Characterization and Computation of Correlated Equilibria in Polynomial Games

Noah Stein

Joint work with Profs. Pablo Parrilo and Asuman Ozdaglar

Outline

- Introduction
- Characterizing correlated equilibria in finite games
- Characterizing correlated equilibria in continuous games [new]
- An overview of SDP / SOS methods
- Computing correlated equilibria in polynomial games with SDP relaxations [new]

Game theoretic setting

- Standard strategic (normal) form game
- Players (rational agents) numbered $i = 1, \ldots, n$
- Each has a set C_i of strategies s_i
- Players choose their strategies simultaneously
- Rationality: Each player seeks to maximize his own utility function $u_i : C \to \mathbb{R}$, which represents all his preferences over outcomes

Goals

- Much is known theoretically about correlated equilibria: existence (under some topological assumptions), relation to Nash equilibria, etc.
- Efficient algorithms are known for computing these when the strategy sets are finite
- Much less on games with infinite strategy sets
- Our goal is to find classes of infinite games for which correlated equilibria can be computed, as well as associated algorithms

Previous work

- Definition of correlated equilibrium [Aumann 1974]
- Rationality argument for playing correlated equilibria [Aumann 1987]
- Elementary existence proof [Hart & Schmeidler 1989]
- Efficient algorithms for computing correlated equilibria of finite games [Papadimitriou 2005]
- Sum of squares techniques [Parrilo 2000, ...]
- Algorithm for computing minimax strategies of polynomial games [Parrilo 2006]

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 games with finite strategy sets

- $u_i(t_i, s_{-i}) - u_i(s)$ is change in player i's utility when strategy t_i replaces s_i in $s = (s_1, \ldots, s_n)$

- A probability distribution π is a **correlated equilibrium** if

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

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

- No player has an incentive to deviate from his recommended strategy r_i

LP characterization

- A probability distribution π is a correlated equilibrium if and only if

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

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

- Proof: Use definition of conditional probability, pull denominator out of sum, and cancel it (reversible).

- Set of correlated equilibria of a finite game is a polytope

Departure function characterization

- A probability distribution π is a correlated equilibrium if and only if

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

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

- Proof: Let $t_i = \zeta_i(s_i)$ on previous slide and sum over $s_i \in C_i$. Conversely, define $\zeta_i(r_i) = r_i$ for all $r_i \neq s_i$ and $\zeta_i(s_i) = t_i$, then cancel terms.

- Interpretation as Nash equilibria of extended game

Continuous games

- Finitely many players (still)
- Strategy spaces C_i are compact metric spaces
- Utility functions $u_i : C \to \mathbb{R}$ are continuous
- Examples:
  - Finite games: $|C_i| < \infty$, u_i arbitrary
  - Polynomial games: $C_i = [-1, 1]$, u_i polynomial
  - ...
- Main property of continuous games: Correlated (and Nash) equilibria always exist

Defining correlated equilibria in continuous games

- Definition in literature: The probability measure π is a correlated equilibrium if

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

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
- Unknown function ζ_i inside u_i
- Is there a characterization without these problems?

An instructive failed attempt

- The following "characterization" fails:

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

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

- Holds for *any* continuous probability distribution π
- This condition is much weaker than correlated equilibrium

Simple departure functions

- A probability measure π is a correlated equilibrium if and only if

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

for all i and all simple (having finite range) measurable departure functions ζ_i

- Proof: Approximate any ζ_i as a pointwise limit of simple functions ξ_i^k (possible because C_i is compact metric). Then $u_i(\xi_i^k(s_i), s_{-i}) \to u_i(\zeta_i(s_i), s_{-i})$ pointwise by continuity of u_i and the result follows by Lebesgue's dominated convergence theorem.

No departure functions

- A probability measure π is a correlated equilibrium if and only if

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

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

- Equivalently, $-\mu_{i,t_i}$ is a positive measure for all i, t_i
- Proof: The above integral is the corresponds to the departure function given by $\zeta_i(r_i) = r_i$ for $r_i \notin B_i$ and $\zeta_i(r_i) = t_i$ for $r_i \in B_i$. Conversely, the integral for any simple departure function is a finite sum of terms of this type for some values of B_i and t_i.

Integration against test functions

- A probability measure π is a correlated equilibrium if and only if

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

for all i, $t_i \in C_i$, and $f_i : C_i \to [0, \infty)$ in some sufficiently rich class of test functions, e.g.

  – Measurable characteristic functions

  – Measurable functions

  – Continuous functions

  – Squares of polynomials (if $C_i \subset \mathbb{R}^k$)

Semidefinite programming

- A **semidefinite program (SDP)** is an optimization problem of the form

$$\begin{aligned} \min \quad & L(S) && \leftarrow L \text{ is a given linear functional} \\ \text{s.t.} \quad & T(S) = v && \leftarrow T \text{ is a given linear transformation,} \\ & && \quad v \text{ is a given vector} \\ & S \succeq 0 && \leftarrow S \text{ is a symmetric matrix} \\ & && \quad \text{of decision variables} \end{aligned}$$

- SDPs generalize linear programs and can be solved efficiently using interior point methods

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 an SDP

Multivariate SOS + SDP

- A polynomial $p(x_1, \ldots, x_k)$ is ≥ 0 for all $x \in \mathbb{R}^k$ iff there is an integer $r(p) \geq 0$ such that

$$\left(x_1^2 + \ldots + x_k^2\right)^r p(x_1, \ldots, x_k)$$

  is a sum of squares of polynomials $q_l(x_1, \ldots, x_k)$

- Sequence of sufficient SDP conditions characterizing multivariate nonnegative polynomials exactly "in the limit"

- Similar conditions for nonnegativity on $[-1, 1]^k, \ldots$

- Such conditions are generally not exact for a fixed r independent of p, with some important exceptions

Moments of measures + SDP

- If τ is a measure on $[-1,1]$, then for a polynomial p:

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

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

Moments of multivariate measures

- If τ is a measure on $[-1,1]^k$, then for a polynomial $p(x_1,\ldots,x_k)$:

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

- Requiring these conditions for all p up to a fixed degree gives a necessary semidefinite condition the joint moments of a measure must satisfy.
- Exact "in the limit"

Polynomial games

- Strategy space is $C_i = [-1, 1]$ for all players i
- Utilities u_i are multivariate polynomials
- Finitely supported equilibria always exist, with explicit bounds on support size [1950s; SOP 2006]
- Minimax strategies and values can be computed by semidefinite programming [Parrilo 2006]

Computing corr. equil. in poly. games: Naive attempt (LP)

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

Computing corr. equil. in poly. games by SDP relaxation

- Describe moments of measures π on $[-1,1]^n$ with a sequence of necessary SDP conditions which become sufficient in the limit
- For fixed d, we want to use SDP to express

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

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

- Shown above: for each d this is a necessary condition for π to be a correlated equilibrium which becomes sufficient as $d \to \infty$

Computing corr. equil. in poly. games by SDP relaxation (II)

- Define $S_i^d \in \mathbb{R}^{d \times d}[s_i]$ by $\left[S_i^d\right]_{jk} = s_i^{j+k}$, $0 \leq j, k < d$
- A polynomial is a square of a polynomial if and only if it can be written as $c^\top S_i^d c$ for some $c \in \mathbb{R}^d$
- Let $M_i^d(t_i) = \int S_i^d [u_i(t_i, s_{-i}) - u_i(s)] d\pi \in \mathbb{R}^{d \times d}[t_i]$
- Then we wish to constrain π to satisfy $c^\top M_i^d(t_i) c \leq 0$ for all $c \in \mathbb{R}^d$ and $t_i \in [-1, 1]$, i.e. $M_i^d(t_i) \preceq 0$ for $t_i \in [-1, 1]$

Computing corr. equil. in poly. games by SDP relaxation (III)

- $M_i^d(t_i)$ is a matrix whose entries are univariate polynomials in t_i with coefficients which are affine in the decision variables of the problem (the joint moments of π)
- This can be expressed exactly as an SDP constraint for any d (this is one of those special cases in which the condition is exact, even though there are multiple variables, t_i and c)
- Putting it all together we get a nested sequence of SDPs converging to the set of correlated equilibria

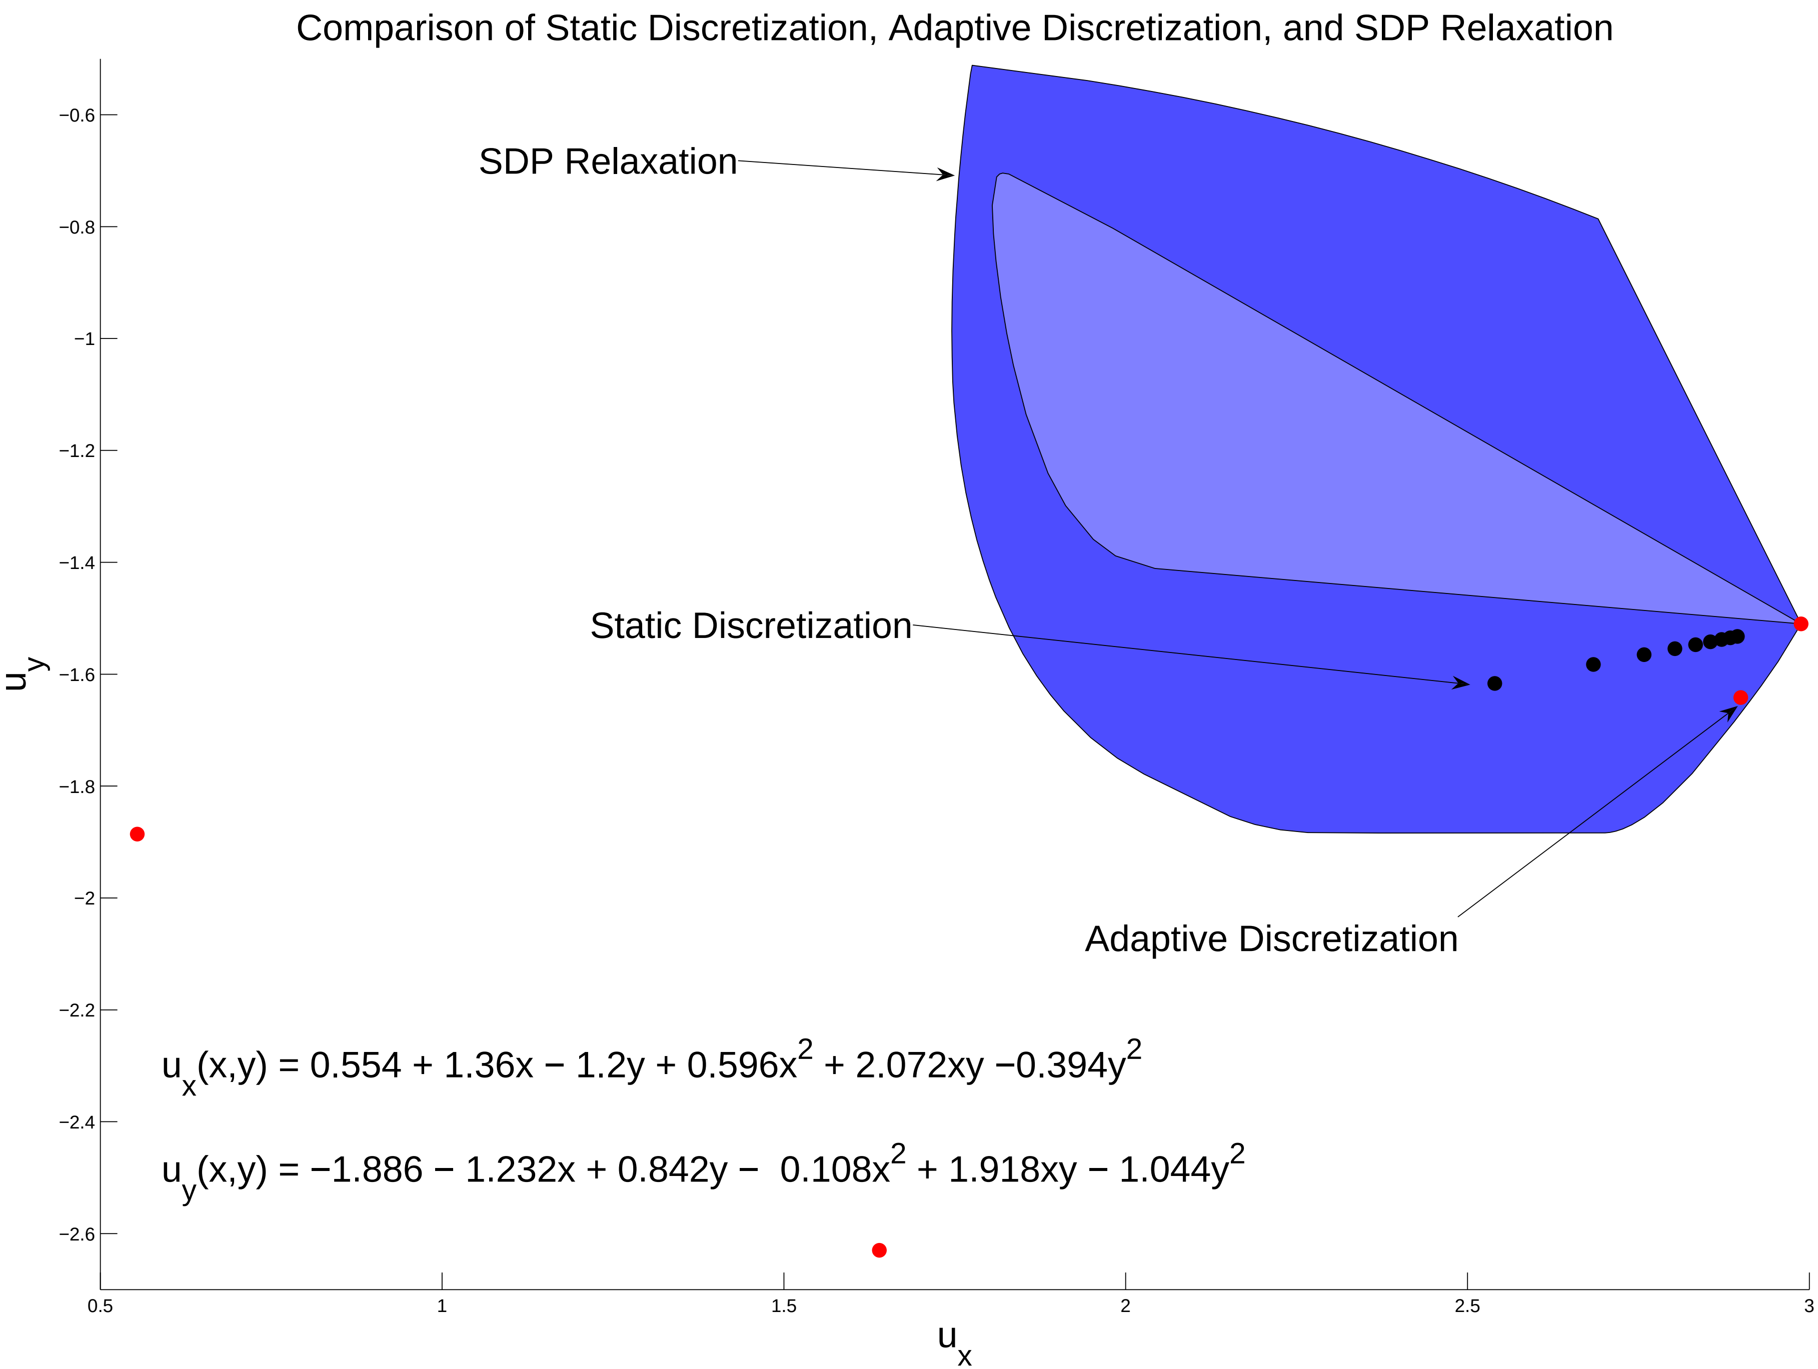
Comparison of Static Discretization, Adaptive Discretization, and SDP Relaxation
SDP Relaxation
Static Discretization
Adaptive Discretization
$u_x(x,y) = 0.554 + 1.36x - 1.2y + 0.596x^2 + 2.072xy - 0.394y^2$
$u_y(x,y) = -1.886 - 1.232x + 0.842y - 0.108x^2 + 1.918xy - 1.044y^2$
u_y
u_x
−0.6
−0.8
−1
−1.2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

Conclusions

- New characterizations of correlated equilibria in infinite games
- First algorithms for computing correlated equilibria in any class of infinite games

Acknowledgements

- My advisors: Profs. Asuman Ozdaglar and Pablo Parrilo
- Prof. Muhamet Yildiz

Time for a better discretization algorithm?

Approximate correlated equilibria

- The probability measure π is an ϵ-correlated equilibrium if

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

  for all i and all (measurable) departure functions ζ_i

- Same as definition of correlated equilibrium if $\epsilon = 0$
- We will be interested in the case when the support of π is a finite set

Characterizing finitely supported ϵ-correlated equilibria

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

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

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

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

for all i.

Applying SOS / SDP

- Given a polynomial game and a finite support set $\tilde{C}_i \subset [-1,1]$ for each player, the condition that π be a probability measure on $\tilde{C}$ and an ϵ-correlated equilibrium can be written in an SDP
- First constraint says a univariate polynomial in t_i with coefficients linear in the $\pi(s)$ and ϵ_{i,s_i} is ≥ 0 on $[-1,1]$, hence is expressible exactly in an SDP
- Second constraint is linear, so usable in SDP
- Conditions to make π a prob. measure also linear
- Always feasible if ϵ can vary

Adaptive discretization

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

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

- Let ϵ^k and π^k be an optimal solution
- If $\epsilon^k = 0$ then halt
- Otherwise, compute $\tilde{C}_i^{k+1}$ as described on next slide and repeat

Adaptive discretization (II)

- Steps to compute $\tilde{C}^{k+1}$
  - 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 $\tilde{C}_i^k$ to get $\tilde{C}_i^{k+1}$
  - For $j \neq i$, let $\tilde{C}_j^{k+1} = \tilde{C}_j^k$
- By construction $\tilde{C}_i^{k+1} \supsetneq \tilde{C}_i^k$ because if not then $\epsilon = 0$, a contradiction

Proof that $\epsilon^k \to 0$

- If not, there is a subseq. along which $\epsilon^{k_l} \geq \epsilon > 0$.
- There is some player i for whom $\tilde{C}_i^{k_l+1} \supsetneq \tilde{C}_i^{k_l}$ for infinitely many l.
- Assume WLOG that $\tilde{C}_i^{k_l+1} \supsetneq \tilde{C}_i^{k_l}$ for all l.
- Cover C_i with finitely many open balls B_i^j such that $|u_i(s_i, s_{-i}) - u_i(t_i, s_{-i})| \leq \frac{\epsilon}{2}$ for all $s_{-i} \in C_{-i}$ whenever s_i and t_i are in the same B_i^j.
- For some l, all the balls B_i^j which will ever contain a point of some $\tilde{C}_i^{k_l}$ already do.

Proof that $\epsilon^k \to 0$ (II)

- Let $t_{i,s_i} \in \tilde{C}_i^{k_l+1}$ make the ϵ^{k_l}-correlated equilibrium constraints tight.
- Let $r_{i,s_i} \in \tilde{C}_i^{k_l}$ be in the same B_i^j as t_{i,s_i} for each s_i.
- Then we get a contradiction:

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

Adaptive discretization (III)

- We did not use the polynomial structure of the u_i in the convergence proof, just continuity
- Used polynomiality to convert the optimization problem into an SDP
- Can also do this conversion if the u_i are rational or even piecewise rational (and continuous)
- Solutions of such games are surprisingly complex – the Cantor measure arises as the unique Nash equilibrium of a game with rational u_i [Gross 1952]
- Now we have a way to approximate these!