Commitment to Correlated Strategies

Dima

[Vincent Conitzer, Dmytro Korzhyk. Commitment to Correlated Strategies. In AAAI-2011.]

Games in Normal Form

Player 2

		L	R
Player 1	U	1, 1	3, 0
	D	0, 0	2, 1

A player's **strategy** is a distribution over the player's actions

An **outcome of the game** is an entry in the matrix

A **strategy profile** is a pair of strategies (pure or randomized)

Nash equilibrium

An NE is a strategy profile in which no player has an incentive to deviate.

Player 2

Player 1		L	R
	U	1, 1	3, 0
	D	0, 0	2, 1

Computing a Nash Equilibrium

Iterated dominance works in this case

		Player 2	
		L	R
Player 1	U	1, 1	3, 0
	D	0, 0	2, 1

Dominated

Generally, there is no known polytime algorithm

[PPAD-completeness: Daskalakis, Goldberg & Papadimitriou '06; Chen & Deng '06; NP-hardness of NE with certain properties: Gilboa & Zemel '89; Conitzer & Sandholm '08]

Stackelberg model

- Suppose the row player (the leader) can **commit** to a strategy

Follower

Leader		L	R
	U	1, 1	3, 0
	D	0, 0	2, 1

The leader benefits from commitment!

Commitment to a mixed strategy

- Suppose the leader commits to (2/3 Down, 1/3 Up)

			Follower: L	Follower: R
Leader	1/3	U	1, 1	3, 0
	2/3	D	0, 0	2, 1

- Commitment to a mixed strategy benefits the leader even more
- The optimal strategy to commit to is (50%-eps, 50%+eps)
- Can be computed in polytime [Conitzer & Sandholm ’06, von Stengel & Zamir ‘10]

Applications of the Stackelberg model

- Resource allocation for
  airport security
  [Pita et al., AI Magazine ‘09]

- Scheduling of federal
  air marshals
  [Tsai et al., AAMAS ‘09]

- GUARDS system for TSA
  resource allocation
  [Pita et al., AAMAS ’11]

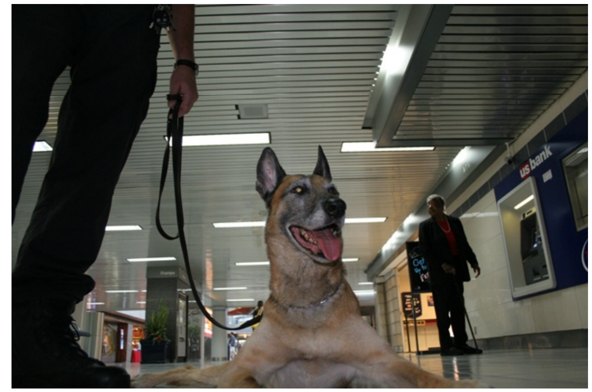

Photo STL airport

Photo AP

LP1: Computing a Stackelberg strategy

[Conitzer and Sandholm ’06, von Stengel and Zamir ‘10]

- Given the leader’s strategy $\mathrm{P}(s_1)$, the follower maximizes $\mathbf{E}[u_2|\mathrm{P}(s_1)]$
- There is always a pure-strategy best response
- Idea: write an LP for each best-response s_2^*, choose the max leader’s utility among the feasible LPs

Objective: leader’s utility

$$Maximize \sum_{s_1} u_1(s_1, s_2^*) p(s_1)$$

Subject to the follower’s rationality

$$\forall s_2: \quad \sum_{s_1} u_2(s_1, s_2^*) p(s_1) \geq \sum_{s_1} u_2(s_1, s_2) p(s_1)$$

$$\sum_{s_1} p(s_1) = 1$$

New idea: Commitment to a correlated strategy

- The leader draws from a distribution over the outcomes

Follower

Leader	L	R
U	1, 1 (40%)	3, 0 (20%)
D	0, 0 (10%)	2, 1 (30%)

- The follower only gets to know the column
- The follower should have no incentive to deviate
- We will look for a correlated strategy that maximizes the leader’s utility

Equivalence to Stackelberg

Proposition 1. There exists an optimal correlated strategy to commit to in which the follower always gets the same recommendation.

Proof of Proposition 1

Break the correlated strategy into two components:

$$P(s_1, s_2) = P(s_2)\, P(s_1 | s_2)$$

Follower's rationality: each s_2 is a best-response to $P(s_1|s_2)$

The leader can rearrange $P(s_2)$ without breaking the follower's rationality condition

Set $P(s_2^*) = 1$,
where s_2^* maximizes $E[u_1|s_2]$

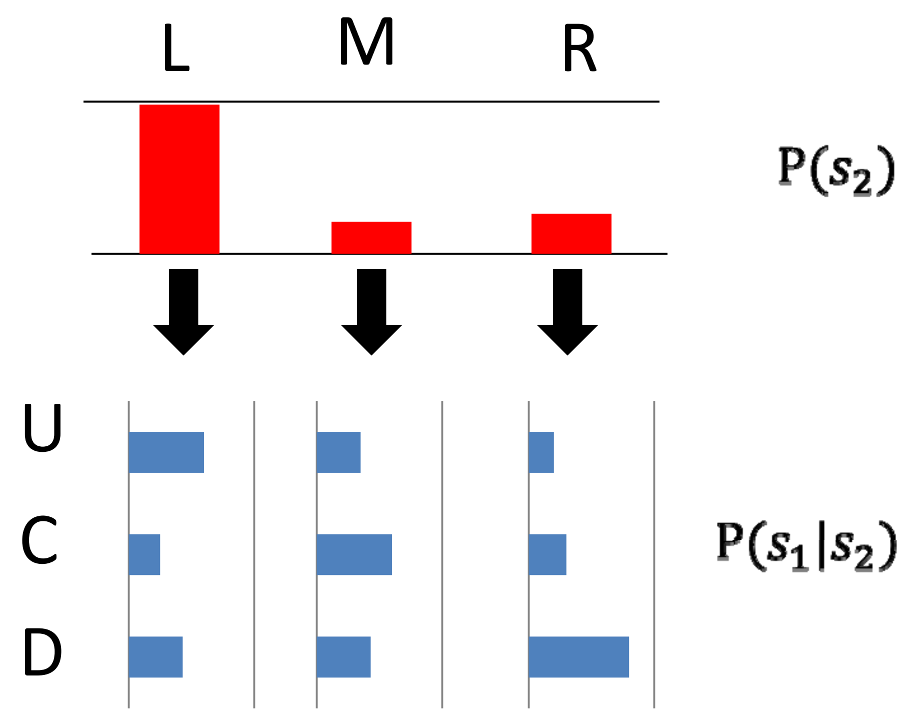

LP2 for computing an optimal correlated strategy to commit to

Objective – leader's utility

$$Maximize \sum_{s_1} \sum_{s_2} u_1(s_1, s_2) p(s_1, s_2)$$

Follower's rationality

$$\forall s_2, s_2': \sum_{s_1} u_2(s_1, s_2) p(s_1, s_2) \geq \sum_{s_1} u_2(s_1, s_2') p(s_1, s_2)$$

$$\sum_{s_1} \sum_{s_2} p(s_1, s_2) = 1$$

- We can use this LP to compute an optimal Stackelberg strategy!

Experimental evaluation

- The single LP actually runs faster than LP1, MIP for many game classes (on 50x50 games, using CPLEX, GAMUT games [Nudelman et al. '04]):

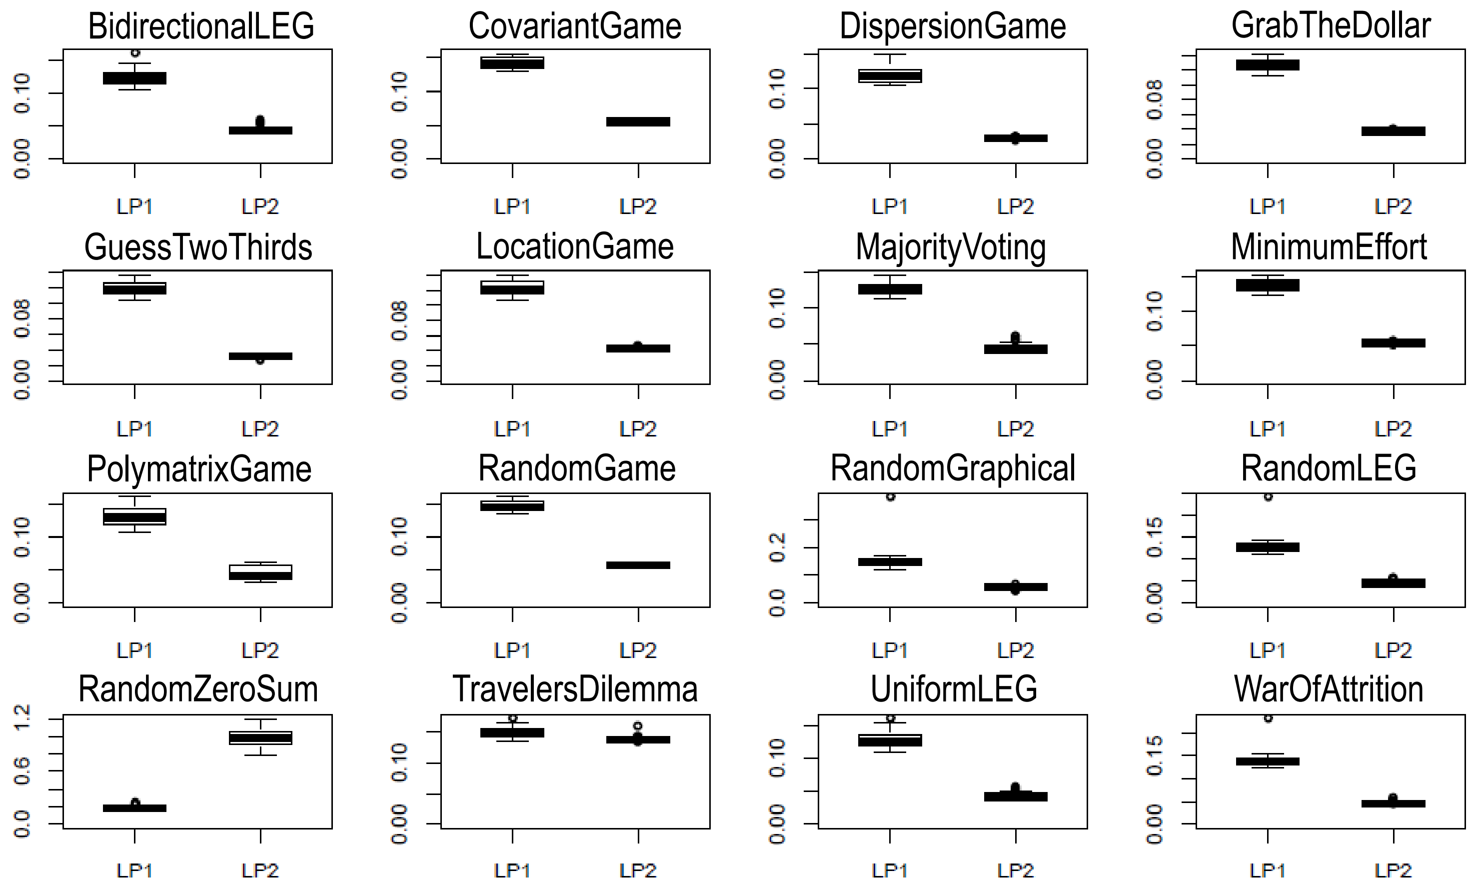

- Downside: the single LP uses more memory.

Correlated equilibrium

- A 3rd party proposes a distribution over the outcomes
- Now, both players should have no incentive to deviate

Player 1's rationality $\forall s_1, s_1': \sum_{s_2} u_1(s_1, s_2) p(s_1, s_2) \geq \sum_{s_2} u_1(s_1', s_2) p(s_1, s_2)$

Player 2's rationality $\forall s_2, s_2': \sum_{s_1} u_2(s_1, s_2) p(s_1, s_2) \geq \sum_{s_1} u_2(s_1, s_2') p(s_1, s_2)$

$$\sum_{s_1} \sum_{s_2} p(s_1, s_2) = 1$$

Stackelberg vs Correlated Equilibrium

[von Stengel and Zamir '10]

Corollary 1. The leader's Stackelberg utility is at least as high as the leader's utility in any correlated equilibrium of the game.

Commitment to correlated strategies for n>2 players

- A generalization of LP2
- With n>2 players, the optimal correlated strategy to commit to may not be a product distribution

Game class \♯ players	**2**		**3**		**4**	
	P	D	P	D	P	D
BidirectionalLEG	1	.96	.9	.86	.84	.84
CovariantGame	1	.48	.64	.6	.68	.68
DispersionGame	1	1	1	1	1	1
GuessTwoThirdsAve	1	1	0	0	0	0
MajorityVoting	1	.88	1	1	1	1
MinimumEffortGame	1	1	1	1	1	1
RandomGame	1	.42	.16	.08	.02	.02
RandomGraphicalGame	1	.4	.22	.1	.02	.02
RandomLEG	1	1	.92	.92	.02	.02
TravelersDilemma	1	0	1	1	.02	.02
UniformLEG	1	.96	.88	.86	.02	.02

P= product distribution

D= degenerate distribution

Example: Commitment to a correlated strategy with 3 players

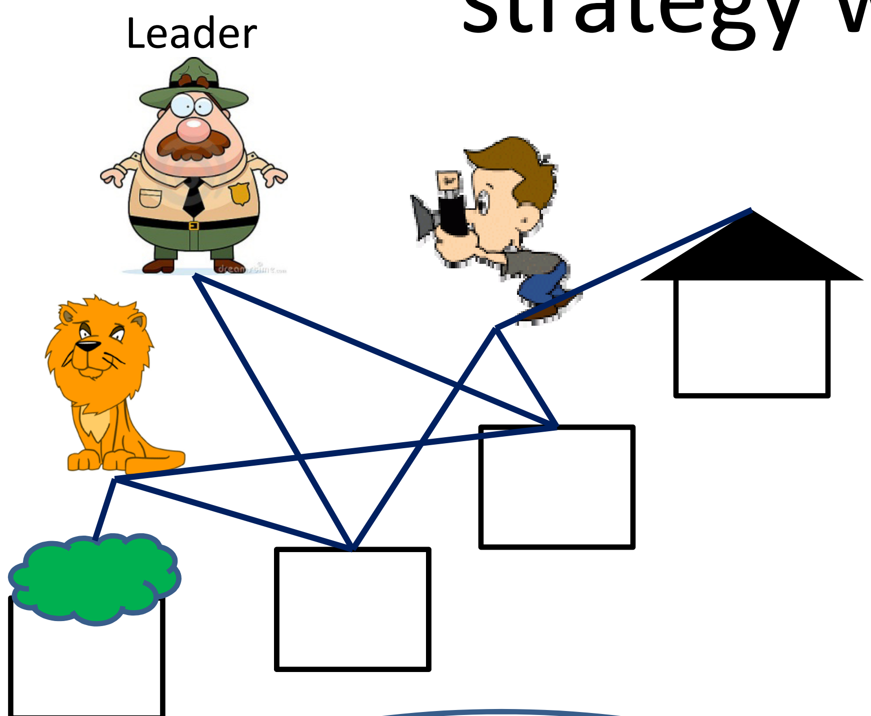

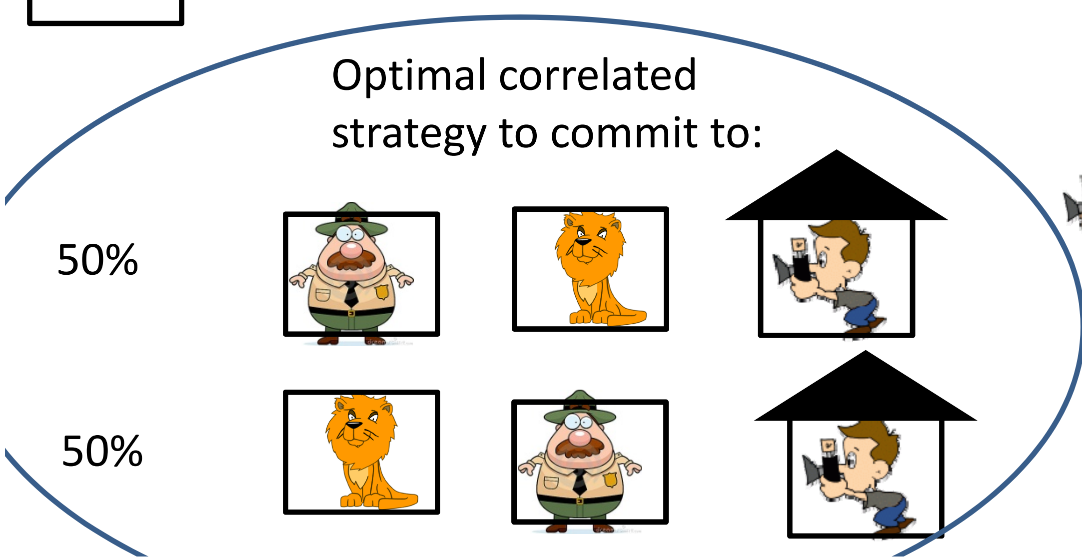

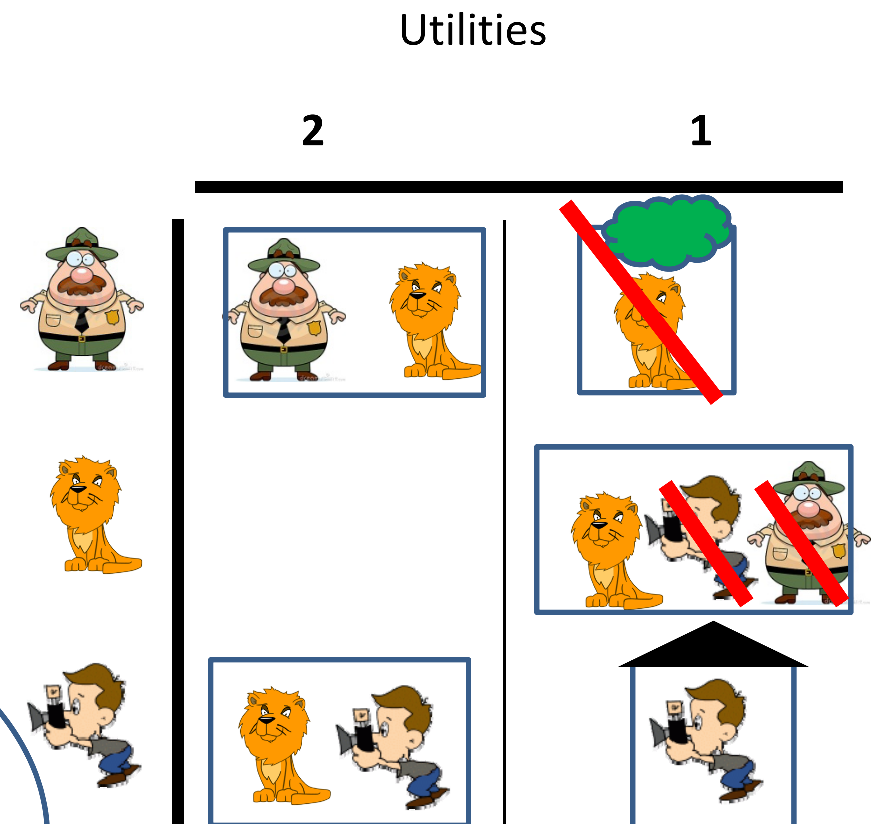

Different from Stackelberg and CE!

Advantages of commitment to a correlated equilibrium

- Same as Stackelberg for n=2 players
- Well defined for n>2 players (unlike the Stackelberg model)
- Easy to compute for any number of players

Overview of contributions

- A single LP for Stackelberg (for 2 players)
- Easy proof of relationship between Stackelberg and correlated equilibrium (for 2 players)
- Model for commitment to a correlated strategy (for n≥2 players)
- LP for commitment to a correlated strategy (for n≥2 players)

Thank you!