Polynomial games: Computation of Nash and correlated equilibria

Noah Stein

Joint work with Profs. Pablo Parrilo and

Asuman Ozdaglar

Outline

- Nash equilibria
  - Finite-dimensional structure
  - Computation in zero-sum games via SOS/SDP
  - Non-zero-sum games and PPAD (conjectures)
- Correlated equilibria
  - In finite games
  - Def. and characterizations in poly. games
  - No finite dimensional characterization
  - Computation – comparison of three methods
- Conclusions and future work

Polynomial games

- Definition
  - n players, strategic form
  - Set of strategies $S_i = [-1, 1] \subset \mathbb{R}$ for each player
  - Utilities $u_i : [-1, 1]^n \to \mathbb{R}$ are polynomials
- Notation
  - Strategy for i^{th} player: $s_i \in S_i$
  - Strategy profile: $s \in S = \prod_i S_i$
  - Without player i: $s_{-i} \in \prod_{j \neq i} S_j, \quad s = (s_i, s_{-i})$
  - Probability measure over S_i: $\sigma_i \in \Delta(S_i)$

Structure
[Dresher, Karlin, Shapley 1950's]

- Utility under random (mixed) strategy is expected utility [von Neumann - Morgenstern assumption]
- Only finitely many moments matter

$$u_i(\sigma_j, s_{-j}) = \int \sum_{\alpha} c_\alpha s_j^{\alpha_j} s_{-j}^{\alpha_{-j}} d\sigma_j(s_j) = \sum_{\alpha} c_\alpha \left(\int s_j^{\alpha_j} d\sigma_j(s_j) \right) s_{-j}^{\alpha_{-j}}$$

- Players can think about choosing moments $\int s_j^k d\sigma_j(s_j)$ instead of choosing σ_j directly
- Any such moments correspond to a measure with support size at most 1 more than the j-degree of u_i

Nash equilibria

- A mixed strategy profile σ is a **Nash equilibrium** if $u_i(\sigma) \geq u_i(\tau_i, \sigma_{-i})$ for all $\tau_i \in \Delta(S_i)$
- For polynomial games this is a finite dimensional problem in the moment spaces.
- In fact we can describe the set of moments of Nash equilibria with explicit polynomial inequalities
- The Nash equilibrium strategies of zero-sum games ($n = 2$, $u_2 = -u_1$) are given by solutions to the **minimax problem**:

$$\min_{\sigma_2 \in \Delta(S_2)} \max_{\sigma_1 \in \Delta(S_1)} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Delta(S_2)} \max_{s_1 \in S_1} u_1(s_1, \sigma_2)$$

Example

Payoffs:

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

Value: -0.48

Optimal mixed strategies:

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

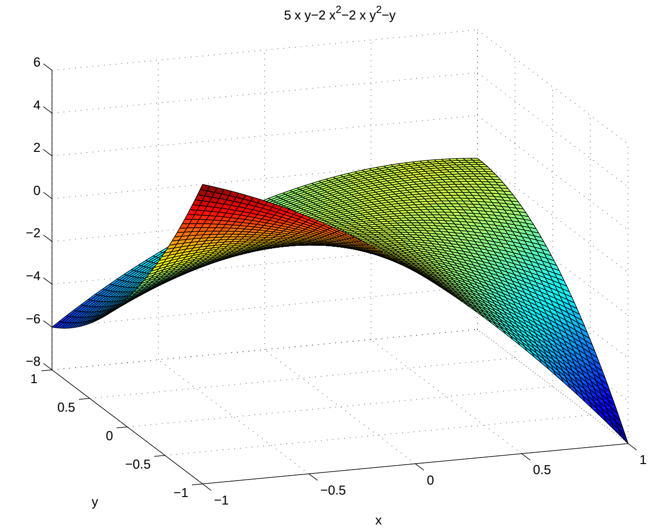

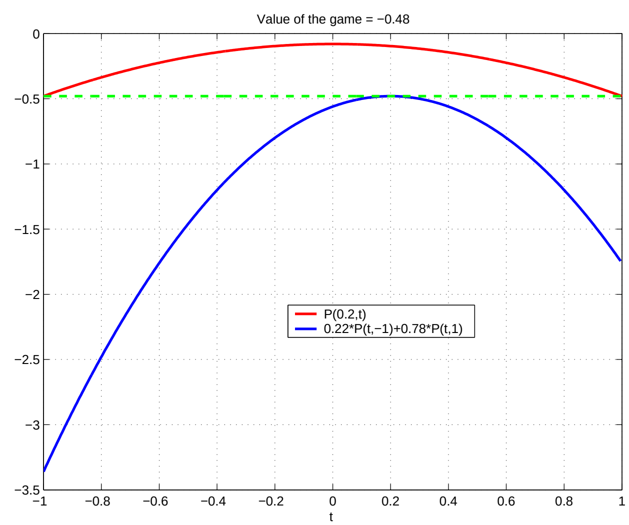

Computing minimax strategies [Parrilo 2006]

- $u_x(x,y) = \sum_{j,k} a_{jk} x^j y^k = -u_y(x,y)$
- A minimax strategy τ for player y solves

$$
\begin{array}{lll}
\min & \beta & \\
\text{s.t.} & \tau & \text{is a prob. measure on } [-1,1] \\
 & \tau_k = \int_{-1}^{1} y^k d\tau & \text{for } k \leq [y\text{-degree of } u_x] \\
 & \sum_{j,k} a_{jk} x^j \tau_k \leq \beta & \text{for all } x \in [-1,1]
\end{array}
$$

- Must describe polynomials nonnegative on $[-1,1]$ as well as moments of measures on $[-1,1]$

Sums of squares + SDP

- A polynomial $p(x)$ is ≥ 0 for all $x \in \mathbb{R}$ iff it is a sum of squares of polynomials q_k (SOS)

$$p(x) = \sum q_k^2(x) \text{ for all } x \in \mathbb{R}$$

- A polynomial $p(x)$ is ≥ 0 for all $x \in [-1, 1]$ iff there are SOS polynomials $s(x), t(x)$ such that

$$p(x) = s(x) + (1 - x^2)t(x)$$

- Coefficients of SOS polynomials can be described in a semidefinite program (SDP)

Moments of measures + SDP

- For all polynomials p, must have

$$\int p^2(x)d\tau(x) \geq 0 \quad \text{and} \quad \int (1-x^2)p^2(x)d\tau(x) \geq 0$$

- $\tau_0, \ldots, \tau_{2m}$ are the moments of a measure τ on $[-1, 1]$ (i.e. $\tau_k = \int y^k d\tau$) iff

$$\begin{bmatrix} \tau_0 & \tau_1 & \tau_2 \\ \tau_1 & \tau_2 & \tau_3 \\ \tau_2 & \tau_3 & \tau_4 \end{bmatrix} \succeq 0, \begin{bmatrix} \tau_0 - \tau_2 & \tau_1 - \tau_3 \\ \tau_1 - \tau_3 & \tau_2 - \tau_4 \end{bmatrix} \succeq 0 \quad (m = 2 \text{ case})$$

- Moments of measures on $[-1, 1]$ can be described in a semidefinite program

Higher dimensions

- What if we want to characterize, e.g., polynomials which are nonnegative on $[-1, 1]^k$ or joint moments of measures on $[-1, 1]^k$?
- There are a sequence of sufficient SDP conditions for polynomial nonnegativity starting with SOS which approach an exact condition
- Similarly there is a sequence of necessary SDP conditions for a list of numbers to be joint moments of a measure which are exact in the limit

Non-zero-sum games

- How hard is computing Nash equilibria of general polynomial games?
- Conjecture: PPAD-complete, i.e., same as finite games
- We have most of a proof that the problem is in PPAD, modulo some details about the polynomial-time solvability of SDPs
- No progress on a completeness proof, but it would be surprising if polynomial games were easier to solve than finite games

Chicken and correlated equilibria

(u_1, u_2)	Wimpy	Macho
Wimpy	$(4, 4)$	$(1, 5)$
Macho	$(5, 1)$	$(0, 0)$

- Nash equilibria (self-enforcing independent distrib.)
  - (M, W) yields utilities $(5, 1)$; (W, M) yields $(1, 5)$
  - $\left(\frac{1}{2}W + \frac{1}{2}M, \frac{1}{2}W + \frac{1}{2}M\right)$ yields expected utility $(2\frac{1}{2}, 2\frac{1}{2})$
- Correlated equilibria (self-enforcing joint distrib.)
  - e.g. $\frac{1}{2}(W, M) + \frac{1}{2}(M, W)$ yields $(3\frac{1}{2}, 3\frac{1}{2})$
  - $\frac{1}{3}(W, W) + \frac{1}{3}(W, M) + \frac{1}{3}(M, W)$ yields $(3\frac{2}{3}, 3\frac{2}{3})$

Correlated equilibria in finite games

- $u_i(t_i, s_{-i}) - u_i(s)$ is change in player i's utility when strategy t_i replaces s_i in $s = (s_1, \ldots, s_n)$
- A prob. distribution π is a **correlated equilibrium** if

$$\sum_{\{s:\ s_i = r_i\}} [u_i(t_i, s_{-i}) - u_i(s)]\pi(s) \leq 0$$

for all players i and all strategies $r_i, t_i \in S_i$

- Linear ineq. in variables $\pi(s) \Rightarrow$ linear program

Defining CE in infinite games

- Definition in literature:

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

  for all i and all (measurable) departure functions ζ_i

- Equivalent to above def. if strategy sets are finite
- Quantifier ranging over large set of functions
- Utilities composed with these complicated functions
- Is there a characterization which looks more like the finite case and doesn't have these problems?

An instructive failed attempt

- The following "characterization" fails:

$$\int_{\{s:\ s_i=r_i\}} [u_i(t_i, s_{-i}) - u_i(s)]\pi(s) \le 0$$

  for all players i and all strategies $r_i, t_i \in S_i$

- Holds for *any* continuous probability distribution π

- This condition is much weaker than correlated equilibrium

New equivalent definitions of CE

- These conditions are equivalent to the departure function definition

- For all i, all $t_i \in S_i$, and all $-1 \leq a_i \leq b_i \leq 1$,

$$\int_{\{a_i \leq s_i \leq b_i\}} [u_i(t_i, s_{-i}) - u_i(s)] d\pi \leq 0$$

- For all i, all $t_i \in S_i$, and all polynomials p,

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

Finite-dim'l characterization?

- Conditions on finitely many moments equivalent to a measure being a correlated equilibrium?
- If so, convexity $\Rightarrow$ all extreme points of correlated eq. set have uniformly bounded finite support
- Counterexample: mixed extension of matching pennies
  - Large family of extreme points with arbitrarily large finite support and infinite support constructed using ergodic theory

Examples of extreme correlated equilibrium supports

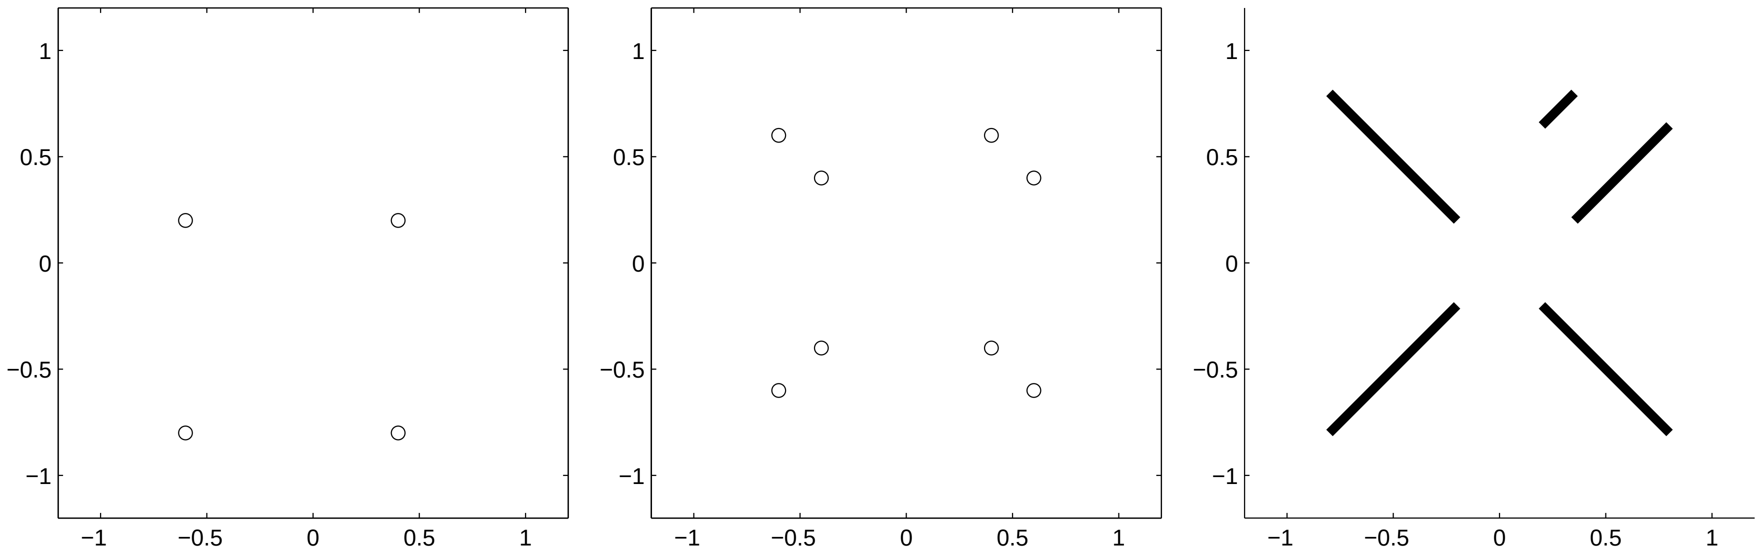

- $n = 2, \quad S_1 = S_2 = [-1, 1]$
- $u_1(s_1, s_2) = s_1 s_2 = -u_2(s_1, s_2)$

1: Static discretization (LP)

- Intended as a benchmark to judge other techniques
- Ignore polynomial structure
- Restrict strategies to finite sets $\tilde{S}_i \subset S_i$
- Compute exact correlated eq. of approximate game
- This is a sequence of LPs which converges (slowly!) to the set of correlated equilibria as the discretization gets finer.

2a: Adaptive discretization (SDP)

- Given $\tilde{S}_i^k$, optimize the following (as an SDP)

$$\begin{aligned} \min \quad & \epsilon \\ \text{s.t.} \quad & \pi \quad \text{is an } \epsilon\text{-correlated equilibrium supported} \\ & \qquad \text{on } \prod \tilde{S}_i^k \text{ which is a correlated equilibrium} \\ & \qquad\quad \text{when deviations are restricted to } \tilde{S}_i^k \end{aligned}$$

- Let ϵ^k and π^k be optimal (we're done if $\epsilon^k = 0$)
- Compute $\tilde{S}_i^{k+1}$ (following slides) and repeat
- Convergence theorem: $\epsilon^k \to 0$

2b: Finitely supported ϵ-correlated equilibria

- A probability measure π with finite support contained in $\prod \tilde{S}_i$ is an ϵ-correlated equilibrium if

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

for all i, $s_i \in \tilde{S}_i$, and $t_i \in S_i$ and

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

for all i.

2c: Adaptive discretization update steps

- Steps to compute $\tilde{S}^{k+1}$
  - For some player i, the ϵ-correlated equilibrium constraints are tight
  - Find values of t_i making these tight (free with SDP duality), add these into $\tilde{S}_i^k$ to get $\tilde{S}_i^{k+1}$
  - There are finitely many such values by polynomiality
  - For $j \neq i$, let $\tilde{S}_j^{k+1} = \tilde{S}_j^k$
- Intuitively, this adds "good" strategies for player i

3: Moment relaxation (SDP)

- No discretization
- Sequence of SDP constraints to describe moments of measures π on $[-1,1]^n$
- For fixed d, use SDP to express

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

  for all i, $t_i \in [-1,1]$, and polys. p of degree $\leq d$

- Get a nested sequence of SDPs converging to the set of correlated equilibria

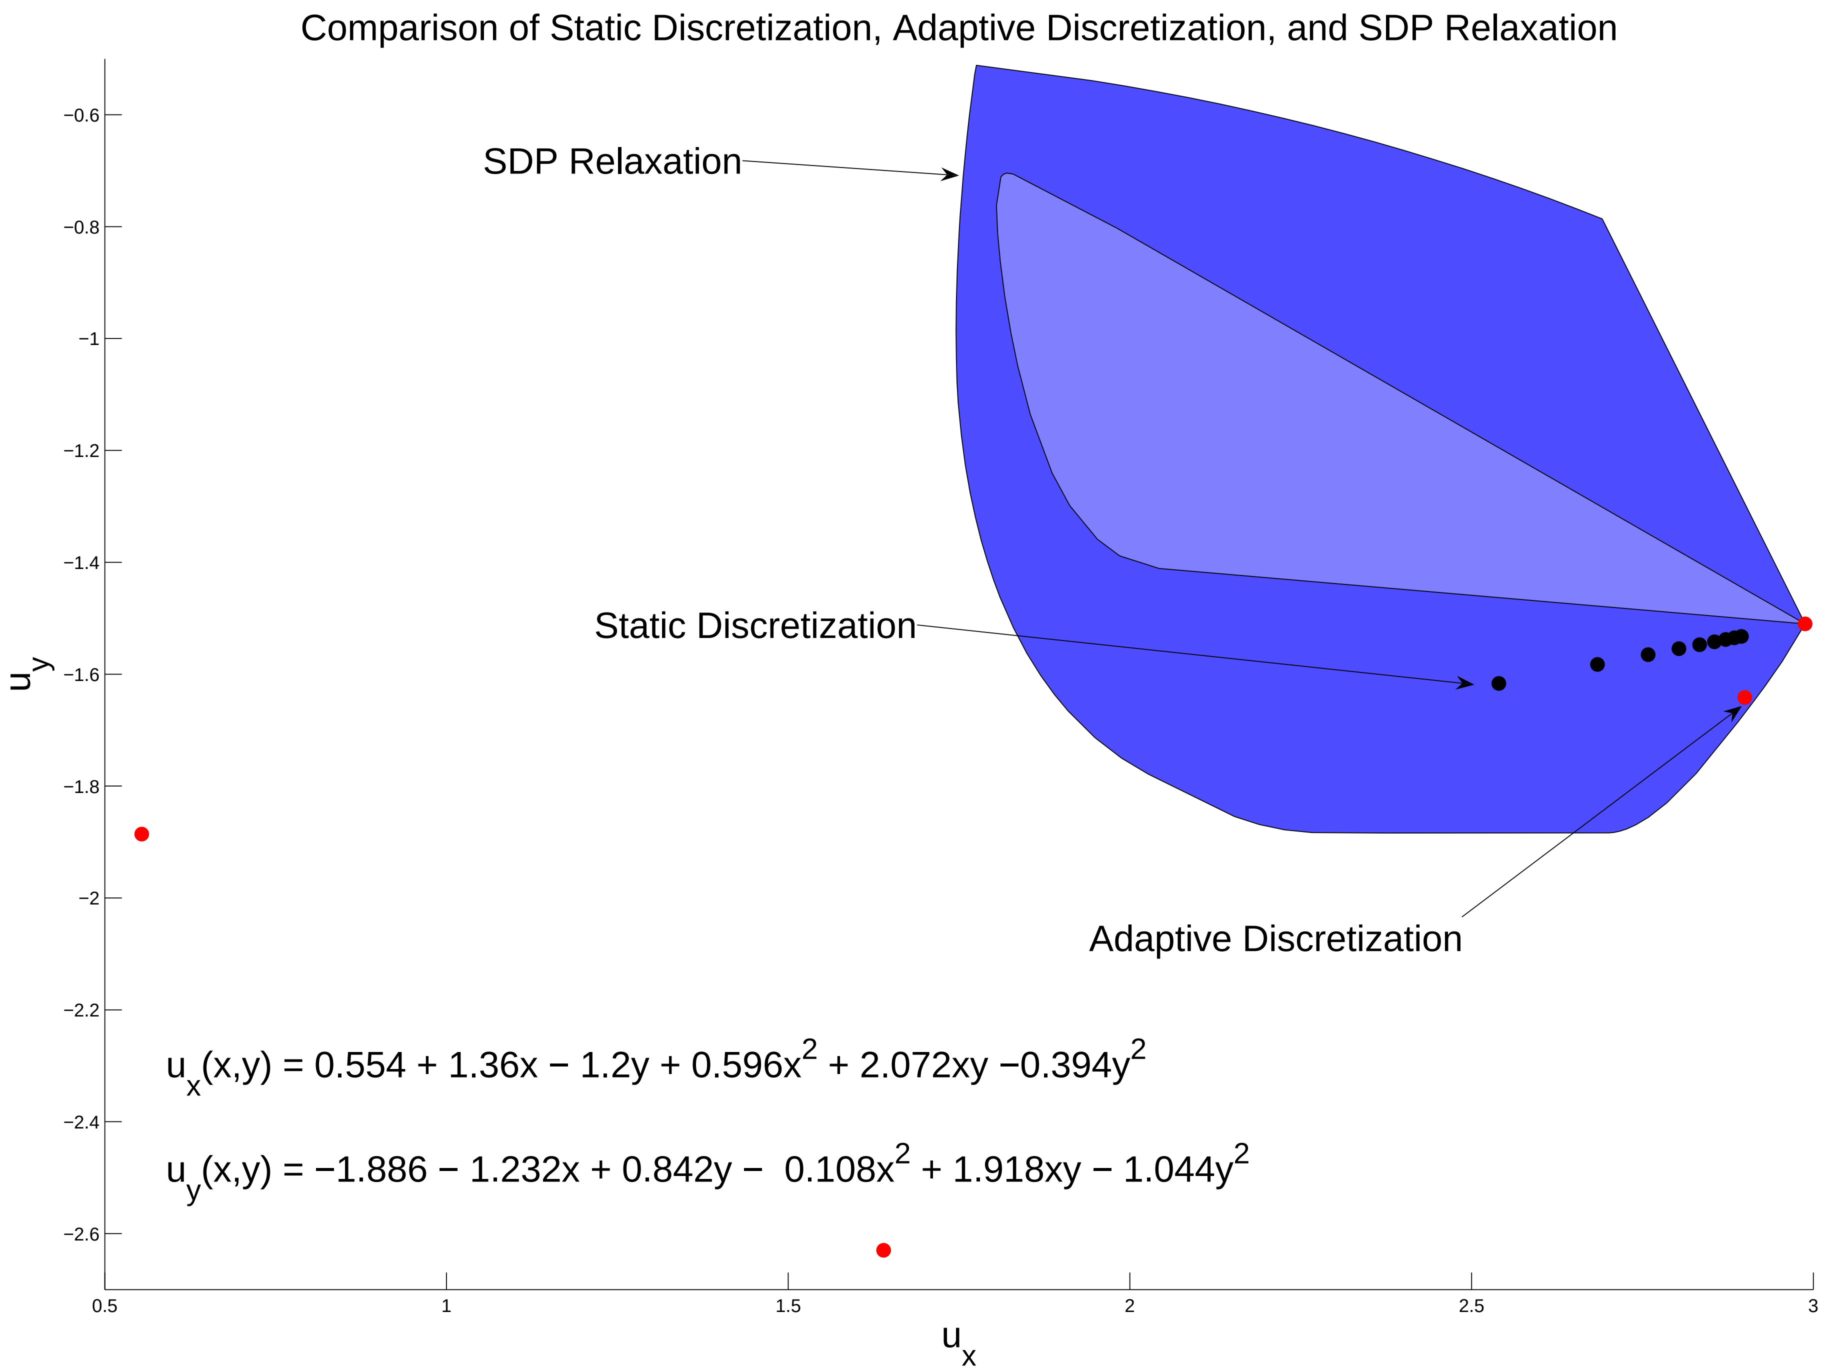
Comparison of Static Discretization, Adaptive Discretization, and SDP Relaxation
SDP Relaxation
Static Discretization
Adaptive Discretization
u_x(x,y) = 0.554 + 1.36x − 1.2y + 0.596x^2 + 2.072xy −0.394y^2
u_y(x,y) = −1.886 − 1.232x + 0.842y − 0.108x^2 + 1.918xy − 1.044y^2
u_y
u_x
−0.6
−0.8
−1
−1.2
−1.4
−1.6
−1.8
−2
−2.2
−2.4
−2.6
0.5
1
1.5
2
2.5
3

Future work

- PPAD-completeness proof
- Convergence rate of adaptive discretization
- Steering adaptive discretization toward a “good” equilibrium
- Finite algorithm for computing correlated equilibria