leftrightarrow \quad p(x) = z(x) + w(x)(1-x^2),$$

where $z(x), w(x) \in \mathbb{R}[x]$ are sums of squares. ■

Dualizing the previous conditions, or invoking classical results on moment spaces (e.g. [9], [17]), we have a similar characterization for \mathcal{M}_n .

Lemma 2.6: The vector $\mu = [\mu_0, \mu_1, \dots, \mu_n]^T$ is a valid set of moments for a probability measure in $[-1, 1]$ if and only if

$$\begin{aligned} \mu_0 &= 1 \\ \mathcal{H}(\mu) &\succeq 0 \\ L_1^T \mathcal{H}(\mu) L_1 - L_2^T \mathcal{H}(\mu) L_2 &\succeq 0. \end{aligned} \quad (6)$$

Proof: This result follows from the previous lemma, by the duality between nonnegative polynomials and moment spaces, and the fact that $[-1, 1]$ is compact. Alternatively, a direct proof is presented, e.g., in [9]. ■

We remark that these SDP characterizations of univariate nonnegative polynomials and moments have been used, in different contexts, in earlier work of Shor [18], Nesterov [13], and Parrilo [14]. On the dual side, the moment approach has been described by Bertsimas and Popescu [2] and Lasserre [11].

Remark 2.7: It is possible to give slightly improved formulations of Lemmas 2.5 and 2.6, depending on whether the degree of the polynomial (or the number of moments) is even or odd. For simplicity, we do not discuss these here, but notice that the formulations presented will always give correct results.

E. SDP reformulation

The results in the previous section directly provide semidefinite representations of the sets \mathcal{P}_n and \mathcal{M}_m appearing in (5). For concreteness, we present below the explicit semidefinite program that results from combining these expressions:

$$\begin{aligned} & \min_{\gamma, \mu, Z, W} \gamma \\ \text{s.t.} \quad & \begin{cases} \mathcal{H}(\mu) \succeq 0 \\ L_1^T \mathcal{H}(\mu) L_1 - L_2^T \mathcal{H}(\mu) L_2 \succeq 0 \\ \mathcal{H}^*(Z + L_1 W L_1^T - L_2 W L_2^T) = \gamma e_1 - P \mu \\ e_1^T \mu = 1 \\ Z, W \succeq 0, \end{cases} \end{aligned} \quad (7)$$

where $\mu \in \mathbb{R}^{m+1}$, $P \in \mathbb{R}^{(n+1) \times (m+1)}$ contains the coefficients p_{ij} of the polynomial $P(x, y)$, and $e_1 \in \mathbb{R}^{m+1}$ has all entries equal to zero, except for a one in the first row.

As it will become clear from its dual, the solution of the SDP (7) exactly corresponds to the value of the game, and the moments of the optimal strategies.

F. Duality

It is well known that in zero-sum games there is a natural relationship between the game-theoretic role of the two players, and the convex duality properties of the corresponding optimization problem. Informally, going to the dual optimization problem is equivalent to switching the role of the players.

Allowing for a sign change of the optimal value, the dual SDP problem can be written as:

$$\begin{aligned} & \min_{\alpha, \nu, A, B} \alpha \\ \text{s.t.} \quad & \begin{cases} \mathcal{H}(\nu) \succeq 0 \\ L_1^T \mathcal{H}(\nu) L_1 - L_2^T \mathcal{H}(\nu) L_2 \succeq 0 \\ \mathcal{H}^*(A + L_1 B L_1^T - L_2 B L_2^T) = \alpha e_1 + P^T \nu \\ e_1^T \nu = 1 \\ A, B \succeq 0. \end{cases} \end{aligned} \quad (8)$$

This expression agrees exactly with the primal SDP (7), except for the change in the payoff matrix from P to $-P^T$, which naturally corresponds to the same game, but from the point of view of the other player. We have therefore a perfect correspondence between the primal and dual sides of the game, given by the mapping $(P, Z, W, \mu, \gamma) \leftrightarrow (-P^T, A, B, \nu, \alpha)$. In view of the original polynomial payoff defined by (1), the mapping that takes the matrix P into $-P^T$ corresponds to $P(x, y) \leftrightarrow -P(y, x)$.

We can then formally state our main result as follows.

Theorem 2.8: Consider a two-player zero-sum polynomial game in $[-1, 1] \times [-1, 1]$, with payoff as described in (1). The value of the game, and optimal mixed strategies, can be obtained by solving the primal-dual SDP pair given by (7)-(8).

This result provides a quite natural and complete generalization of the classical linear programming solution of zero-sum finite games. Furthermore, it shares the same desirable properties of self-duality.

G. Recovering the optimal strategies

The decision variables of the SOS/SDP problems presented earlier were the moments of the mixed strategies. The corresponding measures can be easily recovered from the optimal primal-dual solutions of the SDP, in particular the matrices $\mathcal{H}(\mu)$ and $\mathcal{H}(\nu)$. By complementary slackness (or a simple game theoretic argument), it is easy to see that the support of the atomic measures will be given by the zeros of the polynomials with coefficients given by $\gamma e_1 - P\mu$ and $\alpha e_1 + P^T\nu$. The corresponding weights can be obtained by solving a linear system.

The procedure to recover a univariate atomic measure from its moments is classical, and can be found, e.g., in Shohat and Tamarkin [17], Karlin and Shapley [9], Devroye [3], among others. We omit the details from this version of the paper due to space reasons.

H. Pure strategy solutions

In many applications (e.g., networking and economics) it is of interest of study conditions under which a game is guaranteed to possess optimal solutions that are pure, i.e., are not probabilistic mixtures of strategies. A well-known result about continuous games (e.g., [1, Theorem 4.5]) is the following:

Theorem 2.9: Consider the polynomial game in $[-1, 1] \times [-1, 1]$ described by $P(x, y)$. If $P(x, y)$ is strictly concave in x for each $y \in [-1, 1]$ and strictly convex in y for each $x \in [-1, 1]$, then both players have optimal strategies that are pure.

For games of this kind, then the primal/dual SDP pair (7)-(8) will admit rank one solutions (which will be unique, if the game has a unique equilibria). It should be noticed, however, that if it is known that the game has this convex-concavity property, there is no computational advantage in working on the moment spaces rather than in the natural strategy spaces of the game.

III. EXAMPLES

We present next two simple examples. The first game has as solutions only pure strategies, while the other one (extracted from the literature) require mixed strategies in order to achieve the optimal value of the game.

Example 3.1: Consider the payoff function on $[-1, 1] \times [-1, 1]$ given by:

$$P(x, y) = 2xy^2 - x^2 - y$$

What are the optimal strategies of the two players, and what is the value of the game?

The SDP (7) can be formulated as:

$$\begin{aligned} & \text{maximize } \gamma \\ & \text{s.t. } \begin{cases} \begin{bmatrix} 1 & \mu_1 \\ \mu_1 & \mu_2 \end{bmatrix} \succeq 0, & 1 - \mu_2 \succeq 0, \\ \gamma + x^2 - 2x\mu_2 + \mu_1 = \begin{bmatrix} 1 \\ x \end{bmatrix}^T Z \begin{bmatrix} 1 \\ x \end{bmatrix} + (1 - x^2)W, \\ Z \succeq 0, W \succeq 0. \end{cases} \end{aligned}$$

Solving the SDPs as described above, we directly obtain:

$$\begin{aligned} \gamma &= \alpha^4 - \alpha & \mu_1 &= \alpha & \mu_2 &= \alpha^2 \\ Z &= \begin{bmatrix} \alpha^4 & -\alpha^2 \\ -\alpha^2 & 1 \end{bmatrix} & W &= 0, \end{aligned}$$

where $\alpha = 4^{-\frac{1}{3}} \approx 0.62996$. The value of the game is therefore $\gamma \approx -0.47247$.

For this simple game, we can easily extract the optimal strategies, by decomposing the matrices Z, W . In this case, the solution is given by pure strategies, where Player 1 always chooses $x = \alpha^2$, and Player 2 always chooses $y = \alpha$.

The fact that the optimal strategies are pure is to be expected, in view of Theorem 2.9 and the fact that the payoff function is concave in x . This equilibrium corresponds to the unique saddle-point of the function $P(x, y)$ on the domain of the function. We notice that since in this case there is a unique solution in pure strategies, the saddle point could have also been computed in a more direct way by using the ‘‘reaction curve’’ method; see e.g. [12], [1].

Example 3.2: The payoff function is given by:

$$P(x, y) = 5xy - 2x^2 - 2xy^2 - y,$$

which is neither convex nor concave. After solving the problem using the SOS/SDP optimization approach described, we obtain:

$$\begin{aligned} \gamma &= -0.48 & \mu_1 &= 0.56 & \mu_2 &= 1 \\ Z &= \begin{bmatrix} 0.08 & -0.4 \\ -0.4 & 2 \end{bmatrix} & W &= 0. \end{aligned}$$

This corresponds to a mixed strategy, where Player 1 always picks $x = 0.2$, and Player 2 chooses $y = 1$ with probability 0.78, and $y = -1$ with probability 0.22. The corresponding conditional payoffs are shown in Figure 1, where we can graphically verify that these measures indeed corresponds to the equilibrium of the game.

IV. EXTENSIONS AND GENERALIZATIONS

The results presented here are a natural starting point for a wide variety of extensions and generalizations. We outline some of them below, in order of generality.

A. Extensions

There are several possible variations that are a quite straightforward generalization of these results. Among them, we mention:

- Strategy sets that are finite unions of intervals. As mentioned earlier, the results directly extend to this setting.

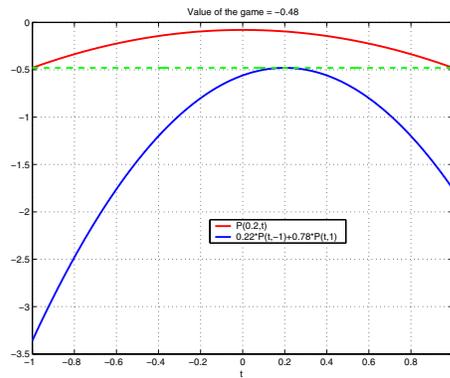


Fig. 1. Solution for Example 3.2. The upper curve represents the expected payoff of Player 2, if Player 1 plays the optimal strategy. Conversely, the lower curve is the expected payoff of Player 1, when Player 2 plays optimally. We can see that the minimum of the first curve is equal to the maximum of the second one, giving the value of the game $\gamma = -0.48$. The points at which these extrema are achieved give the support of the respective optimal measures.

- Discrete-continuous games, where one of the players has a finite set of pure strategies, while the other player has a continuum (e.g., an interval).
- Trigonometric games. Here the payoff is a trigonometric polynomial of the decision variables given by angles θ_1, θ_2 , corresponding to the two players.

In all these cases, the solution proceeds along the exact same lines, since we can use SOS methods to provide SDP characterization of nonnegative trigonometric polynomials or univariate nonnegative rational functions (and their dual moment cones). Details, which are more or less standard, are omitted from this version of the paper for space reasons.

B. Semialgebraic games

For simplicity of exposition, in this paper we have described in detail only the univariate case. Our solution can be easily extended to a much more general setup, with general strategy sets $\mathcal{X} \subset \mathbb{R}^n, \mathcal{Y} \subset \mathbb{R}^m$ that are basic semialgebraic, i.e., are given by a finite number of polynomial equations and inequalities, and a polynomial payoff. In this case, however, as opposed to the setting discussed earlier, the problem of determining the value of the game immediately becomes NP-hard, since both polynomial nonnegativity and the recognition of valid moment sequences are hard problems. However, as in [14], [15], [11] we can approximate arbitrarily tightly these conditions using results based on the Positivstellensatz or Schmüdgen/Putinar representations. In some cases (“Hilbert” games), we can obtain guaranteed solutions, or hard bounds on the value of the game. The details are omitted from this version for space reasons.

V. CONCLUSIONS AND FUTURE WORK

In this paper we presented an exact semidefinite programming solution to the computation of the optimal value and optimal strategies of zero-sum polynomial games. The approach relies on the connections between nonnegative

polynomials, sums of squares, and semidefinite programming. The developments closely parallel the classical linear programming-based solution of finite games, and shares many of the same appealing convexity and self-duality properties. We have also described how the techniques extend to the much more general semialgebraic case. Some further extensions of these techniques to the nonzero sum case have been developed in [19].

Acknowledgments: I would like to thank Prof. Alexandre Megretski, for his help with the translation of [20], and to Noah Stein, for bringing Karlin’s quote in Section II-B to my attention.

REFERENCES

- [1] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Number 23 in Classics in Applied Mathematics. SIAM, Philadelphia, PA, 1999.
- [2] D. Bertsimas and I. Popescu. Optimal inequalities in probability theory: a convex optimization approach. *SIAM J. Optim.*, 15(3):780–804 (electronic), 2005.
- [3] L. Devroye. *Nonuniform random variate generation*. Springer-Verlag, New York, 1986.
- [4] M. Drescher, S. Karlin, and L. S. Shapley. Polynomial games. In *Contributions to the Theory of Games*, Annals of Mathematics Studies, no. 24, pages 161–180. Princeton University Press, Princeton, N. J., 1950.
- [5] D. Fudenberg and J. Tirole. *Game theory*. MIT Press, Cambridge, MA, 1991.
- [6] D. Gale and O. Gross. A note on polynomial and separable games. *Pacific J. Math.*, 8:735–741, 1958.
- [7] I. Glicksberg and O. Gross. Notes on games over the square. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 173–182. Princeton University Press, Princeton, N. J., 1953.
- [8] S. Karlin. *Mathematical methods and theory in games, programming and economics. Vol. I: Matrix games, programming, and mathematical economics. Vol. II: The theory of infinite games*. Addison-Wesley Publishing Co., Inc., Reading, Mass.-London, 1959.
- [9] S. Karlin and L. Shapley. *Geometry of moment spaces*, volume 12 of *Memoirs of the American Mathematical Society*. AMS, 1953.
- [10] H. W. Kuhn and A. W. Tucker, editors. *Contributions to the Theory of Games*. Annals of Mathematics Studies, no. 24. Princeton University Press, Princeton, N. J., 1950.
- [11] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.
- [12] R. B. Myerson. *Game theory: analysis of conflict*. Harvard University Press, Cambridge, MA, 1991.
- [13] Y. Nesterov. Squared functional systems and optimization problems. In J.B.G. Frenk, C. Roos, T. Terlaky, and S. Zhang, editors, *High Performance Optimization*, pages 405–440. Kluwer Academic Publishers, 2000.
- [14] P. A. Parrilo. *Structured semidefinite programs and semi-algebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, May 2000. Available at <http://resolver.caltech.edu/CaltechETD:etd-05062004-055516>.
- [15] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Math. Prog.*, 96(2, Ser. B):293–320, 2003.
- [16] B. Reznick. Some concrete aspects of Hilbert’s 17th problem. In *Contemporary Mathematics*, volume 253, pages 251–272. American Mathematical Society, 2000.
- [17] J.A. Shohat and J.D. Tamarkin. *The Problem of Moments*. American Mathematical Society Mathematical surveys, vol. II. American Mathematical Society, New York, 1943.
- [18] N. Z. Shor. Class of global minimum bounds of polynomial functions. *Cybernetics*, 23(6):731–734, 1987. (Russian orig.: Kibernetika, No. 6, (1987), 9–11).
- [19] N. D. Stein, A. E. Ozdaglar, and P. A. Parrilo. Separable and low-rank continuous games. MIT LIDS technical report, 2006.
- [20] V. L. Tiskin. Solution of polynomial games by the method of cut-offs. In *Advances in game theory (Proc. Second All-Union Conf., Vilnius, 1971) (Russian)*, pages 69–72. Izdat. “Mintis”, Vilnius, 1973.