SOS/SDP methods: from optimization to games

Pablo A. Parrilo
Electrical Engineering
and Computer Science
MIT

www.mit.edu/~parrilo

Massachusetts Institute of Technology

E E
C S

Outline

- Optimization vs. games: alternating quantifiers
- Motivation, examples
- Polynomial games, and geometry of solutions
- Nonnegativity, sum of squares, and SDP
- Characterization of mixed strategies
- Putting it all together
- Semialgebraic games
- Future research and conclusions.

Optimization and games

Setting \ # DM	One	Many
Static	Optimization	Game theory
Dynamic	Optimal control	Dynamic games

Abstractly, what matters is the structure of the logical formula. Both existential and universal quantifiers?

$$\begin{cases} \exists x \, P(x) \\ \forall x \, P(x) \end{cases} \qquad \text{vs.} \qquad \begin{cases} \exists x \, \forall y \; P(x,y) \\ \forall x \, \exists y \; P(x,y) \end{cases}$$

- Many versions (minimax, games, robust control, etc.)
- Different levels of the polynomial hierarchy NP/co-NP $= \Sigma_1/\Pi_1$, Σ_2/Π_2, $\ldots$

From optimization to games

- Problems with *two* alternating quantifiers.

$$\exists x \forall y P(x,y) \qquad\qquad \forall x \exists y P(x,y)$$

- We are interested in predicates that are semialgebraic.
- In general, harder than NP-complete, in the second level of the polynomial hierarchy.
  - Game theory
  - Minimax optimization
  - Robust control, robust optimization

  An important feature: we consider *mixed* strategies.
- Sometimes, however, this is irrelevant: optimal strategies are pure.

Game theory 101: finite games

Two-player, zero-sum, finite games, strategic form.

Example: Rock, paper, scissors

P1 \ P2			
	0	-1	1
	1	0	-1
	-1	1	0

P1 chooses rows, P2 chooses columns. *Finite* number of pure strategies.

Notions of equilibria? How to compute "optimal" strategies?

Review: finite games, LP solution

Well-defined equilibrium concept: minimax value of the game.

Optimal strategies are *randomized* (mixed).

For mixed strategies, we have von Neumann's minimax theorem:

$$\max_{x\in\Delta_n} \min_{y\in\Delta_m} x^T P y \quad = \quad \min_{y\in\Delta_m} \max_{x\in\Delta_n} x^T P y.$$

Can compute this by solving the primal-dual pair of LPs

$$\begin{cases} \max & \alpha \\ \text{s.t.} & P^T x \geq \alpha \mathbf{1} \\ & x \in \Delta_n \end{cases} \qquad\qquad \begin{cases} \min & \beta \\ \text{s.t.} & P y \leq \beta \mathbf{1} \\ & y \in \Delta_m \end{cases}$$

Gives a *saddle point* $(x_*, y_*) \in \Delta_n \times \Delta_m$:

$$\max_{x\in\Delta_n} x^T P y_* \quad = \quad x_*^T P y_* \quad = \quad \min_{y\in\Delta_m} x_*^T P y.$$

Infinite strategies

We are interested in games with an *infinite* number of pure strategies.

In particular, the strategy sets will be semialgebraic, defined by polynomial equations and inequalities.

How to characterize and compute optimal strategies?

It is known that under mild conditions, we can always discretize and approximate with a finite game.

However, that can be irrelevant for computation. For instance, the same thing is true for optimization problems...

Example: A guessing game

$$\begin{aligned} \mathcal{X} &= \{x \,|\, -1 \leq x \leq 1\} \\ \mathcal{Y} &= \{y \,|\, -1 \leq y \leq 1\} \\ P(x,y) &= (x-y)^2 \end{aligned}$$

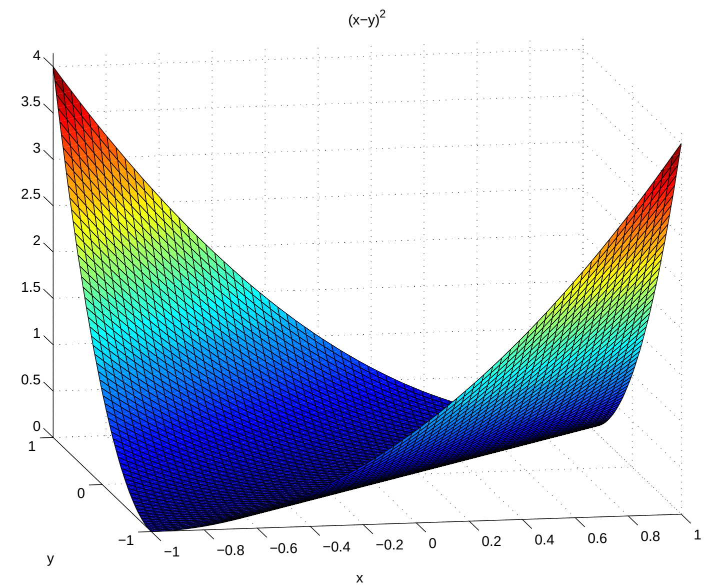

- Zero-sum, infinite number of pure strategies.
- The simplest case, the strategy space is $[-1,1] \times [-1,1]$, and the payoff (to X) is a polynomial function of x and y.

Does the game have a value?
What are the optimal strategies, and how to compute them?

Polynomial games

Introduced by Dresher, Karlin and Shapley (1950). No computational methods available (other than discretization).

- Here, we'll concentrate on two-player, zero-sum games
- Simplest case, strategy space is $[-1, 1] \times [-1, 1]$, and the payoff is a polynomial function of x and y
- The value of the game is well-defined
- Optimal strategies exist. WLOG, finite support.
- Includes polynomial optimization as special case.

A simple game: solution

$$\begin{aligned} \mathcal{X} &= \{x \mid -1 \leq x \leq 1\} \\ \mathcal{Y} &= \{y \mid -1 \leq y \leq 1\} \\ P(x,y) &= (x-y)^2 \end{aligned}$$

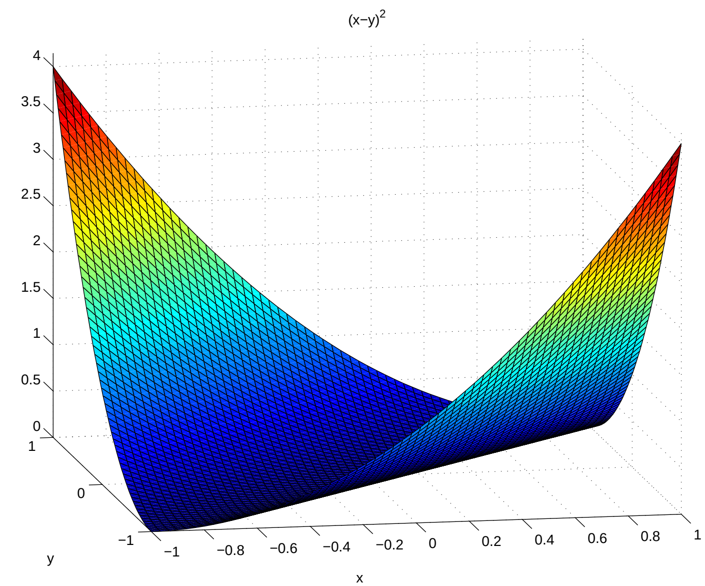

- For *pure* strategies, maxmin $\neq$ minmax. Need to randomize.
- Player X plays $x = -1$ or $x = 1$ with probability $\frac{1}{2}$.
- Player Y always plays $y = 0$.
- The value of the game is 1, no player has incentive to deviate.

Too easy. What about general polynomial payoffs?

Computing the solution

Zero-sum polynomial game in $[-1,1] \times [-1,1]$, payoff function given by

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

The mixed strategies for each player correspond to measures ν, μ.

Similar to the finite case, want to solve

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

where ν, μ are probability measures with support in $[-1,1]$.

Moments and minimax

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

Equivalently,

$$\max_{\nu_i} \min_{\mu_j} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij}\nu_i\mu_j \quad = \quad \min_{\mu_j} \max_{\nu_i} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij}\nu_i\mu_j.$$

where ν_j, μ_j are the *moments*

$$\nu_j := \int_{-1}^{1} x^j d\nu, \qquad \mu_j := \int_{-1}^{1} y^j d\mu.$$

What does the set of moments look like?

Geometry of Moments

The set of moments of measures supported on $[-1, 1]$ is convex.

Convex hull of the curve $(1, t, \ldots, t^d)$.

Not polyhedral (a "spectrahedron").

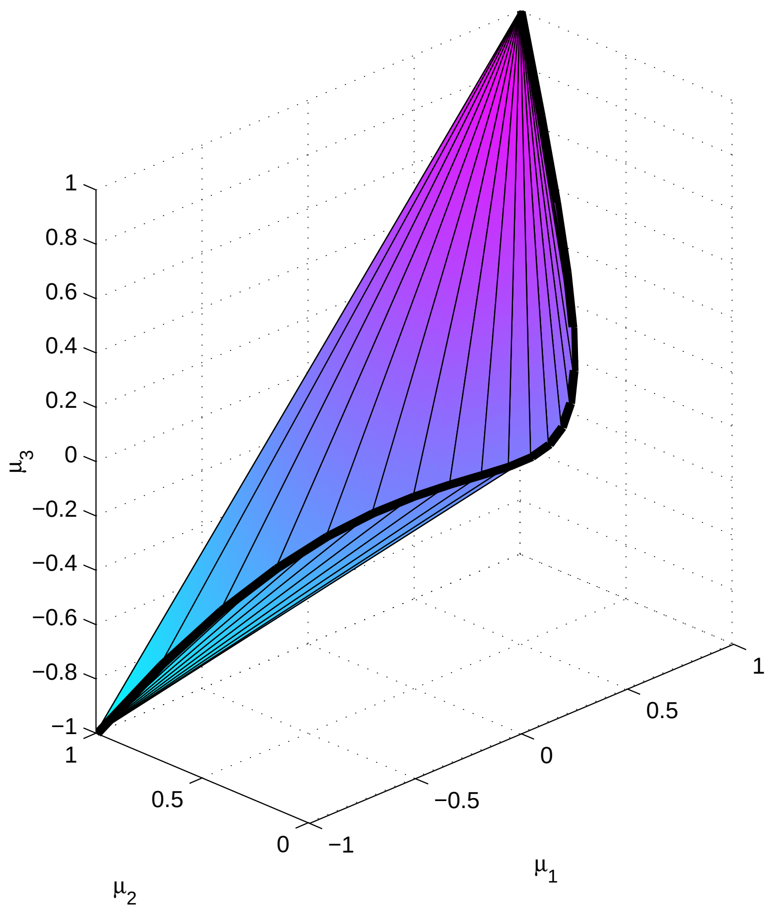

Well-understood geometry.

"Simplicial": every supporting hyperplane yields a simplex.

Related to cyclic polytopes.

Moments and minimax

The game is equivalent to:

$$\max_{\nu_i} \min_{\mu_j} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij} \nu_i \mu_j \quad = \quad \min_{\mu_j} \max_{\nu_i} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij} \nu_i \mu_j .$$

A bilinear function, and two finite dimensional compact convex sets.

Thus, the minimax theorem still applies.

Key fact: The optimal strategies are characterized only in terms of its first m (or n) moments. Higher moments are irrelevant.

How to compute solutions? What's the equivalent of the game LPs?

Computation

Similar derivation as in the finite case: if Player 1 knows the strategy μ of Player 2, he can ensure a payoff of $\max_{x \in [-1,1]} E_\mu[P(x,y)]$.

Thus, Player 2 should be optimizing

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

Since $P(x,y)$ is polynomial, the constraints can be equivalently written in terms of the first m moments of the measure μ:

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

This is a *univariate polynomial* in x.

From minimax to optimization

Thm: (P.) Consider a zero-sum polynomial game in $[-1,1] \times [-1,1]$, with payoff function given by

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

The value of the game, and optimal mixed strategies, can be computed via

$$\min \ \beta \qquad \text{s.t.} \qquad \begin{cases} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij} \mu_j x^i \leq \beta \quad \forall x \in [-1,1] \\ \mu_j \text{ are moments of a } [-1,1] \text{ measure} \end{cases}$$

This is a *convex optimization* problem.

Furthermore, is exactly equivalent to a finite-dimensional SDP problem.

Two obstacles

To solve this, we need to understand how to computationally characterize:

- Polynomials that are *nonnegative* on a given set.
- The *valid moments* of a measure.

Interestingly, the answers turn out to be dual of each other!

We'll explain how. But first, a mini-review about SDP.

Semidefinite programming - background

- A semidefinite program:

$$M(z) := M_0 + \sum_{i=1}^{m} z_i M_i \succeq 0,$$

  where $z \in \mathbb{R}^m$ are decision variables and $M_i \in \mathbb{R}^{n \times n}$ are given symmetric matrices.

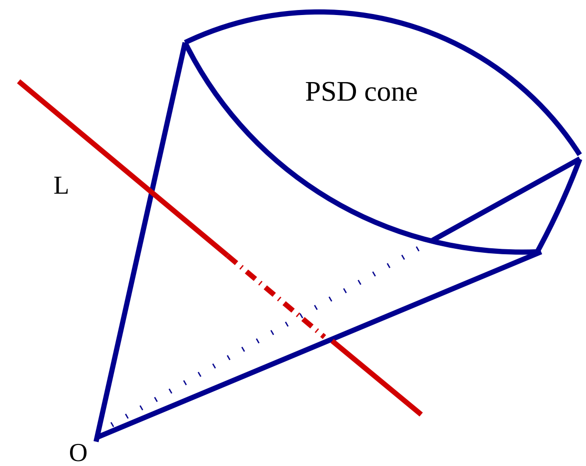

- The intersection of an affine subspace and the self-dual convex cone of positive semidefinite matrices.
- Convex finite dimensional optimization problem.
- A broad generalization of linear programming. Nice duality theory.
- Essentially, solvable in *polynomial time* (interior point, etc.).
- *Many* applications.

Nonnegativity of polynomials

How to check if a given $F(x_1, \ldots, x_n)$ is globally nonnegative?

$$F(x_1, x_2, \ldots, x_n) \geq 0, \quad \forall x \in \mathbb{R}^n$$

In general, a hard question. *NP-hard* if $\deg(F) \geq 4$.
A "simple" sufficient condition: a sum of squares (SOS) decomposition:

$$F(x) = \sum_i f_i^2(x) \qquad \Rightarrow \qquad F(x) \geq 0, \ \forall x \in \mathbb{R}^n.$$

Important properties:

- In some cases (e.g. univariate), it is exact.
- Efficiently computable, using semidefinite programming.

Some notation

For later reference, define a linear operator $\mathcal{H} : \mathbb{R}^{2d-1} \to \mathcal{S}^d$:

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

Its adjoint $\mathcal{H}^* : \mathcal{S}^d \to \mathbb{R}^{2d-1}$ is given by:

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

that "flattens" a matrix by adding along antidiagonals. Furthermore, let

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

SOS and SDP

Lemma 1. *The polynomial* $p(x) = \sum_{k=0}^{2d} p_k x^k$ *is nonnegative (or SOS) if and only if there exist* $S \in \mathcal{S}^{d+1}, S \succeq 0$ *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. Furthermore, letting $m_d := [1, x, \dots, x^d]^T$, for every $S \in \mathcal{S}^{d+1}$ we have

$$p(x) = \langle \mathcal{H}^*(S), m_{2d} \rangle = m_d^T S m_d.$$

Example:

$$x^4 - 2x^2 + 2 = 1 + (1 - x^2)^2 = \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}^T \begin{bmatrix} 2 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ x \\ x^2 \end{bmatrix}$$

Polynomials nonnegative on an interval

Lemma 2. *The polynomial* $p(x) = \sum_{k=0}^{n} p_k x^k$ *is nonnegative in* $[-1, 1]$ *if and only if there exist* $S, T \succeq 0$ *such that*

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

Proof. Follows directly from the relationships between SOS and SDP, and the fact (Fekete) that

$$p(x) \geq 0 \quad \forall x \in [-1, 1] \qquad \Leftrightarrow \qquad p(x) = s(x) + t(x)(1 - x^2).$$

where $s(x), t(x)$ are SOS. ∎

This is an *SDP characterization!*

Valid measures

Lemma 3. *The vector $\mu = [\mu_0, \mu_1, \ldots, \mu_n]^T$ is a valid moments sequence for a measure in $[-1, 1]$ if and only if*

$$\begin{aligned} \mathcal{H}(\mu) &\succeq 0 \\ L_1^T \mathcal{H}(\mu) L_1 - L_2^T \mathcal{H}(\mu) L_2 &\succeq 0 \\ e_1^T \mu &= 1. \end{aligned}$$

This result follows from the previous lemma, by the duality between non-negative polynomials and moment spaces. Also, direct proofs (Hausdorff, Markov, etc).

Example: $\{\mu_0, \mu_1, \mu_2, \mu_3, \mu_4\}$ are valid moments if and only if

$$\begin{bmatrix} \mu_0 & \mu_1 & \mu_2 \\ \mu_1 & \mu_2 & \mu_3 \\ \mu_2 & \mu_3 & \mu_4 \end{bmatrix} \succeq 0, \qquad \begin{bmatrix} \mu_0 - \mu_2 & \mu_1 - \mu_3 \\ \mu_1 - \mu_3 & \mu_2 - \mu_4 \end{bmatrix} \succeq 0, \qquad \mu_0 = 1.$$

Explicit SDP

Recall

$$\min \; \beta \qquad \text{s.t.} \qquad \begin{cases} \sum_{i=0}^{n} \sum_{j=0}^{m} p_{ij} \mu_j x^i \leq \beta \quad \forall x \in [-1, 1] \\ \mu_j \text{ are moments of a } [-1, 1] \text{ measure} \end{cases}$$

We can put all conditions together in an explicit SDP:

$$\min \; \beta \qquad \text{s.t.} \qquad \begin{cases} \mathcal{H}^*(Z + L_1 W L_1^T - L_2 W L_2^T) = \beta \, e_1 - P\mu \\ \mathcal{H}(\mu) \succeq 0 \\ L_1^T \mathcal{H}(\mu) L_1 - L_2^T \mathcal{H}(\mu) L_2 \succeq 0 \\ e_1^T \mu = 1 \\ Z, W \succeq 0. \end{cases}$$

A perfect generalization of the well-known LP for matrix games!

SDP and game-theoretic duality

Self-dual: convex duality corresponds to switching players.

The dual SDP is

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

that corresponds to the mapping $(P, Z, W, \mu, \beta) \leftrightarrow (-P^T, A, B, \nu, \alpha)$.

Thm: Let $\mathcal{G}(P)$ represent the value of the game with payoff function $P(x, y)$. Then,

$$\mathcal{G}(P(x, y)) = \mathcal{G}(-P(y, x)).$$

Obtaining the measures

How to recover the optimal measures from the computed moments?

A very classical problem, with a classical solution (Hausdorff, etc)

The extreme measures are (in general) discrete, with atoms in the zeros of the polynomials $\beta - \int P(x,y)d\mu(y)$ and $\alpha + P(x,y)d\nu(x)$.

To obtain the measures, need to factorize a univariate polynomial, and solve a linear system to find the corresponding weights.

Example

The payoff function is given by:

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

Solving using the SOS approach, we obtain:

$$\mu = \begin{bmatrix} 1 \\ 0.56 \\ 1 \end{bmatrix}, \quad Z = \begin{bmatrix} 0.08 & -0.4 \\ -0.4 & 2 \end{bmatrix}, \quad W = 0.$$

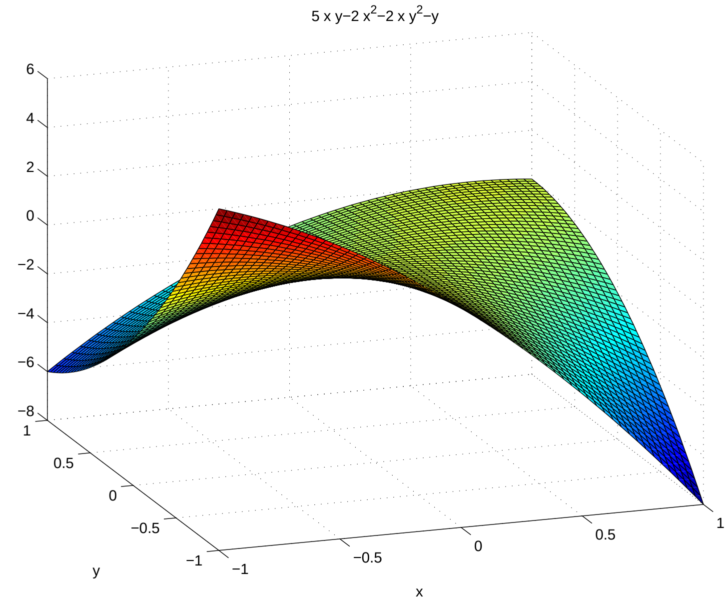

The value of the game is $\beta = -0.48$.
The optimal mixed strategies are

- P1 always picks $x = 0.2$
- P2 plays $y = 1$ with probability 0.78, and $y = -1$ with probability 0.22.

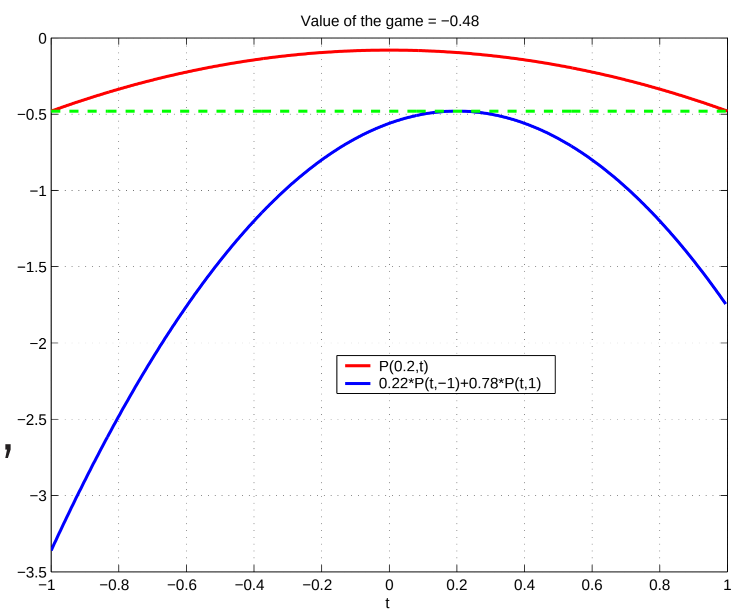

Semialgebraic games

We described in detail only the univariate case. Can extend to semialgebraic strategy sets $\mathcal{X} \subset \mathbb{R}^n, \mathcal{Y} \subset \mathbb{R}^m$.

- Two semialgebraic sets, a polynomial payoff.
- Becomes NP-hard, since both polynomial nonnegativity and moment sequences are.
- However, we can approximate arbitrarily tightly via Positivstellensatz or Schmüdgen/Putinar representations.
- In some cases, exact results ("Hilbert" games) or hard bounds on the value of the game.

General setting

Let $\mathcal{S} \subset E_1, \mathcal{T} \subset E_2$ be proper cones, and $P : E_2 \rightarrow E_1^*$. Then, the game can be solved via:

$$\begin{cases} \max & \alpha \\ \text{s.t.} & P^* x - \alpha\, \mathbf{e}_{\mathcal{T}^*} \in \mathcal{T}^* \\ & x \in \mathcal{S} \\ & \langle \mathbf{e}_{\mathcal{S}^*}, x \rangle = 1 \end{cases} \qquad \begin{cases} \min & \beta \\ \text{s.t.} & \beta\, \mathbf{e}_{\mathcal{S}^*} - P y \in \mathcal{S}^* \\ & y \in \mathcal{T} \\ & \langle \mathbf{e}_{\mathcal{T}^*}, y \rangle = 1 \end{cases}$$

The convex cones $\mathcal{S}, \mathcal{T}$ are now the moments of measures supported on the given semialgebraic sets. The dual cones $\mathcal{S}^*, \mathcal{T}^*$ correspond to nonnegative polynomials in the sets.

Semialgebraic games: solution

The sets $\mathcal{S}, \mathcal{S}^*, \mathcal{T}, \mathcal{T}^*$ can be uniformly approximated via SOS/SDP.

For positive polynomials, under mild compactness conditions (Putinar):

$$p(x) > 0 \text{ on } \{x \in \mathbb{R}^n | g_i(x) \geq 0\} \quad \Leftrightarrow \quad p(x) = s_0(x) + \textstyle\sum_i s_i(x) g_i(x)$$

where the $s_i(x)$ are SOS.

By restricting the degree in the SOS representations, we obtain

- *inner* approximations to $\mathcal{S}^*, \mathcal{T}^*$
- *outer* approximations to $\mathcal{S}, \mathcal{T}$

Converges to the value of the game. However, in general no hard bounds, unless the approximations on either side are exact.

Alternatively, use inner approximations for measures (essentially, discretization).

Other extensions

- Separable rational games
- Correlated equilibria for nonzero sum, or more players? Almost, but complicated by the fact that we cannot reduce the problem to pure moment space.

Summary

- Optimization and games, common setting
- A useful *computational framework* for games with infinite number of pure strategies.
- Relationships between SOS and SDP
- A broad generalization of known successful techniques.
- Unifies numerical and algebraic approaches.