Finding Optimal Mixed Strategies to Commit to in Security Games

Vincent Conitzer
Departments of Computer Science and Economics
Duke University

Co-authors on various parts:
@Duke: Dmytro (Dima) Korzhyk, Josh Letchford, Kamesh Munagala, Ron Parr
@USC: Zhengyu Yin, Chris Kiekintveld, Milind Tambe
@CMU: Tuomas Sandholm

Newsweek National News

The Element of Surprise

To help combat the terrorism threat, officials at Los Angeles Inter Airport are introducing a bold new idea into their arsenal: random of security checkpoints. Can game theory help keep us safe?

WEB EXCLUSIVE

By Andrew Murr
Newsweek
Updated: 1:00 p.m. PT Sept 28, 2007

Security forces work the sidewalk

Sept. 28, 2007 - Security officials at Los Angeles International Airport now have a new weapon in their fight against terrorism: complete, baffling randomness. Anxious to thwart future terror attacks in the early stages while plotters are casing the airport, LAX security patrols have begun using a new software program called ARMOR, NEWSWEEK has learned, to make the placement of security checkpoints completely unpredictable. Now all airport security officials have to do is press a button labeled "Randomize," and they can throw a sort of digital cloak of invisibility over where they place the cops' antiterror checkpoints on any given day.

What is game theory?

- Game theory studies settings where multiple parties (agents) each have
  - different preferences (utility functions),
  - different actions that they can take
- Each agent's utility (potentially) depends on all agents' actions
  - What is optimal for one agent depends on what other agents do
    - Very circular!
- Game theory studies how agents can rationally form beliefs over what other agents will do, and (hence) how agents should act
  - Useful for acting as well as predicting behavior of others

Penalty kick example

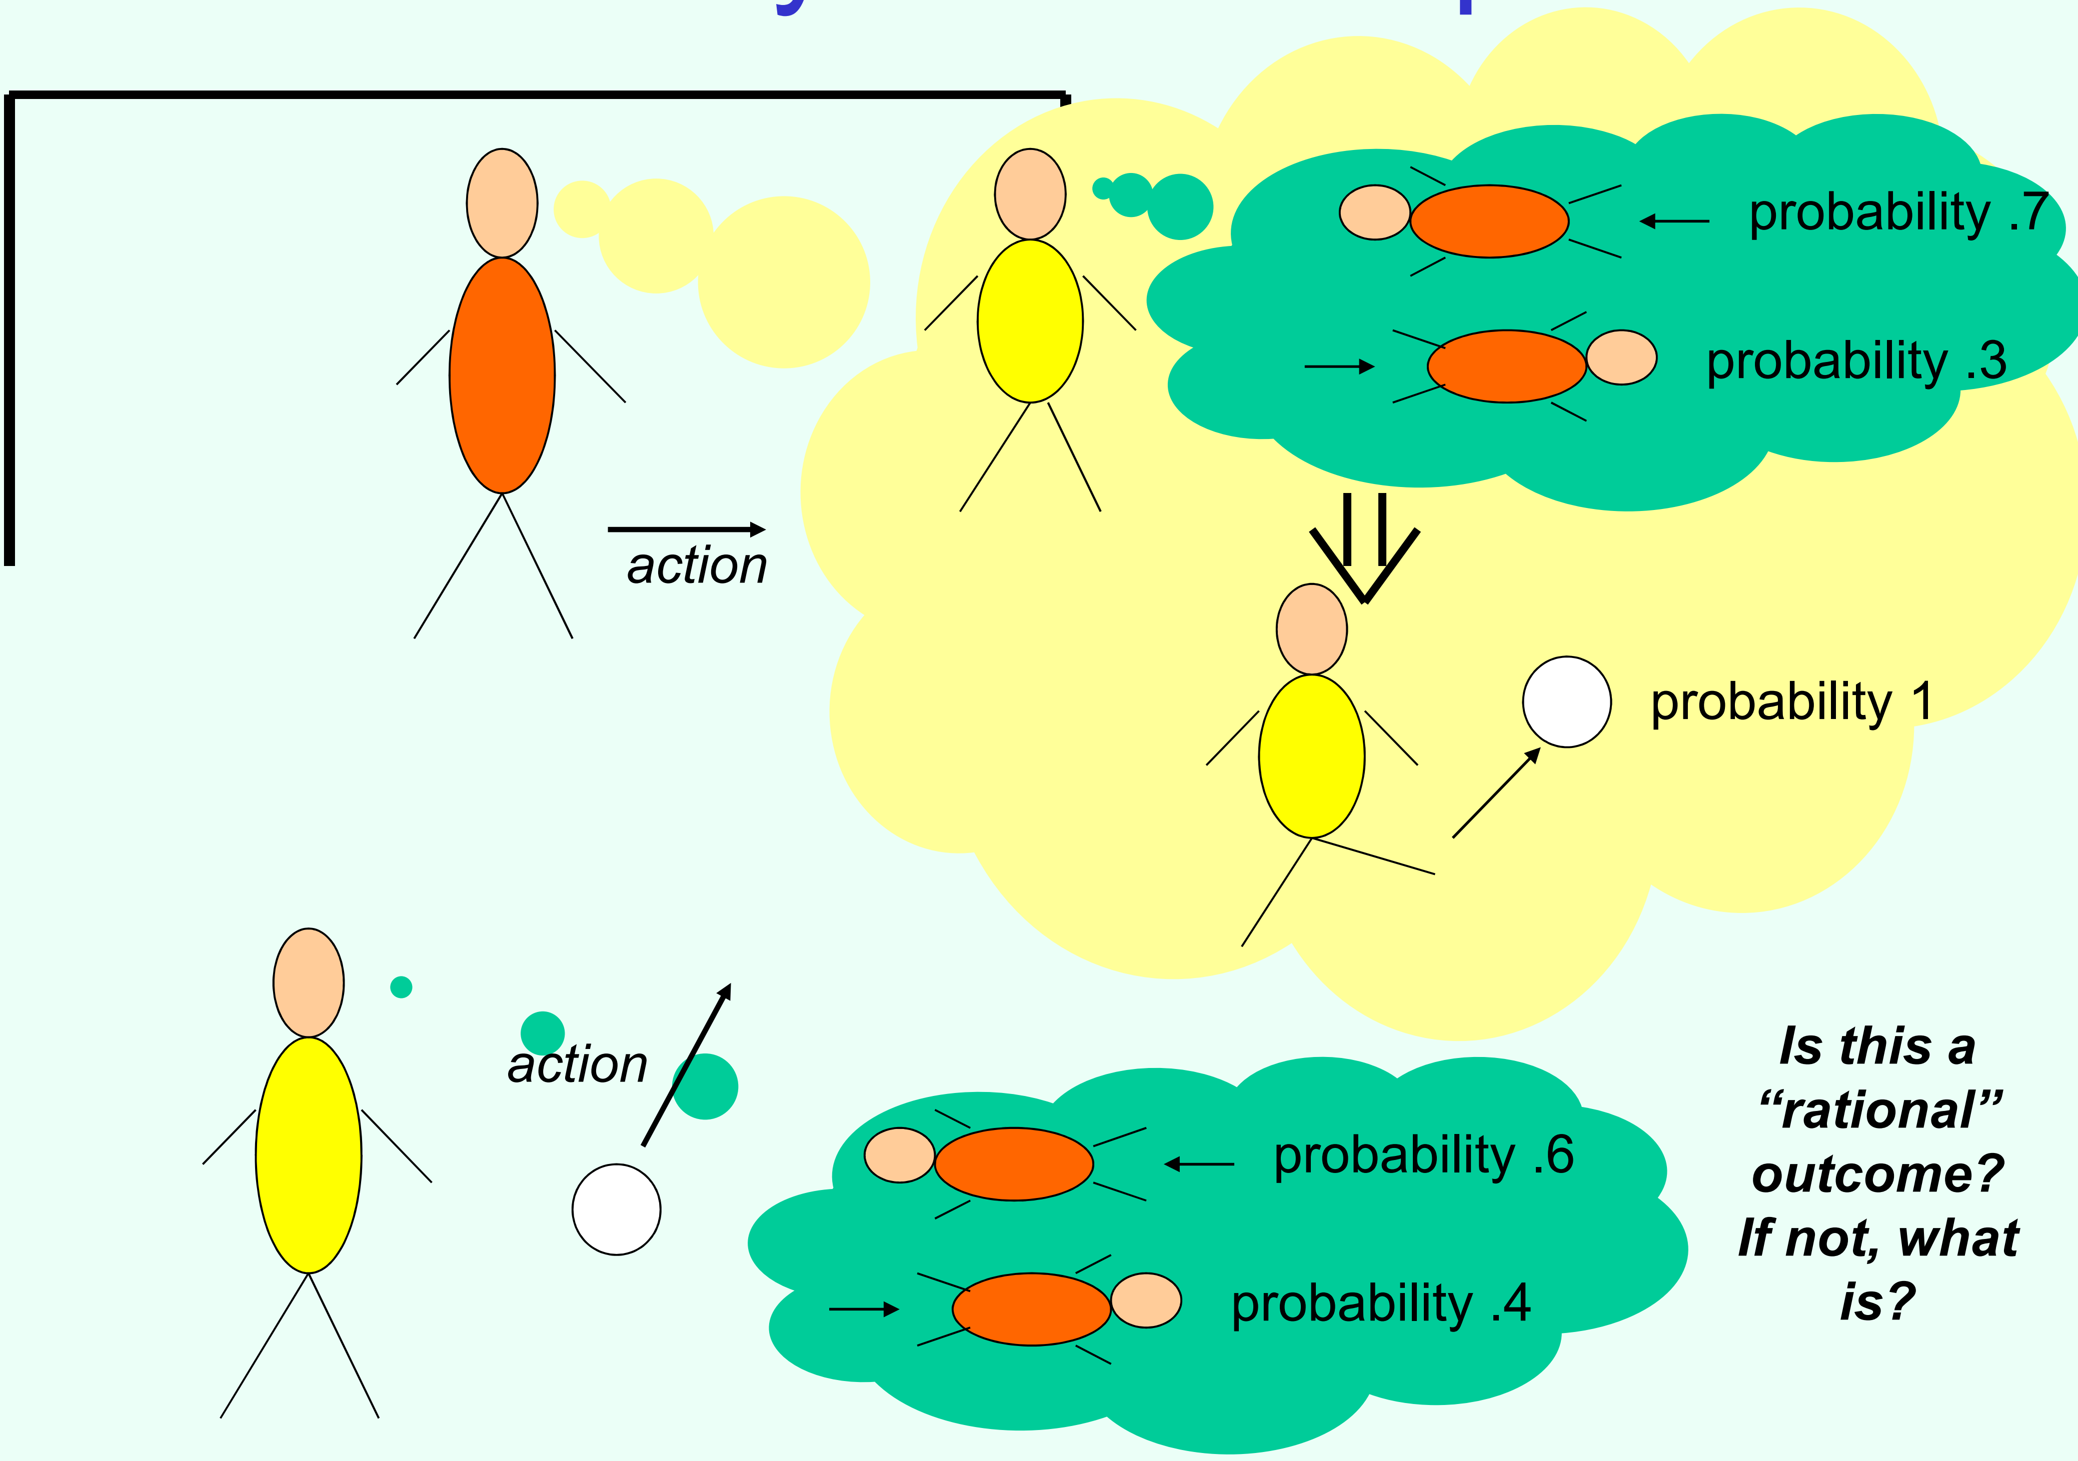

Is this a "rational" outcome? If not, what is?

Rock-paper-scissors

Column player aka. player 2 chooses a column

		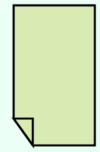	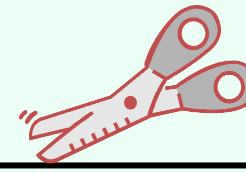
	0, 0	-1, 1	1, -1
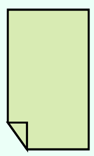	1, -1	0, 0	-1, 1
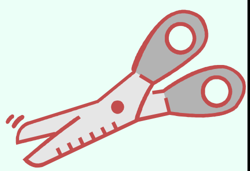	-1, 1	1, -1	0, 0

Row player aka. player 1 chooses a row

A row or column is called an action or (pure) strategy

Row player's utility is always listed first, column player's second

Zero-sum game: the utilities in each entry sum to 0 (or a constant)
Three-player game would be a 3D table with 3 utilities per entry, etc.

Matching pennies (~penalty kick)

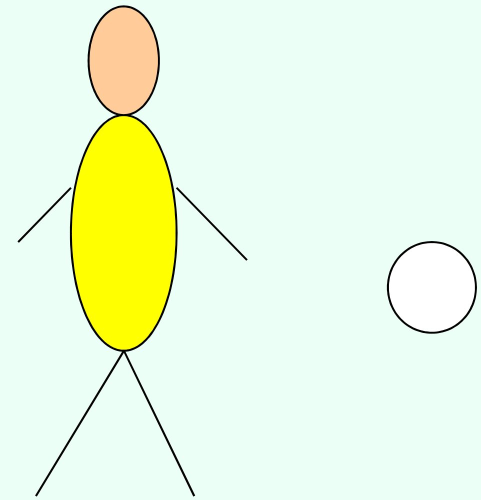

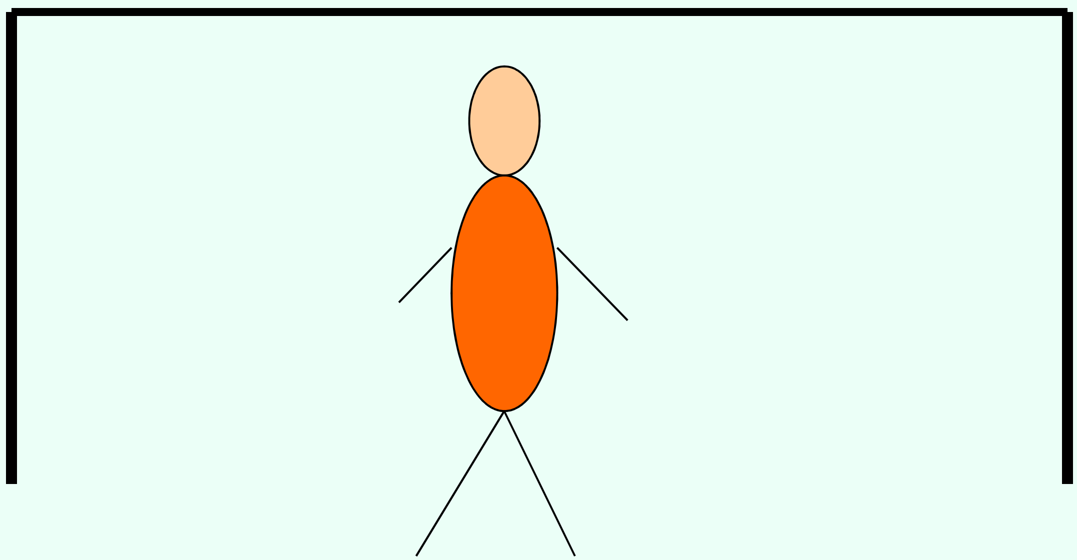

	L	R
L	1, -1	-1, 1
R	-1, 1	1, -1

“Chicken”

- Two players drive cars towards each other
- If one player goes straight, that player wins
- If both go straight, they both die

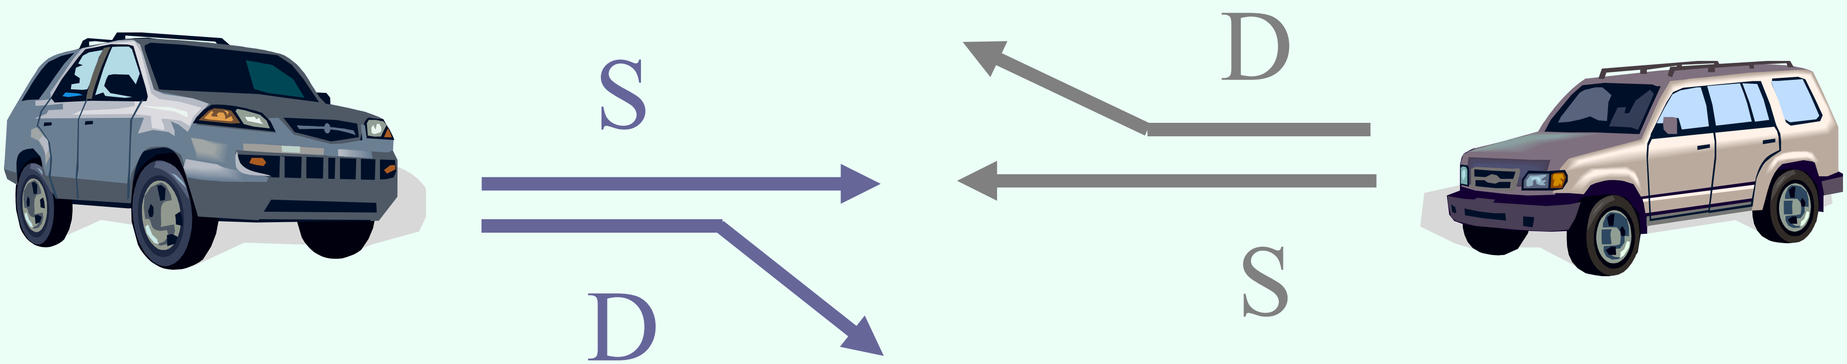

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

not zero-sum

How to play matching pennies

		Them L	R
Us	L	1, -1	-1, 1
	R	-1, 1	1, -1

- Assume opponent knows our strategy…
  - hopeless?
- … but we can use randomization
- If we play L 60%, R 40%...
- … opponent will play R…
- … we get .6*(-1) + .4*(1) = -.2
- What's optimal for us? What about rock-paper-scissors?

Matching pennies with a sensitive target

Them

Us		L	R
	L	1, -1	-1, 1
	R	-2, 2	1, -1

- If we play 50% L, 50% R, opponent will attack L
  - We get .5*(1) + .5*(-2) = -.5
- What if we play 55% L, 45% R?
- Opponent has choice between
  - L: gives them .55*(-1) + .45*(2) = .35
  - R: gives them .55*(1) + .45*(-1) = .1
- We get -.35 > -.5

Matching pennies with a sensitive target

		Them L	R
Us	L	1, -1	-1, 1
	R	-2, 2	1, -1

- What if we play 60% L, 40% R?
- Opponent has choice between
  - L: gives them .6*(-1) + .4*(2) = .2
  - R: gives them .6*(1) + .4*(-1) = .2
- We get -.2 either way
- This is the maximin strategy
  - Maximizes our minimum utility

Let's change roles

Them

		L	R
Us	L	1, -1	-1, 1
	R	-2, 2	1, -1

- Suppose **we** know **their** strategy
- If they play 50% L, 50% R,
  - We play L, we get .5*(1)+.5*(-1) = 0
- If they play 40% L, 60% R,
  - If we play L, we get .4*(1)+.6*(-1) = -.2
  - If we play R, we get .4*(-2)+.6*(1) = -.2
- This is the minimax strategy

von Neumann's minimax theorem [1927]: maximin value = minimax value (~LP duality)

Minimax theorem falls apart in nonzero-sum games

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Let's say we play S
- Most they could hurt us is by playing S as well
- But that is not rational for them
- If we can commit to S, they will play D
  - Commitment advantage

Nash equilibrium [Nash 1950]

- A profile (= strategy for each player) so that no player wants to deviate

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- This game has another Nash equilibrium in mixed strategies – both play D with 80%

The presentation game

		Presenter	
		Put effort into presentation (E)	*Do not put effort into presentation (NE)*
Audience	*Pay attention (A)*	2, 2	-8, -7
	Do not pay attention (NA)	0, -1	0, 0

- Pure-strategy Nash equilibria: (A, E), (NA, NE)
- Mixed-strategy Nash equilibrium:

  ((1/10 A, 9/10 NA), (4/5 E, 1/5 NE))

  - Utility 0 for audience, -7/10 for presenter
  - Can see that some equilibria are strictly better for both players than other equilibria, i.e. some equilibria Pareto-dominate other equilibria

Properties of Nash equilibrium in two-player games

- In zero-sum games, same thing as maximin/minimax strategies
- Any (finite) game has at least one Nash equilibrium [Nash 1950]
- PPAD-complete to compute one Nash equilibrium [Daskalakis, Goldberg, Papadimitriou 2006; Chen & Deng, 2006]
- NP-hard & inapproximable to compute the "best" Nash equilibrium [Gilboa & Zemel 1989; Conitzer & Sandholm 2008]

Nash isn’t optimal if one player can commit

2, 1	4, 0
1, 0	3, 1

Unique Nash equilibrium

- Suppose the game is played as follows:
  - Player 1 commits to playing one of the rows,
  - Player 2 observes the commitment and then chooses a column
- Optimal strategy for player 1: commit to Down

Commitment as an extensive-form game

- For the case of committing to a pure strategy:

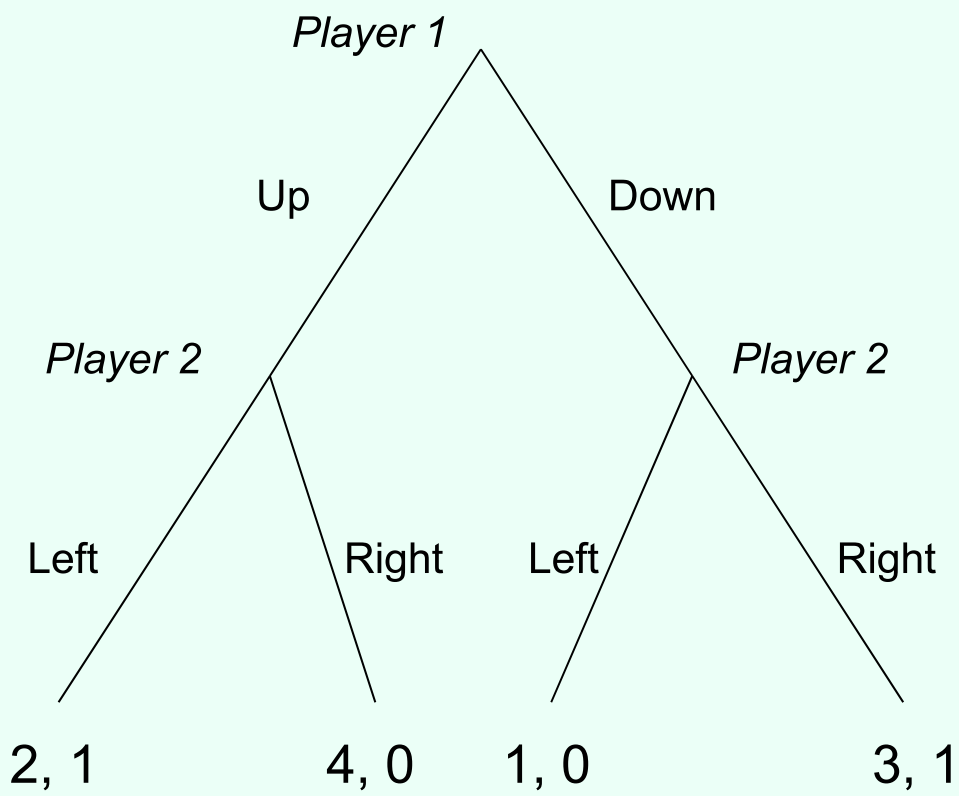

Commitment to mixed strategies

.49 / .5	2, 1	4, 0
.51 / .5	1, 0	3, 1

- Assume follower breaks ties in leader's favor
  - In generic games this is the unique SPNE outcome of the extensive-form game [von Stengel & Zamir 2010]

– We will also refer to this as a Stackelberg strategy

Commitment as an extensive-form game…

- … for the case of committing to a mixed strategy:

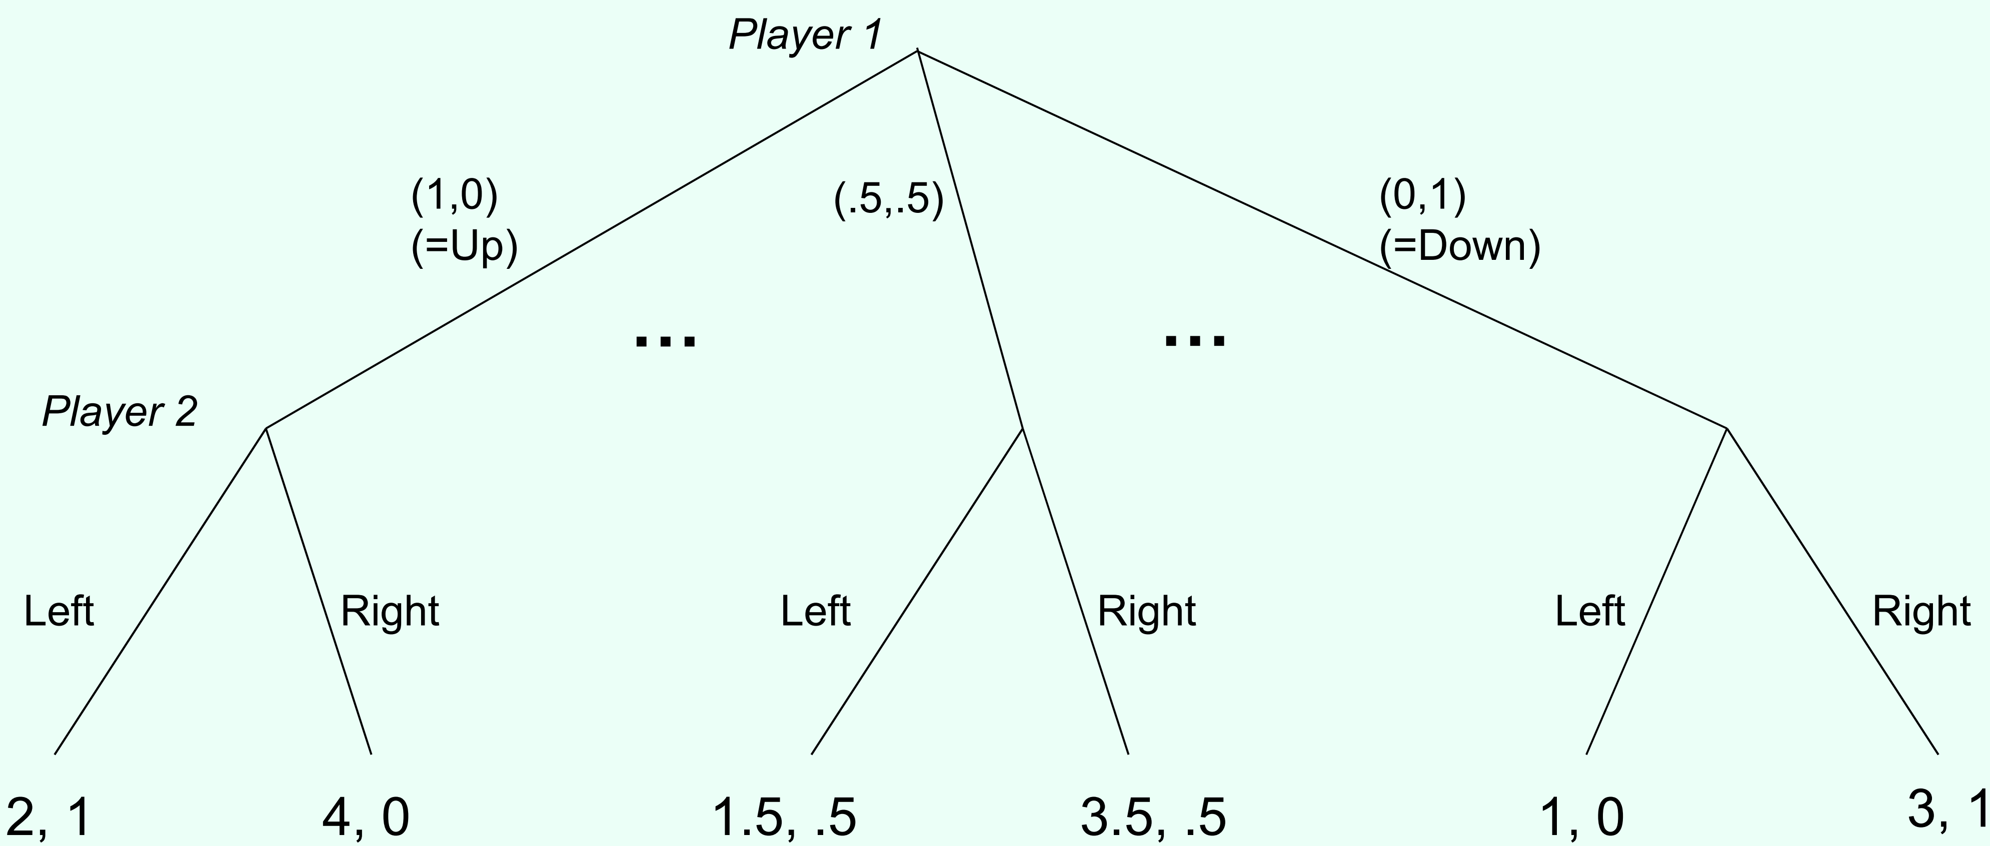

- Economist: Just an extensive-form game, nothing new here
- Computer scientist: Infinite-size game! Representation matters

Computing the optimal mixed strategy to commit to

[Conitzer & Sandholm 2006, von Stengel & Zamir 2010]

- Separate LP for every possible follower's action *t**

$$\text{Maximize} \sum_s p_s U_l(s, t^*)$$ Leader utility

Subject to

$$\sum_s p_s = 1$$ Distributional constraint

$$\forall t : \sum_s p_s U_f(s, t) \leq \sum_s p_s U_f(s, t^*)$$ Follower optimality

- Choose *t** for which the LP is feasible and has the highest objective. The leader plays the corresponding strategy <p_s>.

Easy polynomial-time algorithm for two players

[Conitzer & Sandholm 2006; von Stengel & Zamir 2010]

- For every column t separately, we solve separately for the best mixed row strategy (defined by p_s) that induces player 2 to play t
- maximize $\Sigma_s \, p_s \, u_1(s, t)$
- subject to

  for any t’, $\Sigma_s \, p_s \, u_2(s, t) \geq \Sigma_s \, p_s \, u_2(s, t')$

  $\Sigma_s \, p_s = 1$
- (May be infeasible)
- Pick the t that is best for player 1

Visualization

	L	C	R
U	0,1	1,0	0,0
M	4,0	0,1	0,0
D	0,0	1,0	1,1

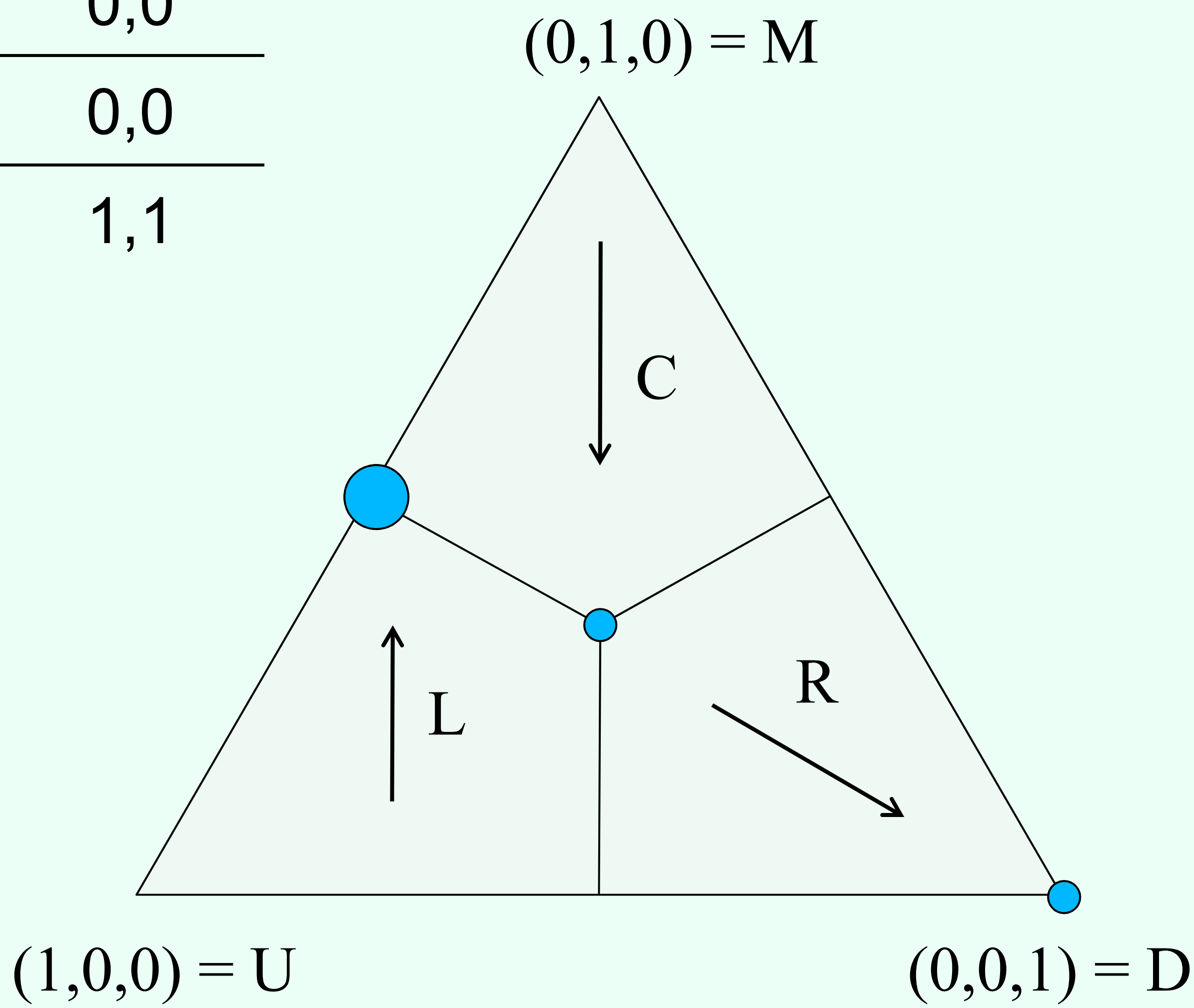

Observations about commitment to a mixed strategy in a two-player game

- Coincides with minimax strategies in zero-sum games
- Leader's payoff always at least as good as in any Nash equilibrium (see [von Stengel & Zamir 2010])
  - Can simply commit to the Nash equilibrium strategy
  - Follower breaks ties in your favor
  - Actually at least as good as any correlated equilibrium
    - Close relationship to LP for correlated equilibrium [Conitzer 2010 draft]
- No equilibrium selection problem
- Natural notion of approximation

(a particular kind of) Bayesian games

leader utilities

2	4
1	3

follower utilities (type 1)

1	0
0	1

probability .6

follower utilities (type 2)

1	0
1	3

probability .4

Multiple types - visualization

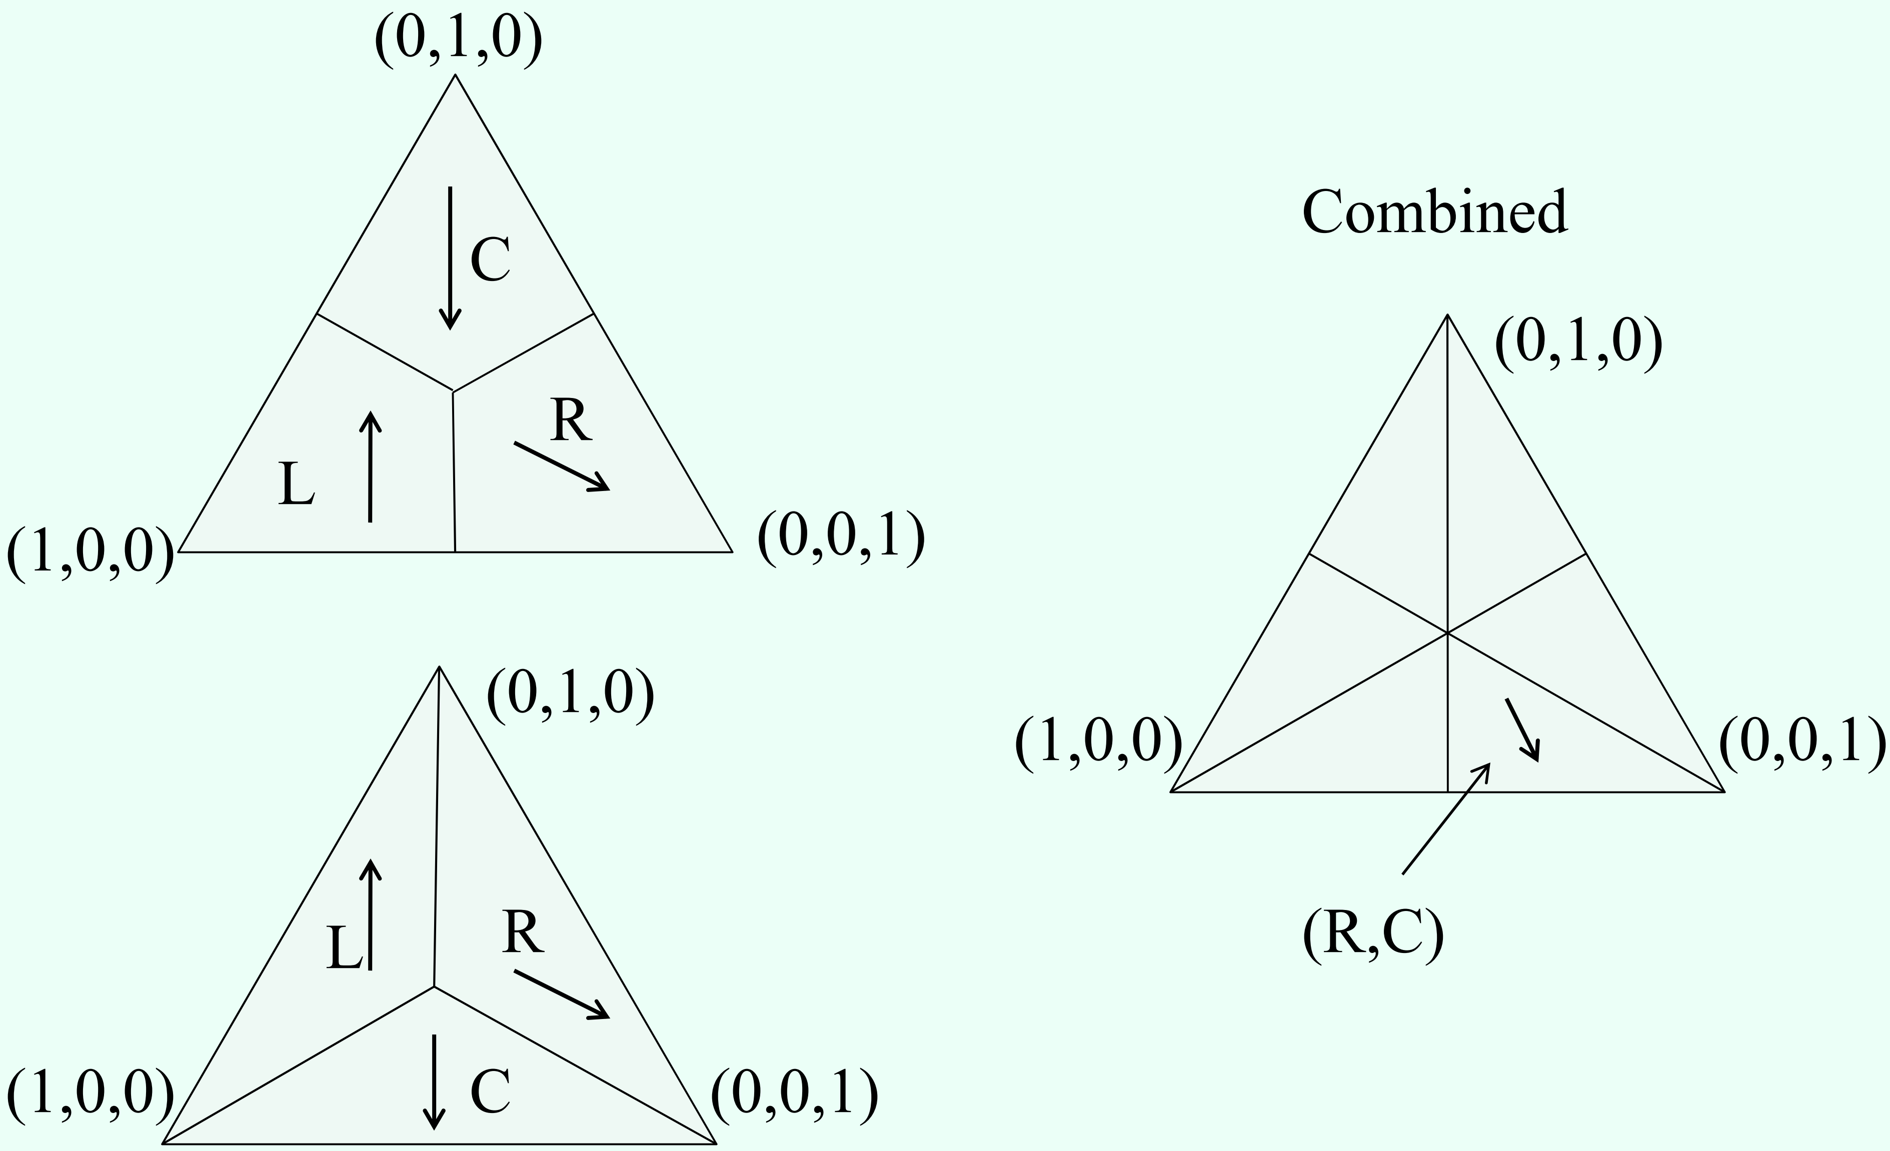

	2 players	$\geq$ **3 players**
normal-form	$O(\#outcomes)$	$O(\#outcomes \cdot \#players)$
Bayesian, 1-type leader	$O(\#outcomes \cdot \#types)$	NP-hard
Bayesian, 1-type follower	NP-hard	NP-hard
Bayesian (general)	NP-hard	NP-hard

Results for commitment to pure strategies. (With more than 2 players, the "follower" is the last player to commit, the "leader" is the first.)

	2 players	$\geq$ **3 players**
normal-form	one LP-solve per follower action	NP-hard
Bayesian, 1-type leader	NP-hard	NP-hard
Bayesian, 1-type follower	one LP-solve per follower action	NP-hard
Bayesian (general)	NP-hard	NP-hard

Results for commitment to mixed strategies. (With more than 2 players, the "follower" is the last player to commit, the "leader" is the first.)

LAX techniques

[Paruchuri et al. 2008, Pita et al. 2009]

- Uses Bayesian games framework
- Mixed integer programming formulation for solving Bayesian games optimally
  - Much faster than converting game to normal form, solving that

(In)approximability

[Letchford, Conitzer, Munagala 2009]

- (#types)-approximation: pick one type uniformly at random, optimize for it using LP approach
  - … or (deterministic) optimize for every type separately, pick best
- Can't do any better in polynomial time, unless P=NP
  - Reduction from INDEPENDENT-SET
- For adversarially chosen types, cannot decide in polynomial time whether it is possible to guarantee positive utility, unless P=NP
  - Again, a MIP formulation can be given

Reduction from independent set

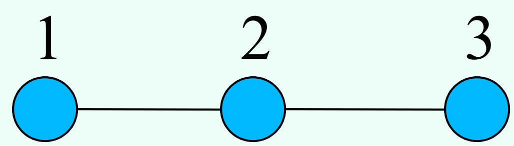

leader utilities

	A	B
a_l^1	1	0
a_l^2	1	0
a_l^3	1	0

follower utilities (type 1)

	A	B
a_l^1	3	1
a_l^2	0	10
a_l^3	0	1

follower utilities (type 2)

	A	B
a_l^1	0	10
a_l^2	3	1
a_l^3	0	10

follower utilities (type 3)

	A	B
a_l^1	0	1
a_l^2	0	10
a_l^3	3	1

Switching topics: Learning

- Single follower type
- Unknown follower payoffs
- Repeated play: commit to mixed strategy, see follower's (myopic) response

	L	R
U	1,?	3,?
D	2,?	4,?

Visualization

	L	C	R
U	0,1	1,0	0,0
M	4,0	0,1	0,0
D	0,0	1,0	1,1

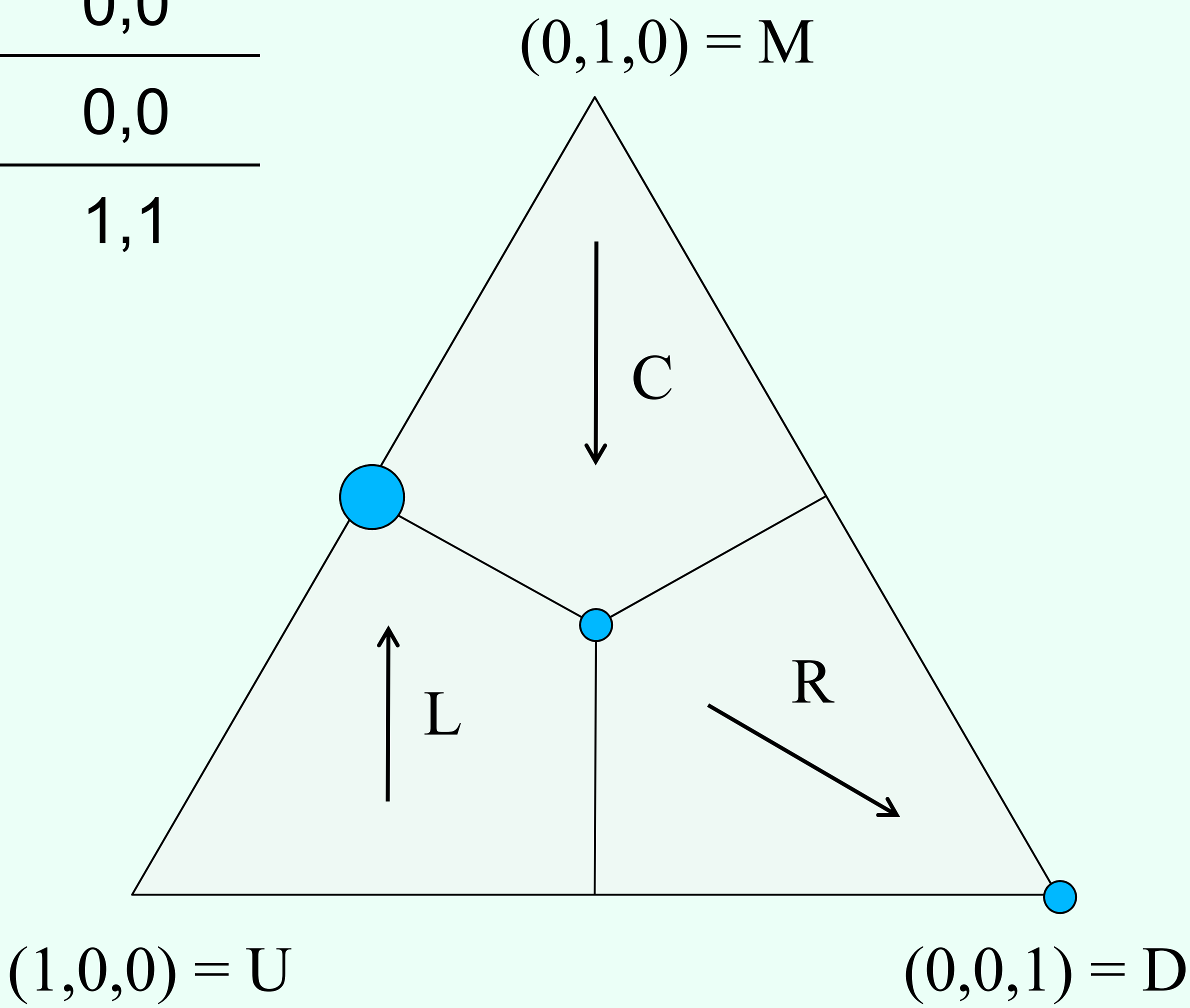

Sampling

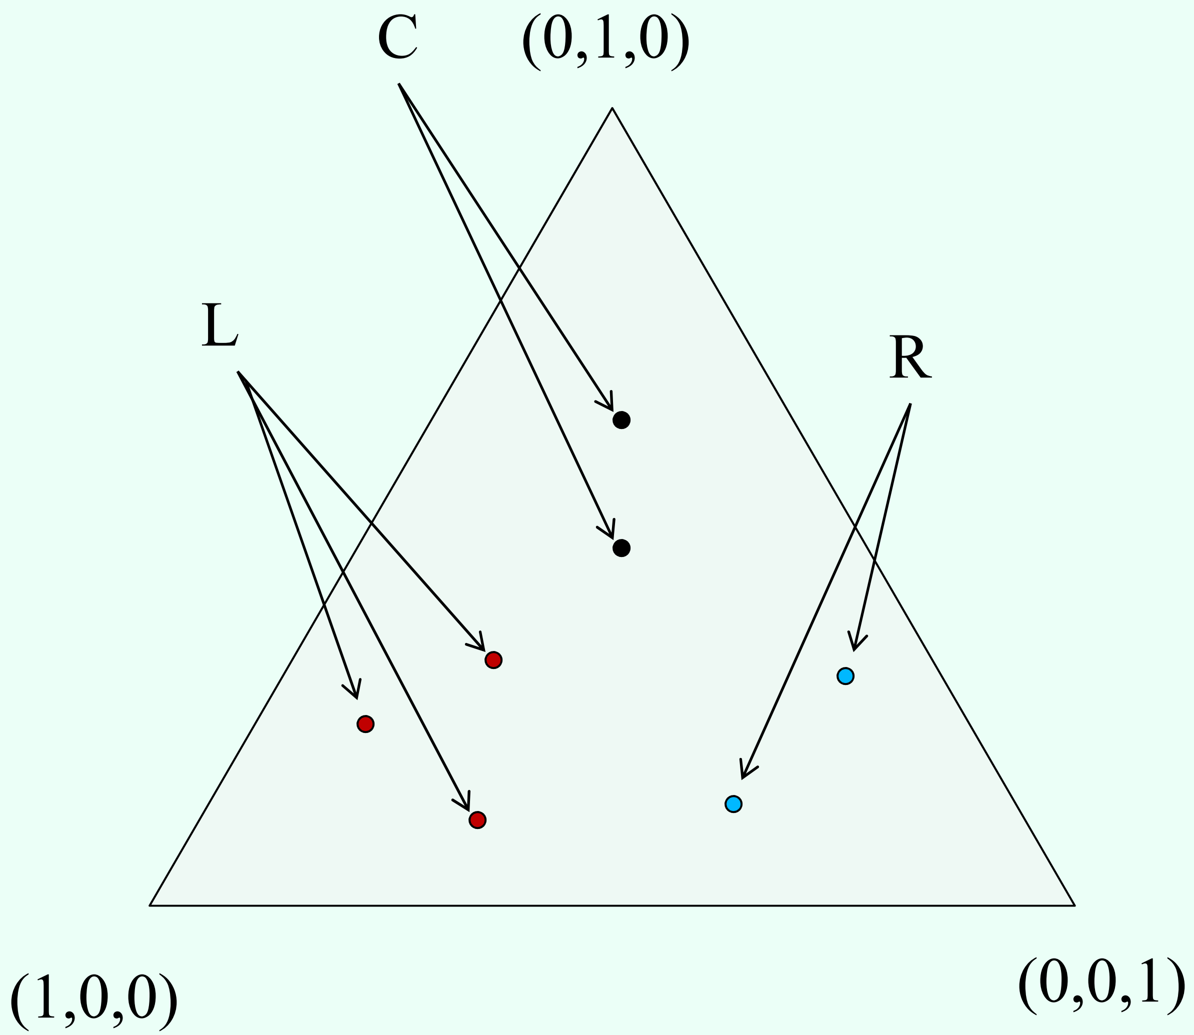

Three main techniques in the learning algorithm

- Find one point in each region (using random sampling)
- Find a point on an unknown hyperplane
- Starting from a point on an unknown hyperplane, determine the hyperplane completely

Finding a point on an unknown hyperplane

Step 1. Sample in the overlapping region

Step 2. Connect the new point to the point in the region that doesn't match

Step 3. Binary search along this line

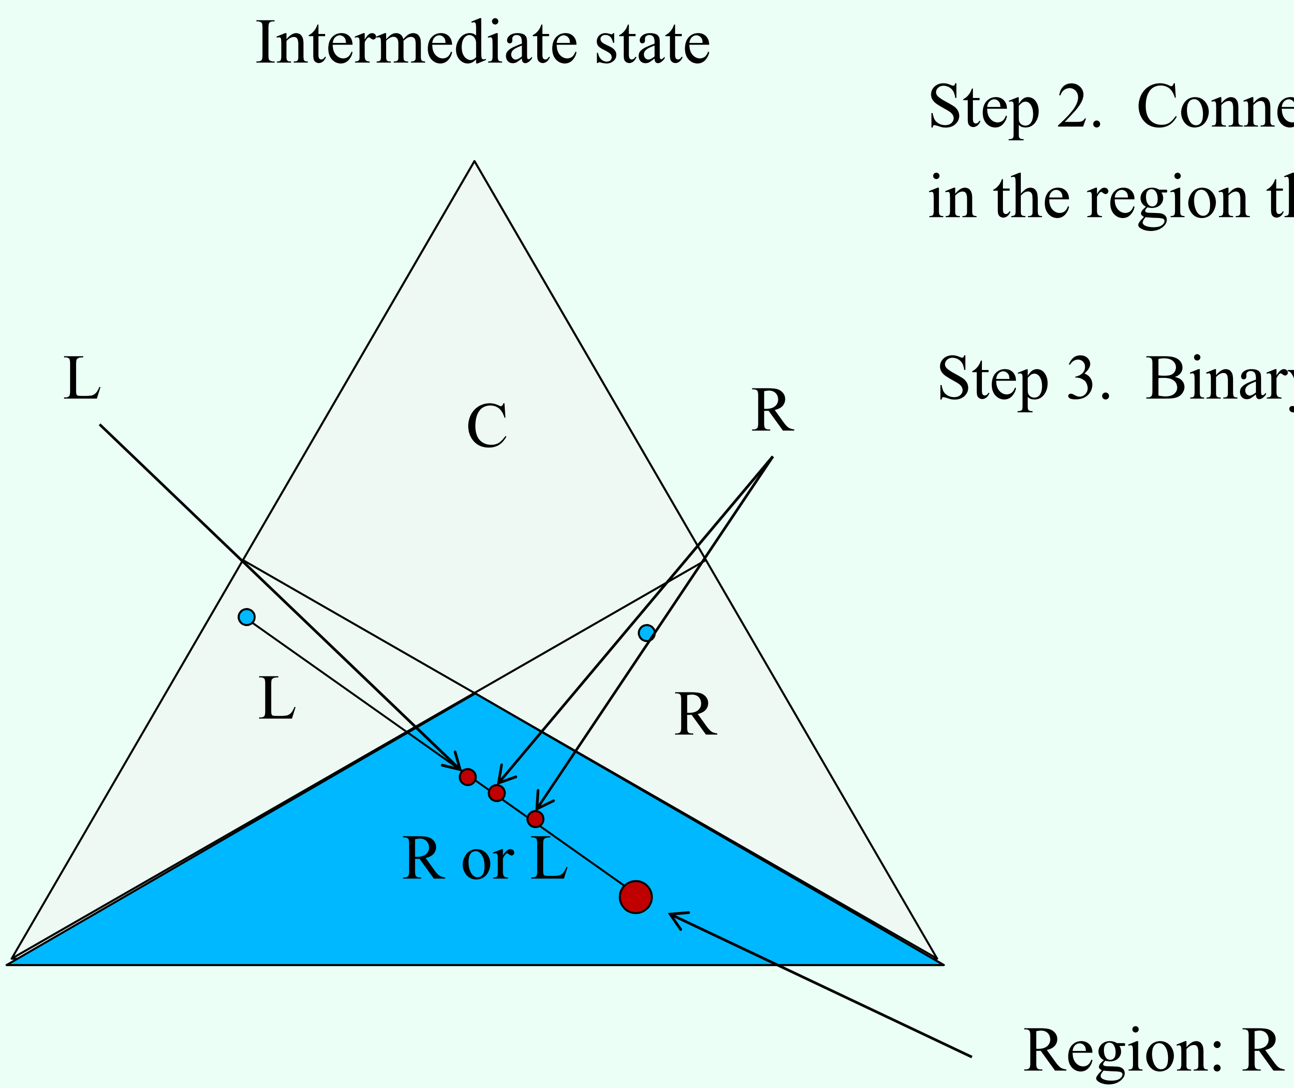

Determining the hyperplane

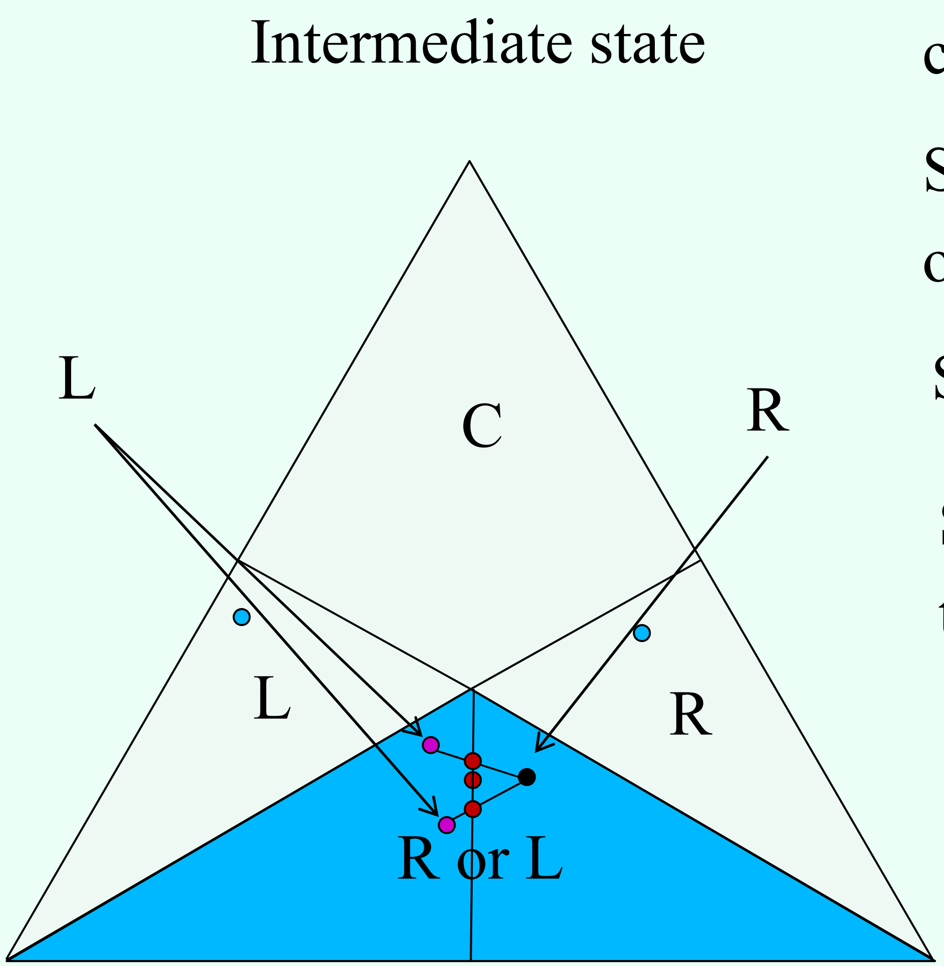

Step 1. Sample a regular d-simplex centered at the point

Step 2. Connect d lines between points on opposing sides

Step 3. Binary search along these lines

Step 4. Determine hyperplane (and update the region estimates with this information)

Bound on number of samples

Theorem. *Finding all of the hyperplanes necessary to compute the optimal mixed strategy to commit to requires* $O(Fk \log(k) + dLk^2)$ *samples*

- F depends on the size of the smallest region
- L depends on desired precision
- k is the number of follower actions
- d is the number of leader actions

Discussion about appropriateness of leadership model in security applications

- Mixed strategy not actually communicated
- Observability of mixed strategies?
  - Imperfect observation?
- Does it matter much (close to zero-sum anyway)?
- Modeling follower payoffs?
  - Sensitivity to modeling mistakes
- Human players… [Pita et al. 2009]

2, 1	4, 0
1, 0	3, 1

Computing optimal strategies to commit to in extensive-form games [Letchford & Conitzer 2010]

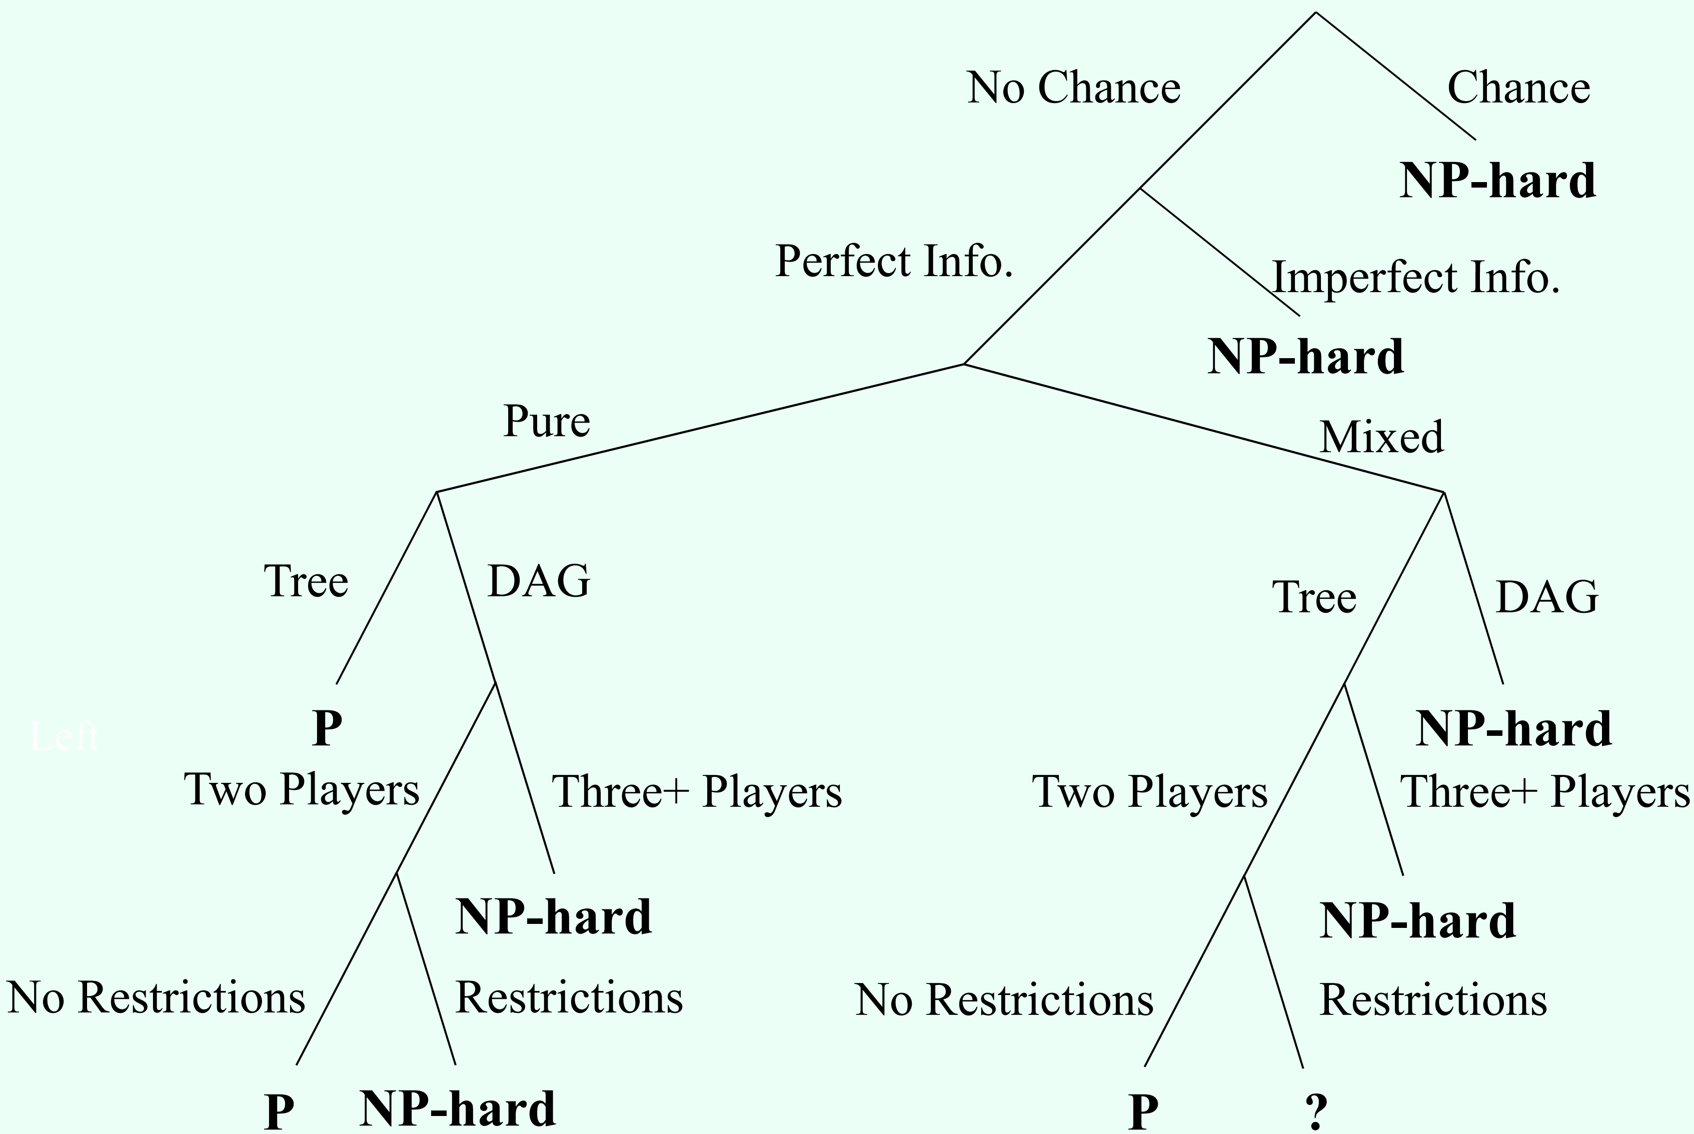

A problem for scaling to (some) real applications

- So far, we have assumed that we can enumerate all the defender pure strategies
- Not feasible in some applications
  - Federal Air Marshals [Tsai et al. 2009]
  - Protecting a city [Tsai et al. 2010]
  - …
- Problem: each possible allocation of resources is a pure strategy
  - Combinatorial explosion

Security resource allocation games

[Kiekintveld et al. 2009]

- Set of targets *T*
- Set of security resources Ω available to the defender (leader)
- Set of schedules $S \subseteq 2^T$
- Resource ω can be assigned to one of the schedules in $A(\omega) \subseteq S$
- Attacker (follower) chooses one target to attack
- Utilities: $U_d^c(t), U_a^c(t)$ if the attacked target is defended, $U_d^u(t), U_a^u(t)$ otherwise
- $U_d^c(t) \geq U_d^u(t); U_a^c(t) \leq U_a^u(t)$

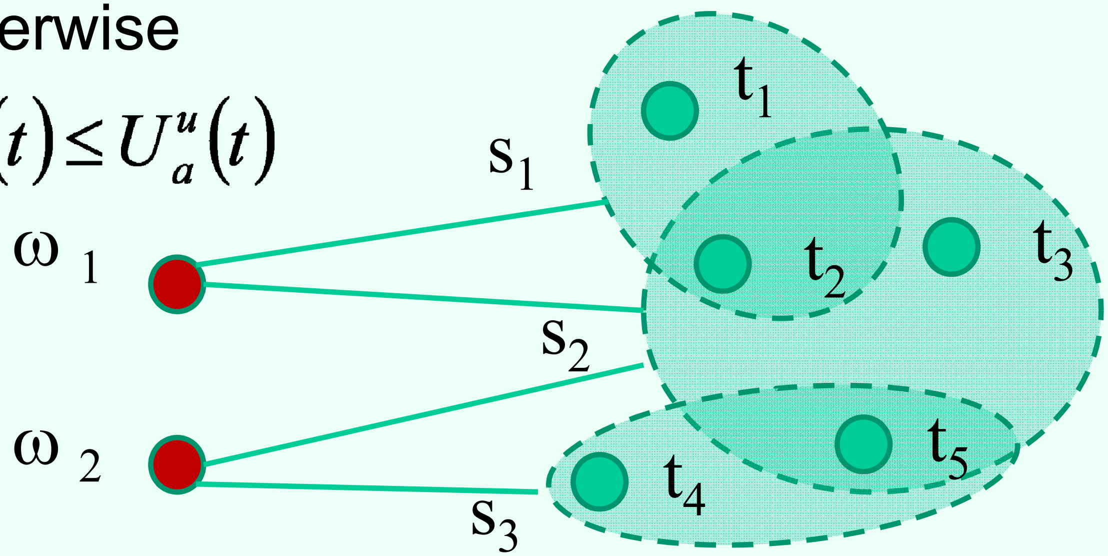

Applications and previous work

- Security checkpoints in airports (implemented at LAX) [Paruchuri et al. 2008, Pita et al. 2009]

- Federal air marshal service [Tsai et al. 2009]

Compact LPs approach

- Motivation: exponential number of pure strategies for the defender, so the standard LP is exponential in size

- Instead, we will find the (marginal) probability $c_{\omega,s}$ of resource ω being assigned to schedule s

Compact LP

- Cf. ERASER-C algorithm by Kiekintveld et al. [2009]
- Separate LP for every possible t* attacked:

$$\text{Maximize } U_d^c(t^*)\sum_{\omega}\sum_{s:t^*\in s} c_{\omega,s} + U_d^u(t^*)\left(1-\sum_{\omega}\sum_{s:t^*\in s} c_{\omega,s}\right)$$

Defender utility

Marginal probability of t^* being defended

Subject to

$$\forall \omega : \sum_{s} c_{\omega,s} \le 1$$

$$\forall t : \sum_{\omega}\sum_{s:t\in s} c_{\omega,s} \le 1$$

Distributional constraints

$$\forall t : U_a^c(t)\sum_{\omega}\sum_{s:t\in s} c_{\omega,s} + U_a^u(t)\left(1-\sum_{\omega}\sum_{s:t\in s} c_{\omega,s}\right)$$

$$\le U_a^c(t^*)\sum_{\omega}\sum_{s:t^*\in s} c_{\omega,s} + U_a^u(t^*)\left(1-\sum_{\omega}\sum_{s:t^*\in s} c_{\omega,s}\right)$$

Attacker optimality

Counter-example to the compact LP

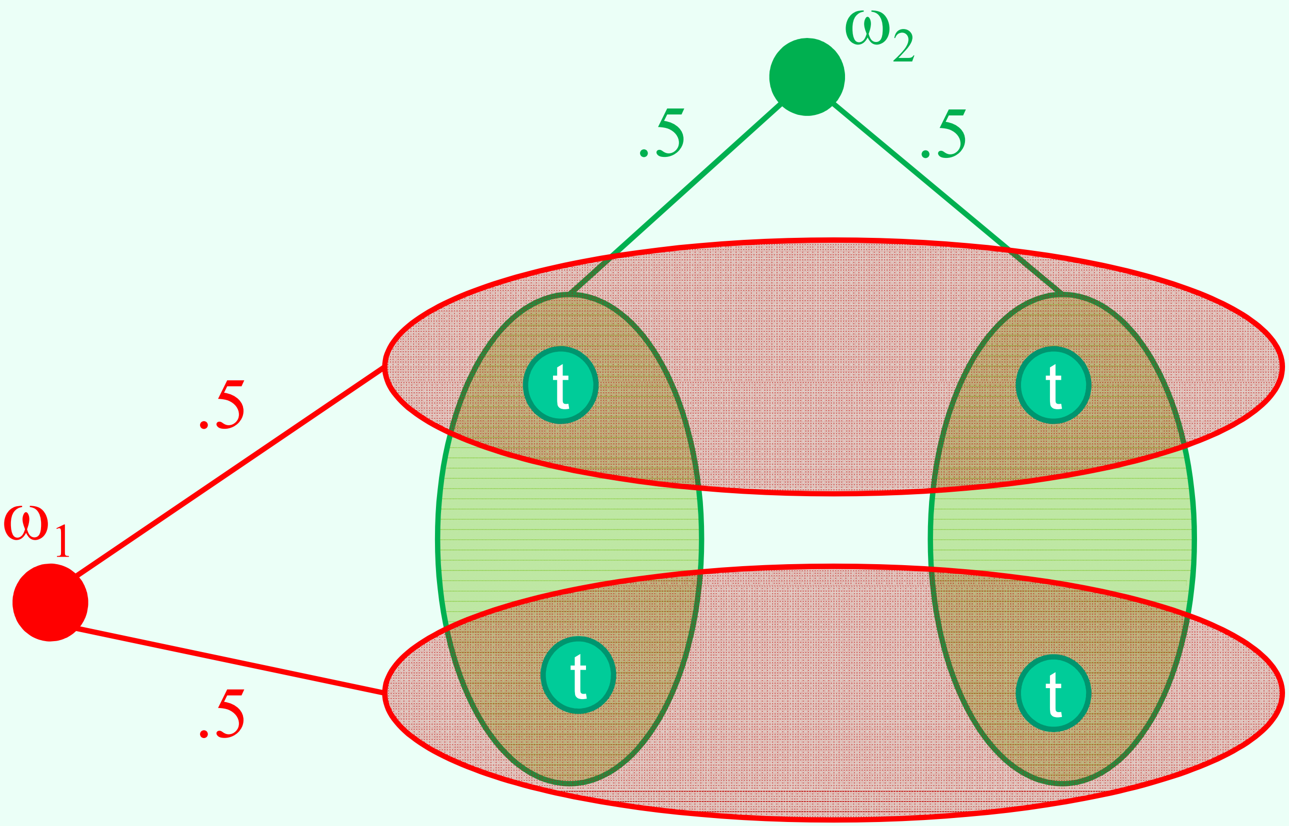

- LP suggests that we can cover every target with probability 1…
- … but in fact we can cover at most 3 targets at a time

Schedules of size 1

- Kiekintveld et al. prove that in this case, there exists a mixed strategy with the given marginal probabilities
- How can we find it?

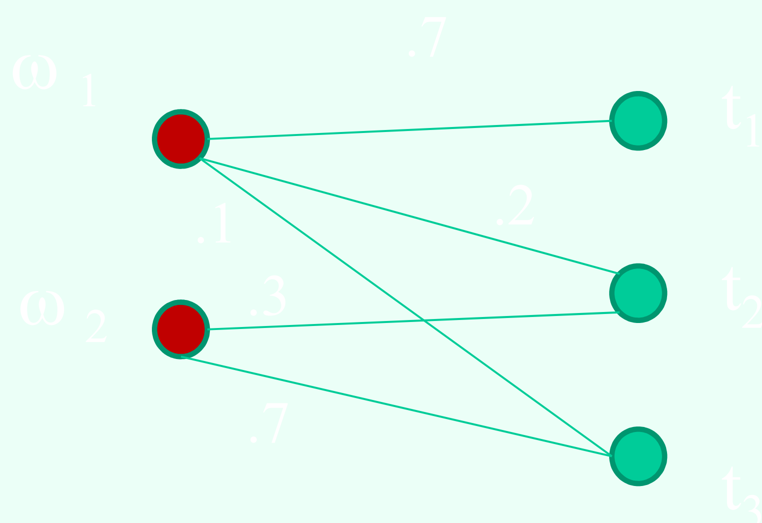

Birkhoff-von Neumann theorem

- Every *doubly stochastic n x n* matrix can be represented as a convex combination of *n x n* permutation matrices

.1	.4	.5
.3	.5	.2
.6	.1	.3

$$= .1 \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix} + .1 \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} + .5 \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} + .3 \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

- Decomposition can be found in polynomial time $O(n^{4.5})$, and the size is $O(n^2)$ [Dulmage and Halperin, 1955]
- Can be extended to *rectangular* doubly *substochastic* matrices

Computing the probabilities for each pure strategy

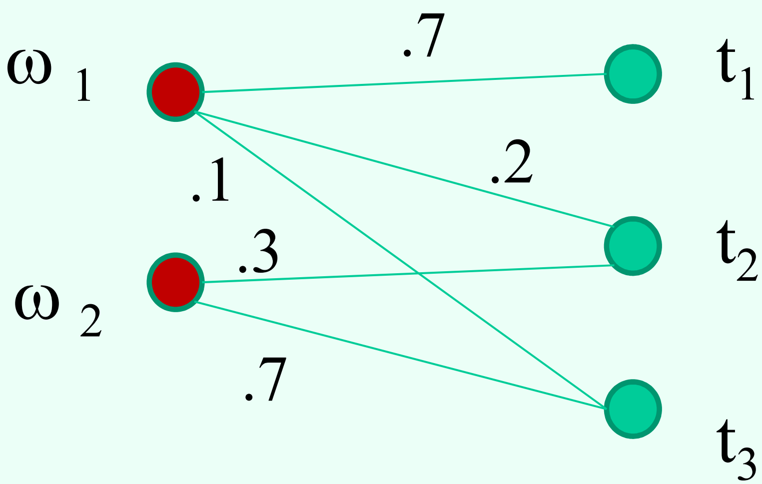

	t_1	t_2	t_3
ω_1	.7	.2	.1
ω_2	0	.3	.7

.1

0	0	1
0	1	0

.2

0	1	0
0	0	1

.2

1	0	0
0	1	0

.5

1	0	0
0	0	1

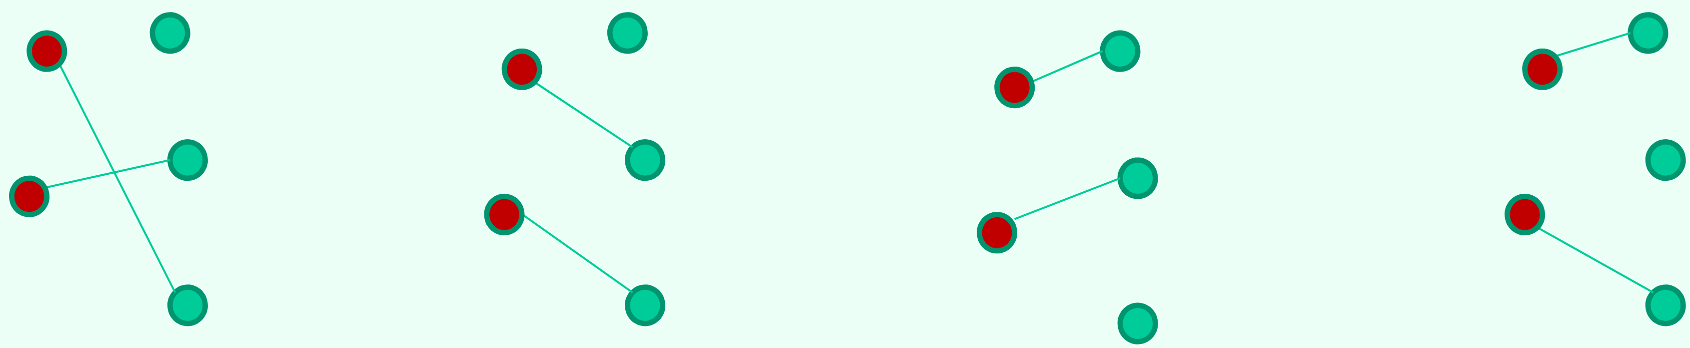

Summary of results

[Korzhyk, Conitzer, Parr 2010]

		Homogeneous Resources	Heterogeneous resources
Schedules	**Size 1**	P	P (BvN theorem)
	Size ≤2, bipartite	P (BvN theorem)	NP-hard (SAT)
	Size ≤2	P (constraint generation)	NP-hard
	Size ≥3	NP-hard (3-COVER)	NP-hard

Is it right to play Stackelberg?

- Typical argument: *attacker can observe realizations of our distribution over time before executing an attack, learn the distribution*
- Is this accurate?
- We show that under certain conditions, it is "safe" to play the Stackelberg strategy [Yin et al. 2010]

Every Stackelberg strategy is also a Nash strategy in security games

- *Theorem:* If any subset of any schedule is also a schedule, then every Stackelberg strategy is also part of a Nash equilibrium

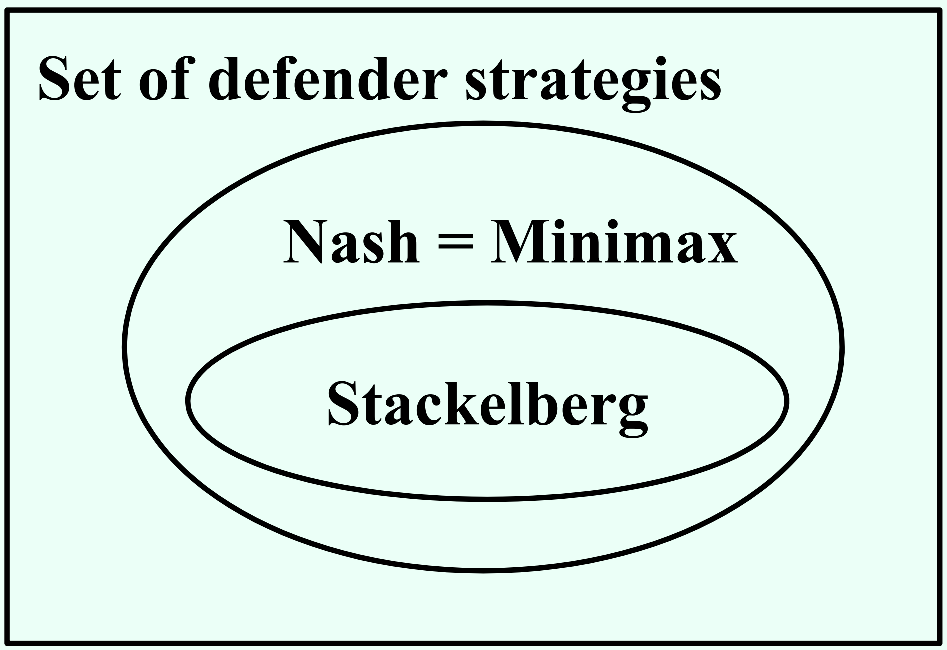

So how do we know we're playing the "right" equilibrium?

- Turns out not to matter:
- *Theorem.* Security games satisfy the interchange property:

  if $<c_1,a_1>$ and $<c_2,a_2>$ are NE profiles, then $<c_1,a_2>$ and $<c_2,a_1>$ are also NE profiles

  – Doesn't hold in general games (e.g., chicken)
- Proof analyzes a related zero-sum game
  – Two-player zero-sum games always have the interchange property

Interchange property in security games

- There is a 1:1 equivalence between NE profiles in general-sum and zero-sum games.
- *Interchange* property of NE in zero-sum games: if $<c_1,a_1>$ and $<c_2,a_2>$ are NE profiles, then $<c_1,a_2>$ and $<c_2,a_1>$ are also NE profiles. This property doesn't hold in general games.
- Interchange property carries over to general-sum security games because of the above equivalence.

Consequence

- When the defender is uncertain whether her strategy is known to the attacker or not, it is safe to play an SSE strategy.
- If the attacker somehow learns the defender’s strategy, the defender gets optimal utility.
- If the attacker does not learn the defender’s strategy, the SSE strategy is as good as any other NE strategy because of the interchange property.

Conclusion

- Desire to address general-sum games in security
- Optimal mixed strategies to commit to ("Stackelberg strategies") have certain conceptual & algorithmic advantages over (say) Nash equilibrium
- Computational challenges remain: Many games have exponential strategy spaces
- Also raises & forces close examination of fundamental game-theoretic questions

Thank you for your attention!

Rock-paper-scissors – Seinfeld variant

 

 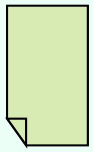

MICKEY: All right, rock beats paper!
(Mickey smacks Kramer's hand for losing)
KRAMER: I thought paper covered rock.
MICKEY: Nah, rock flies right through paper.
KRAMER: What beats rock?
MICKEY: (looks at hand) Nothing beats rock.

		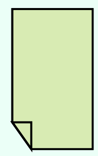	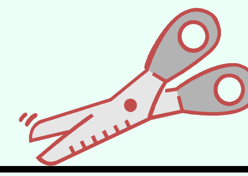
	0, 0	1, -1	1, -1
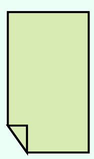	-1, 1	0, 0	-1, 1
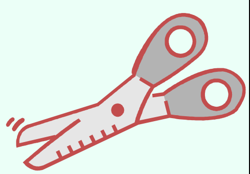	-1, 1	1, -1	0, 0

Dominance

- Player i's strategy s_i **strictly dominates** s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$
- s_i **weakly dominates** s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$; and
  - for some s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

-i = "the player(s) other than i"

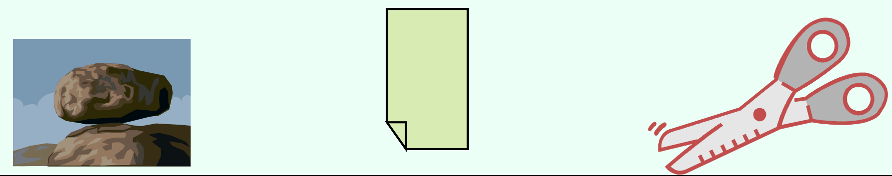

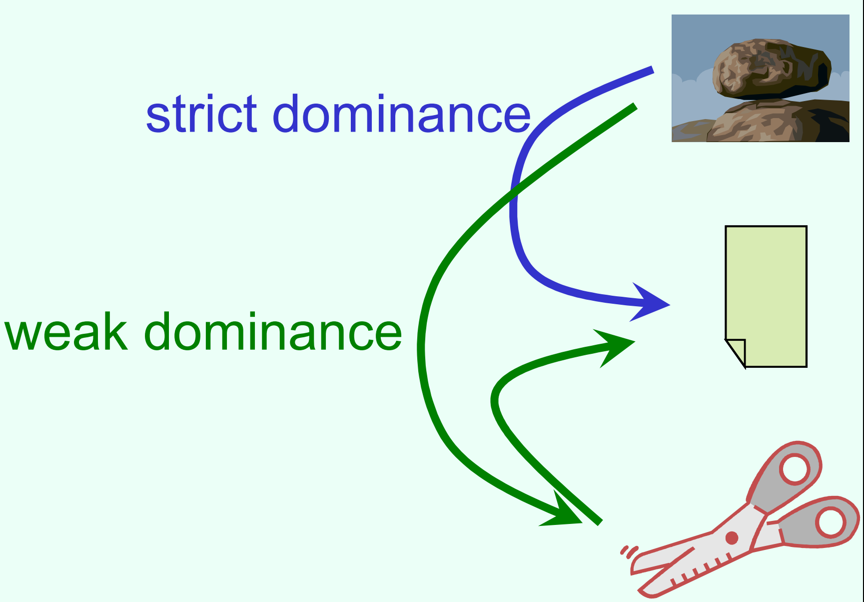

0, 0	1, -1	1, -1
-1, 1	0, 0	-1, 1
-1, 1	1, -1	0, 0

Prisoner's Dilemma

- Pair of criminals has been caught
- District attorney has evidence to convict them of a minor crime (1 year in jail); knows that they committed a major crime together (3 years in jail) but cannot prove it
- Offers them a deal:
  - If both confess to the major crime, they each get a 1 year reduction
  - If only one confesses, that one gets 3 years reduction

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

"Should I buy an SUV?"

purchasing + gas cost

cost: 5

cost: 3

accident cost

cost: 5  cost: 5

cost: 8 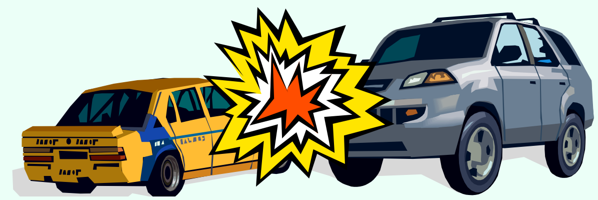 cost: 2

cost: 5  cost: 5

	-10, -10	-7, -11
	-11, -7	-8, -8

“2/3 of the average” game

- Everyone writes down a number between 0 and 100
- Person closest to 2/3 of the average wins
- Example:
  - A says 50
  - B says 10
  - C says 90
  - Average(50, 10, 90) = 50
  - 2/3 of average = 33.33
  - A is closest (|50-33.33| = 16.67), so A wins

Iterated dominance

- Iterated dominance: remove (strictly/weakly) dominated strategy, repeat
- Iterated strict dominance on Seinfeld's RPS:

0, 0	1, -1	1, -1
-1, 1	0, 0	-1, 1
-1, 1	1, -1	0, 0

0, 0	1, -1
-1, 1	0, 0

Iterated dominance: path (in)dependence

Iterated weak dominance is path-dependent: sequence of eliminations may determine which solution we get (if any)
(whether or not dominance by mixed strategies allowed)

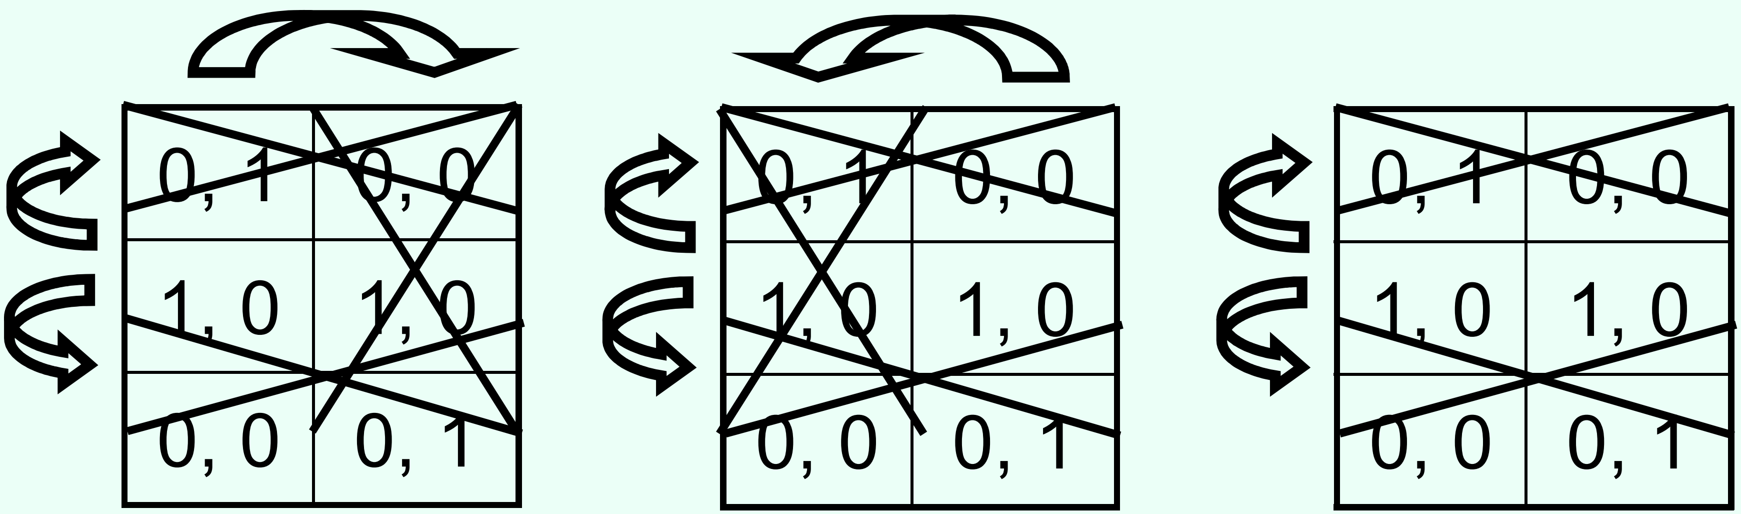

Iterated strict dominance is path-independent: elimination process will always terminate at the same point
(whether or not dominance by mixed strategies allowed)

“2/3 of the average” game revisited

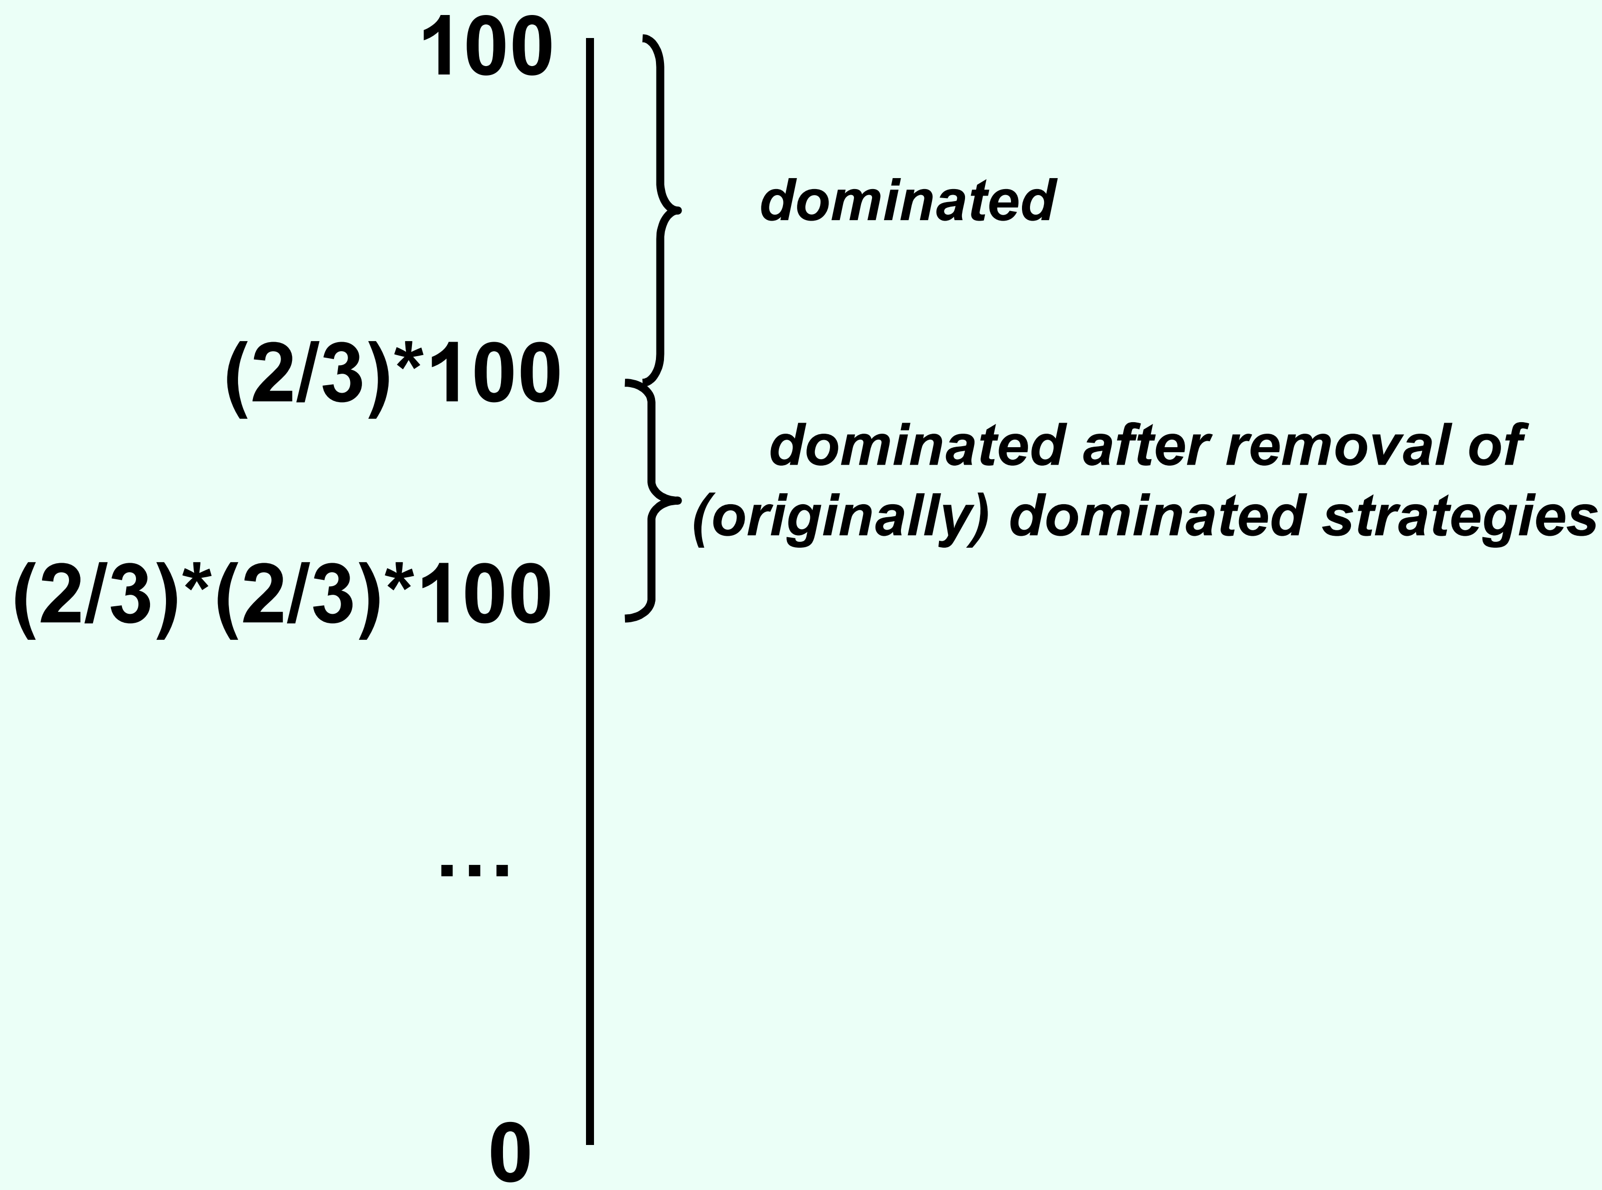

Mixed strategies

- Mixed strategy for player i = probability distribution over player i's (pure) strategies
- E.g. 1/3 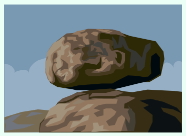 , 1/3 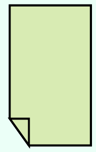 , 1/3 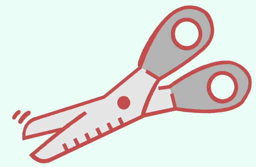
- Example of dominance by a mixed strategy:

1/2	3, 0	0, 0
1/2	0, 0	3, 0
	1, 0	1, 0

Checking for dominance by mixed strategies

- Linear program for checking whether strategy s_i^* is strictly dominated by a mixed strategy:
- maximize ε
- such that:
  - for any s_{-i}, $\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i}) + \varepsilon$
  - $\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} = 1$

- Linear program for checking whether strategy s_i^* is weakly dominated by a mixed strategy:
- maximize $\Sigma_{s_{-i}}(\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} u_i(s_i, s_{-i})) - u_i(s_i^*, s_{-i})$
- such that:
  - for any s_{-i}, $\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i})$
  - $\Sigma_{s_i} \mathbf{p}_{\mathbf{s}_i} = 1$

The presentation game

		Presenter	
		Put effort into presentation (E)	*Do not put effort into presentation (NE)*
Audience	*Pay attention (A)*	4, 4	-16, -14
	Do not pay attention (NA)	0, -2	0, 0

- Pure-strategy Nash equilibria: (A, E), (NA, NE)
- Mixed-strategy Nash equilibrium:

  ((1/10 A, 9/10 NA), (4/5 E, 1/5 NE))

  - Utility 0 for audience, -14/10 for presenter
  - Can see that some equilibria are strictly better for both players than other equilibria, i.e. some equilibria Pareto-dominate other equilibria

A poker-like game

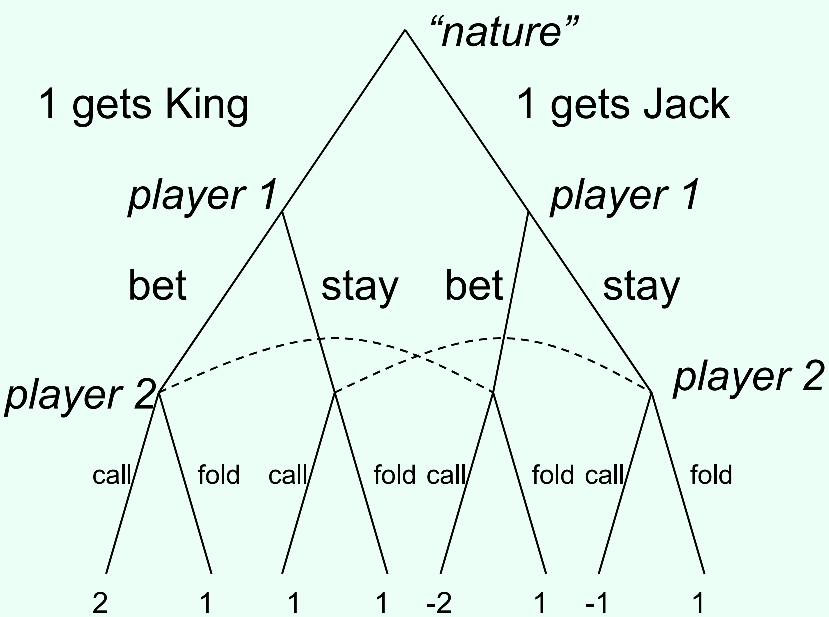

	cc	cf	fc	ff
bb	0, 0	0, 0	1, -1	1, -1
bs	.5, -.5	1.5, -1.5	0, 0	1, -1
sb	-.5, .5	-.5, .5	1, -1	1, -1
ss	0, 0	1, -1	0, 0	1, -1

A poker-like game

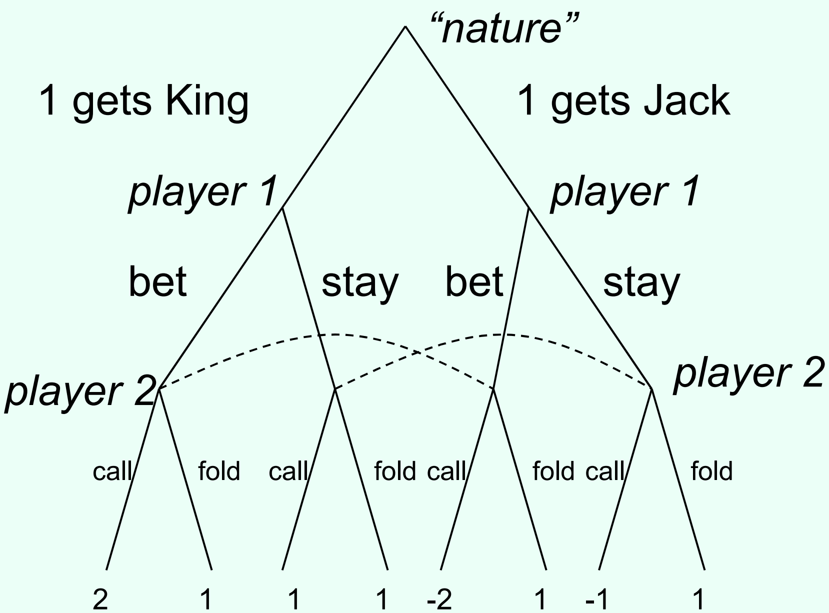

		2/3 cc	cf	1/3 fc	ff
1/3	bb	0, 0	0, 0	1, -1	1, -1
2/3	bs	.5, -.5	1.5, -1.5	0, 0	1, -1
	sb	-.5, .5	-.5, .5	1, -1	1, -1
	ss	0, 0	1, -1	0, 0	1, -1

- To make player 1 indifferent between bb and bs, we need:

  utility for bb = 0*P(cc)+1*(1-P(cc)) = .5*P(cc)+0*(1-P(cc)) = utility for bs

  That is, P(cc) = 2/3
- To make player 2 indifferent between cc and fc, we need:

  utility for cc = 0*P(bb)+(-.5)*(1-P(bb)) = -1*P(bb)+0*(1-P(bb)) = utility for fc

  That is, P(bb) = 1/3

Rock-paper-scissors

		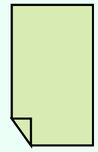	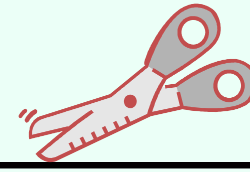
	0, 0	-1, 1	1, -1
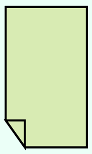	1, -1	0, 0	-1, 1
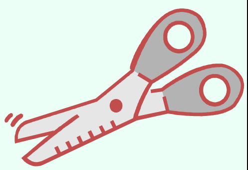	-1, 1	1, -1	0, 0

- Any pure-strategy Nash equilibria?
- But it has a mixed-strategy Nash equilibrium:

  Both players put probability 1/3 on each action
- If the other player does this, every action will give you expected utility 0
  - Might as well randomize

Nash equilibria of “chicken”

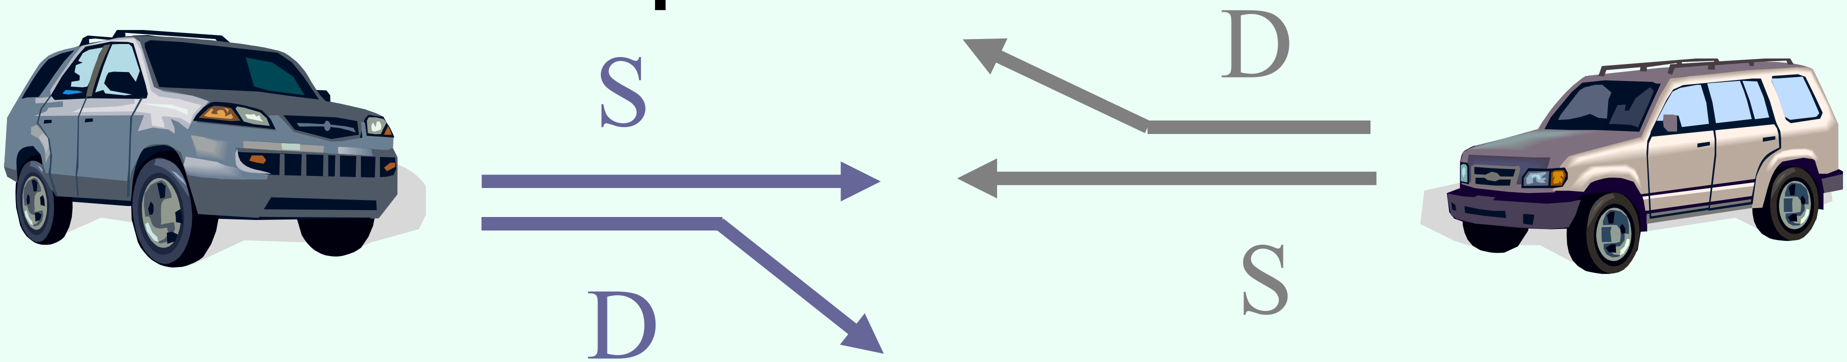

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- (D, S) and (S, D) are Nash equilibria
  - They are pure-strategy Nash equilibria: nobody randomizes
  - They are also strict Nash equilibria: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

Nash equilibria of “chicken”…

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 uses a mixed strategy?
- If a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- So we need to make player 1 indifferent between D and S
- Player 1’s utility for playing D = $-p^c_S$
- Player 1’s utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
  - People may die! Expected utility -1/5 for each player

Ranges for the follower payoffs

- Suppose we just know a range within which each follower payoff lies

	L	R
U	1, [0,3]	2, 3
D	0, [1,3]	1, [1,2]

- NP-hard if payoffs are adversarially drawn
  - We do not know about (in)approximability…
  - … except for a richer variant

Extension of the BvN theorem

- Every *m x n* doubly *substochastic* matrix can be represented as a convex combination of *m x n* matrices with elements from {0, 1} such that every row and column contains "1" in at most one cell.

$$\begin{pmatrix} .1 & .4 & .5 \\ .3 & .5 & .2 \\ .6 & .1 & .3 \end{pmatrix}$$

$$= .1 \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} + .1 \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} + .5 \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} + .3 \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

[backup] Will compact LP work for homogeneous resources?

- Suppose that every resource can be assigned to any schedule.
- We can still find a counter-example for this case:

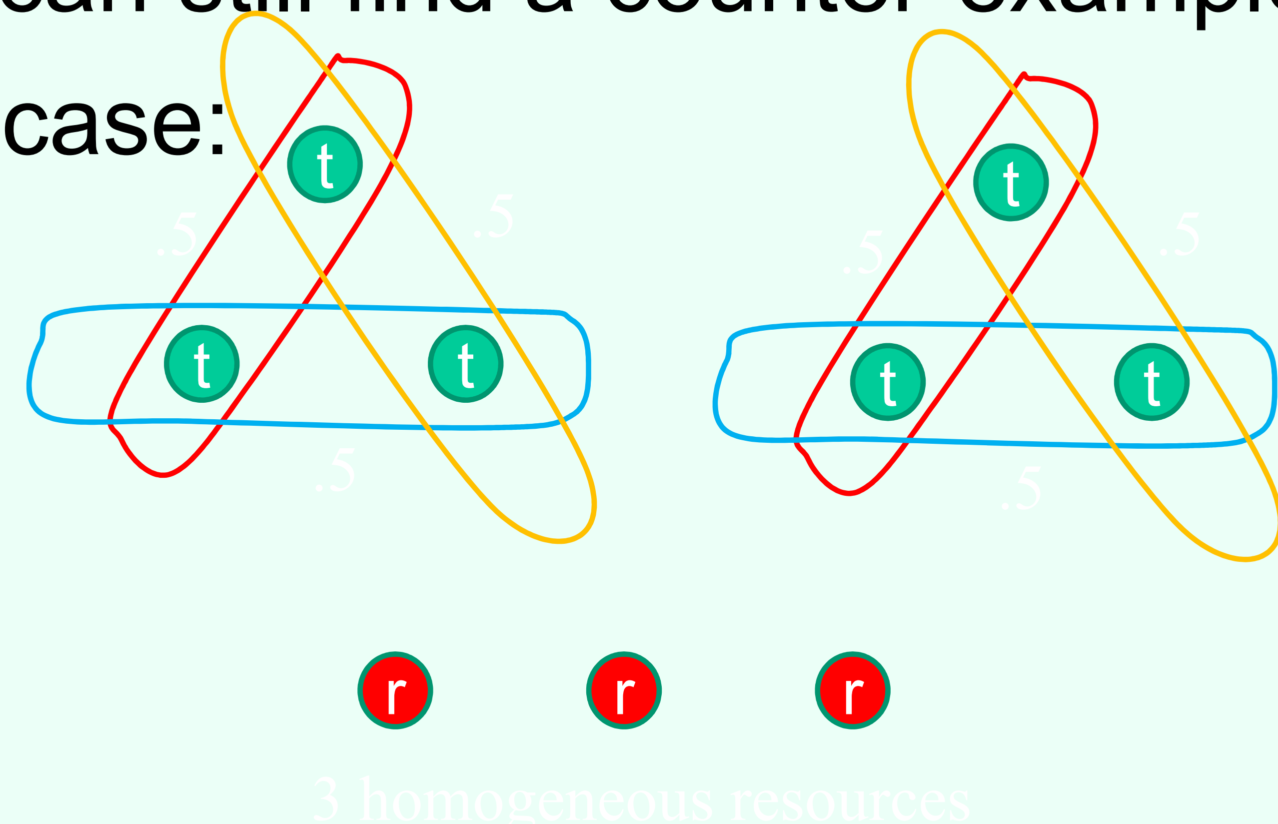

Stackelberg games in extensive form