


Computing Game-Theoretic Solutions

Vincent Conitzer
Duke University

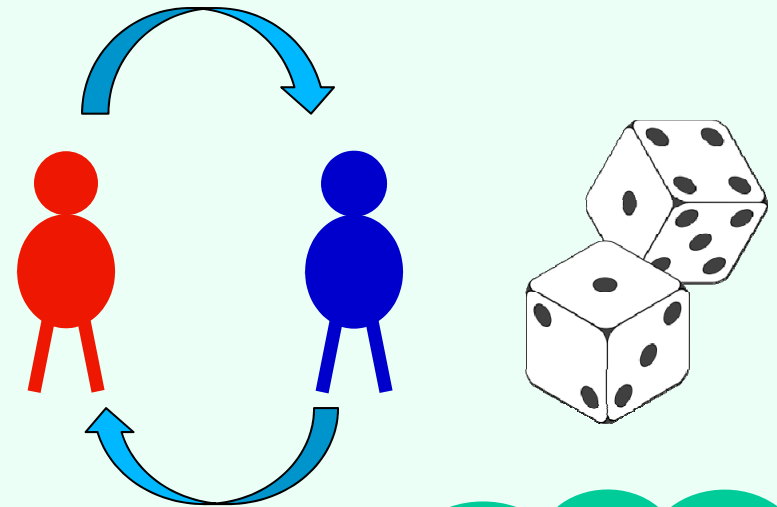
overview article: V. Conitzer. Computing Game-Theoretic Solutions and Applications to Security. Proc. AAAI'12.

Game theory


Multiple **self-interested** agents interacting in the same environment



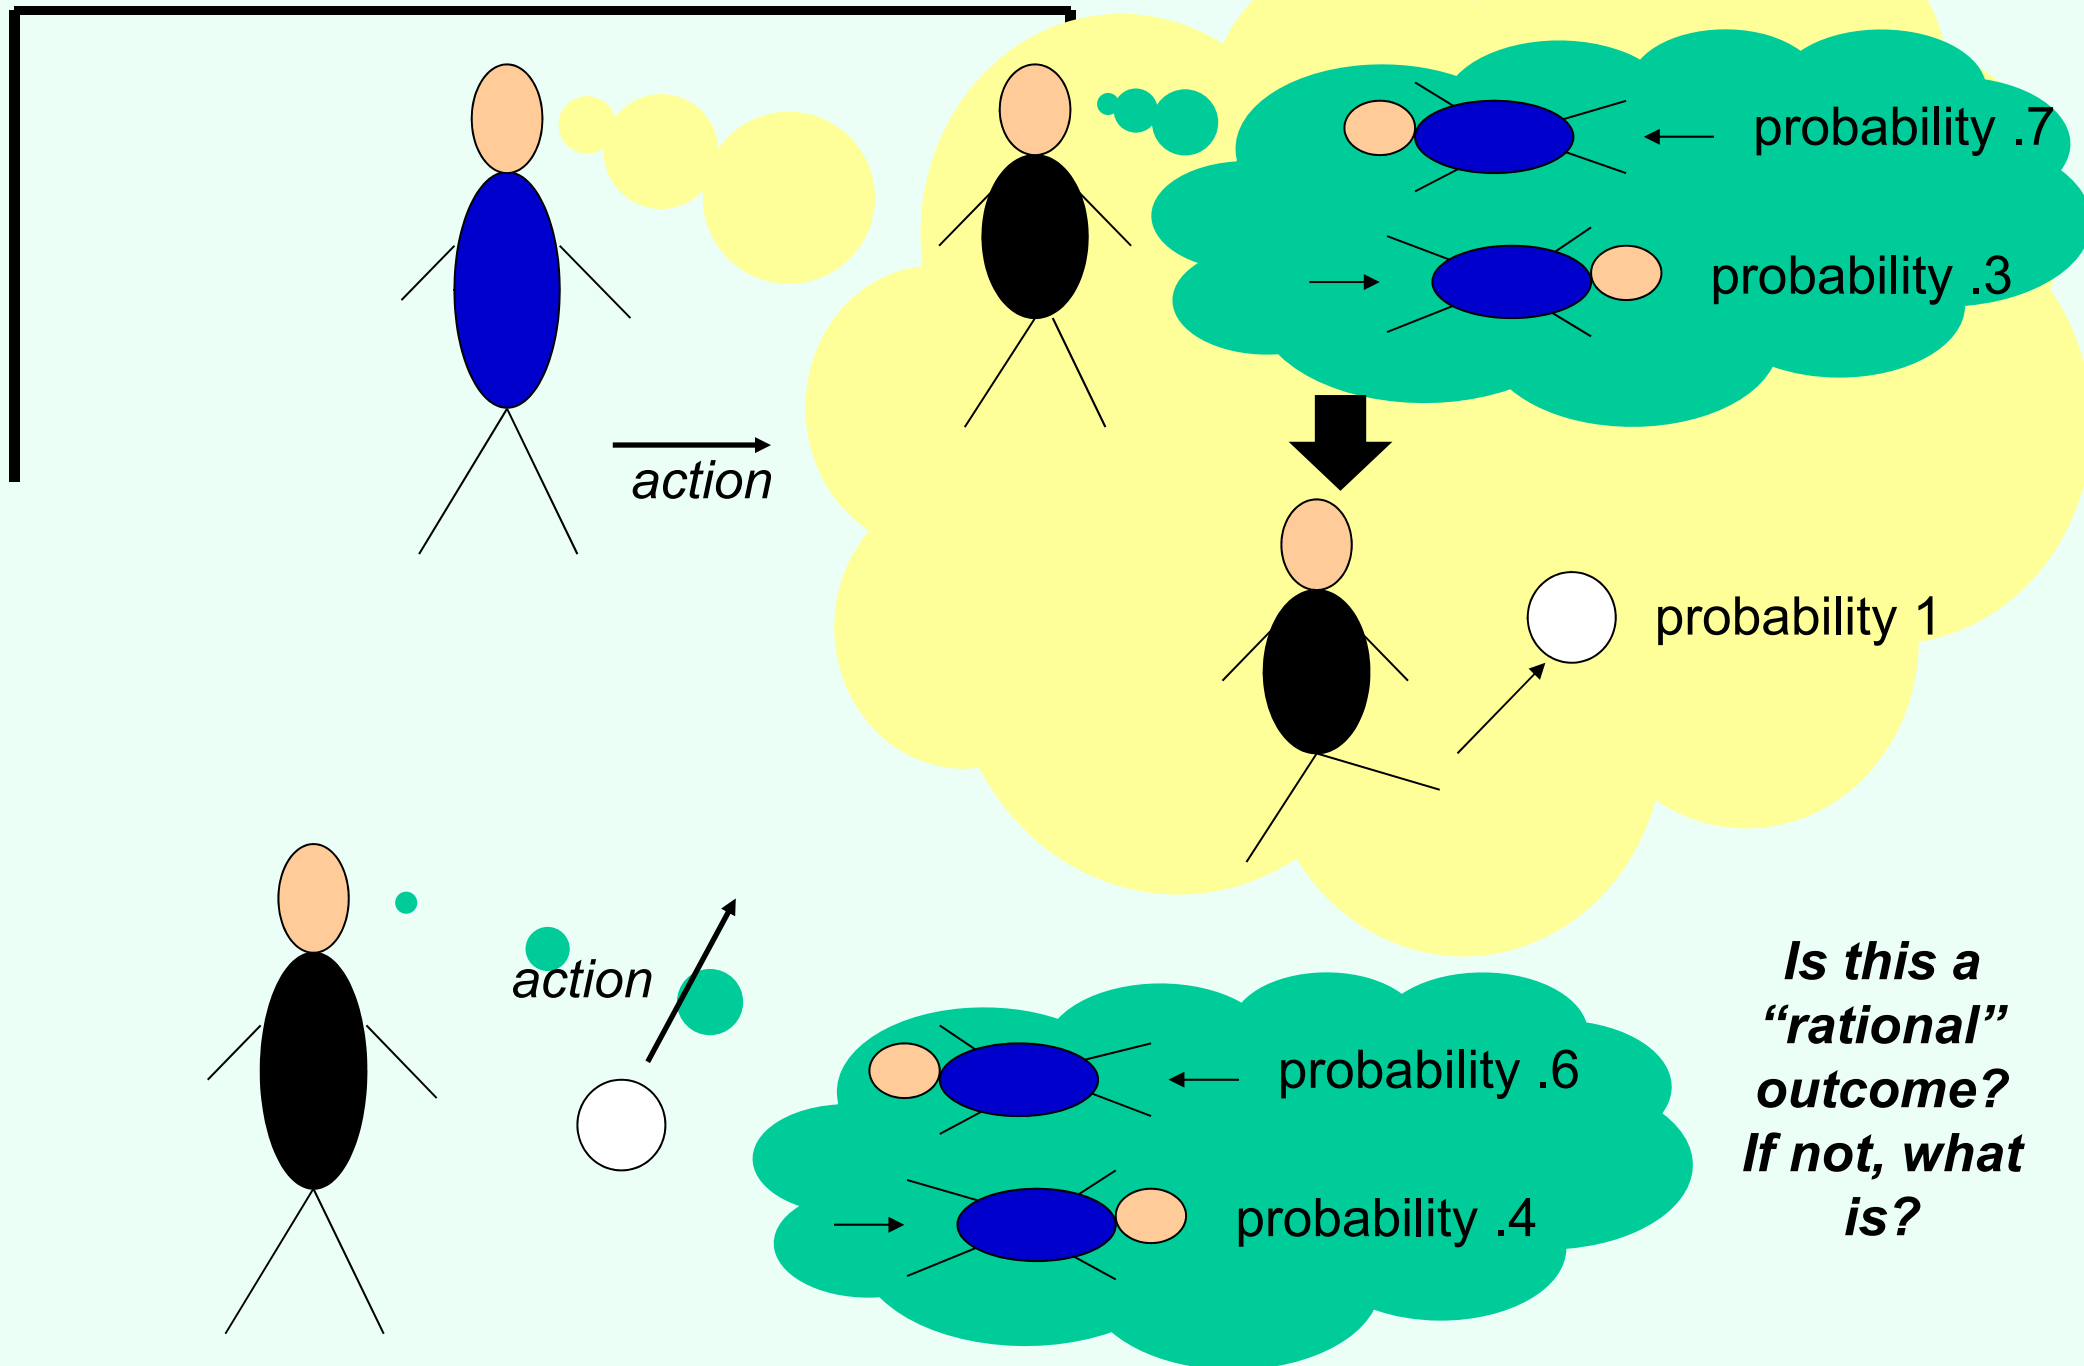
What is an agent to **do**?



What is an agent to **believe**? (What are we to believe?)



Penalty kick example



Multiagent systems

Goal:
Blocked (Room0)




Goal:
Clean (Room0)



Game playing



Real-world security applications



Airport security *Milind Tambe's TEAMCORE group (USC)*

- Where should checkpoints, canine units, etc. be deployed?
- Deployed at LAX airport and elsewhere

Federal Air Marshals

- Which flights get a FAM?

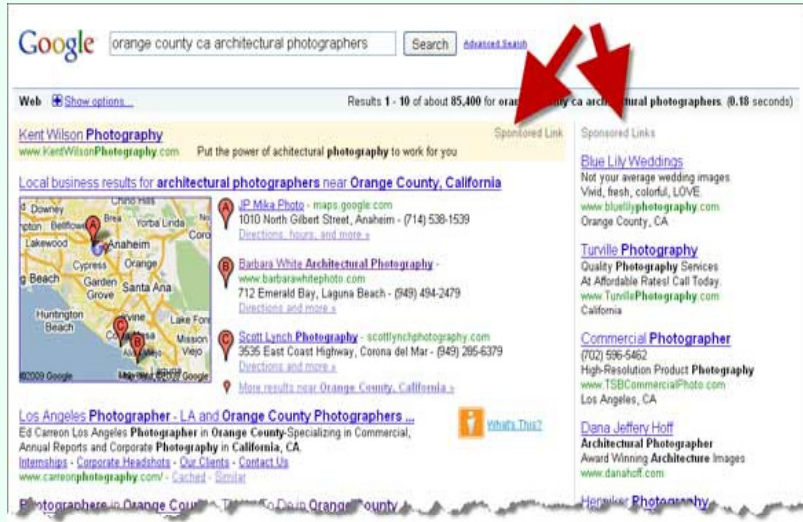


US Coast Guard

- Which patrol routes should be followed?
- Deployed in Boston, New York, Los Angeles




Mechanism design



- 1. Rating: (14 votes cast)
- 2. Rating: (15 votes cast) Thanks for voting!
- 3. Rating: (15 votes cast)
- 4. Rating: (12 votes cast)

Rating/voting systems

Auctions



Kidney exchanges



Prediction markets



Donation matching

Outline

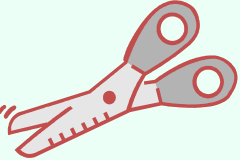
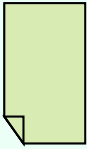

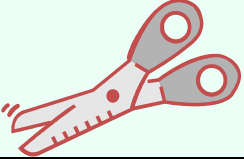
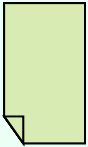

- **Introduction to game theory (from CS perspective)**
 - Representing games
 - Standard solution concepts
 - (Iterated) dominance
 - Minimax strategies
 - Nash and correlated equilibrium
- **Recent developments**
 - Commitment: Stackelberg mixed strategies
 - Security applications
- **Learning in games (time permitting)**
 - Simple algorithms
 - Evolutionary game theory
 - Learning in Stackelberg games

Representing games

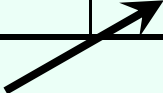
Rock-paper-scissors

Column player aka.
player 2
(simultaneously)
chooses a column

Row player
aka. player
1 chooses a
row



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





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.

Penalty kick

(also known as: matching pennies)



.5
L

.5
R

.5 L

0, 0

-1, 1



.5 R

-1, 1

0, 0



Security example

Terminal A





→
action

Terminal B



↗
action


Security game



	A	B
A	0, 0	-1, 2
B	-1, 1	0, 0

“Chicken”

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




Modeling and representing games

THIS TALK
(unless
specified
otherwise)

2, 2	-1, 0
-7, -8	0, 0

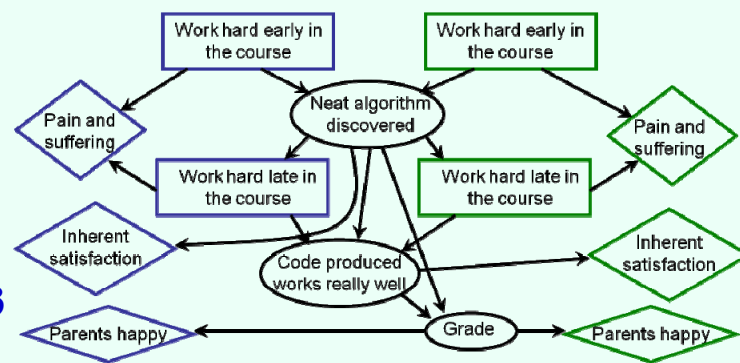



normal-form games



		L	R		L	R	
row player	U	4	6	column player	U	4	6
type 1 (prob. 0.5)	D	2	4	type 1 (prob. 0.5)	D	4	6

		L	R		L	R	
row player	U	2	4	column player	U	2	2
type 2 (prob. 0.5)	D	4	2	type 2 (prob. 0.5)	D	4	2

Bayesian games

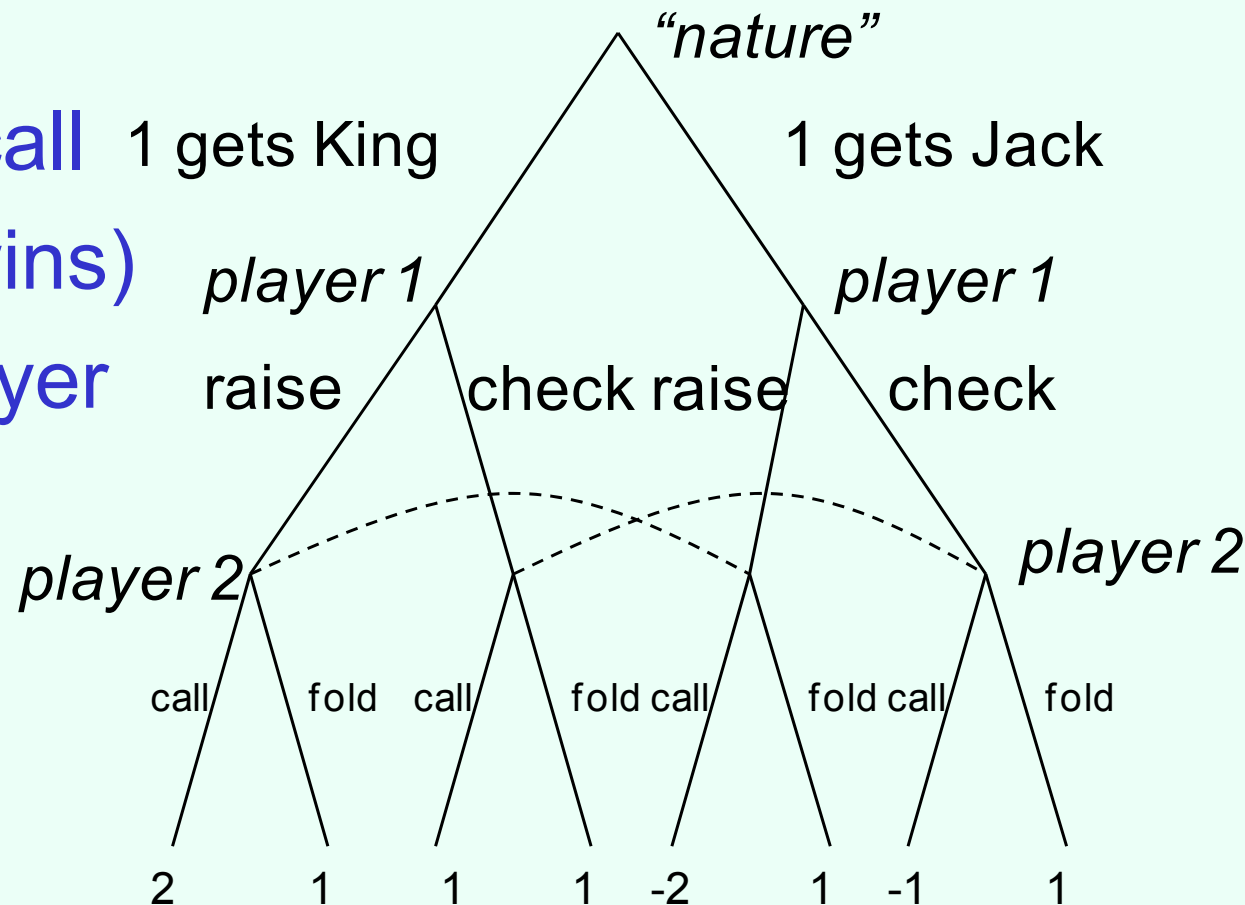


[Kearns, Littman, Singh UAI'01] [Leyton-Brown & Tennenholtz IJCAI'03]
[Bhat & Leyton-Brown, UAI'04]
[Jiang, Leyton-Brown, Bhat GEB'11]

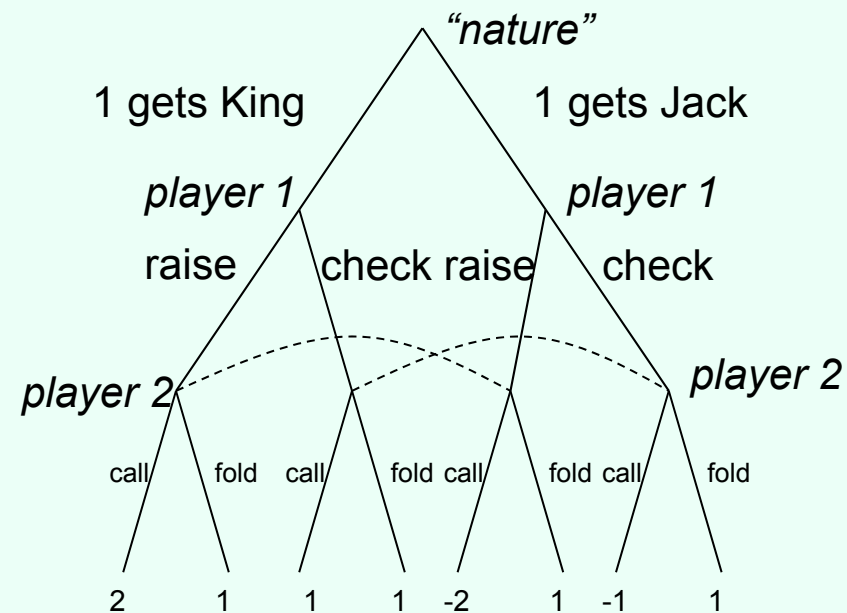
[Koller & Milch. IJCAI'01/GEB'03]

A poker-like game

- Both players put 1 chip in the pot
- Player 1 gets a card (King is a winning card, Jack a losing card)
- Player 1 decides to raise (add one to the pot) or check
- Player 2 decides to call (match) or fold (P1 wins)
- If player 2 called, player 1's card determines pot winner




Poker-like game in normal form





	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1






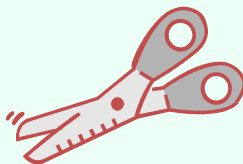
**Our first solution concept:
Dominance**

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"*



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 (additional 2 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, maintenance) cost

accident cost



cost: 5

cost: 5



cost: 5



cost: 3

cost: 8



cost: 2

cost: 5




cost: 5




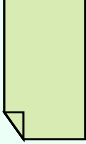
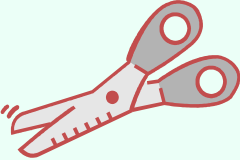
-10, -10	-7, -11
-11, -7	-8, -8

Back to the poker-like game



	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1

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

A blue bracket on the left side of the table groups the first two rows, with a blue arrow pointing from the bracket to the third row.


Usage:
 σ_i denotes a mixed strategy,
 s_i denotes a pure strategy

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

Iterated dominance


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



“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
 - $\text{Average}(50, 10, 90) = 50$
 - $2/3$ of average = 33.33
 - A is closest ($|50 - 33.33| = 16.67$), so A wins

“2/3 of the average” game solved




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)

Leads to various NP-hardness results [Gilboa, Kalai, Zemel Math of
OR '93; C. & Sandholm EC '05, AAI'05; Brandt, Brill, Fischer, Harrenstein TOCS '11]



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



Solving two-player zero-sum games

How to play matching pennies

		<i>Them</i>	
		L	R
<i>Us</i>	L	1, -1	-1, 1
	R	-1, 1	1, -1

- Assume opponent **knows our mixed strategy**
- 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?

A locally popular sport



go for 3 go for 2

defend the 3

defend the 2

0, 0	-2, 2
-3, 3	0, 0

Solving basketball

		<i>Them</i>	
		3	2
<i>Us</i>	3	0, 0	-2, 2
	2	-3, 3	0, 0

- If we 50% of the time defend the 3, opponent will shoot 3
 - We get $.5*(-3) + .5*(0) = -1.5$
- Should defend the 3 more often: 60% of the time
- Opponent has choice between
 - Go for 3: gives them $.6*(0) + .4*(3) = 1.2$
 - Go for 2: gives them $.6*(2) + .4*(0) = 1.2$
- We get -1.2 (the **maximin** value)

Let's change roles

		<i>Them</i>	
		3	2
<i>Us</i>	3	0, 0	-2, 2
	2	-3, 3	0, 0

- Suppose **we** know **their** strategy
 - If 50% of the time they go for 3, then we defend 3
 - We get $.5*(0)+.5*(-2) = -1$
 - Optimal for them: 40% of the time go for 3
 - If we defend 3, we get $.4*(0)+.6*(-2) = -1.2$
 - If we defend 2, we get $.4*(-3)+.6*(0) = -1.2$
 - This is the **minimax** value
- von Neumann's minimax theorem [1928]: maximin value = minimax value
(~ linear programming duality)

Minimax theorem [von Neumann 1928]

- Maximin utility: $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i})$

$$(\text{= - } \min_{\sigma_i} \max_{s_{-i}} u_{-i}(\sigma_i, s_{-i}))$$

- Minimax utility: $\min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$


$$(\text{= - } \max_{\sigma_{-i}} \min_{s_i} u_{-i}(s_i, \sigma_{-i}))$$

- Minimax theorem:

$$\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i}) = \min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$$

- Minimax theorem does not hold with pure strategies only (example?)

Back to the poker-like game, again



		$\frac{2}{3}$ cc	cf	$\frac{1}{3}$ fc	ff
$\frac{1}{3}$	rr	0, 0	0, 0	1, -1	1, -1
$\frac{2}{3}$	rc	.5, -.5	1.5, -1.5	0, 0	1, -1
	cr	-.5, .5	-.5, .5	1, -1	1, -1
	cc	0, 0	1, -1	0, 0	1, -1

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

$$\text{utility for bb} = 0 \cdot P(\text{cc}) + 1 \cdot (1 - P(\text{cc})) = .5 \cdot P(\text{cc}) + 0 \cdot (1 - P(\text{cc})) = \text{utility for bs}$$
 That is, $P(\text{cc}) = \frac{2}{3}$
- To make player 2 indifferent between cc and fc, we need:

$$\text{utility for cc} = 0 \cdot P(\text{bb}) + (-.5) \cdot (1 - P(\text{bb})) = -1 \cdot P(\text{bb}) + 0 \cdot (1 - P(\text{bb})) = \text{utility for fc}$$
 That is, $P(\text{bb}) = \frac{1}{3}$

A brief history of the minimax theorem

Borel
some very
special cases of
the theorem



Émile Borel

1921-1927

von Neumann
complete proof



*John von
Neumann*

1928

Ville
new proof
related to
systems of
linear
inequalities
(in Borel's
book)

1938




*Oskar
Morgenstern*

1944

**von Neumann &
Morgenstern**
*Theory of Games
and Economic
Behavior*
new proof also based
on systems of linear
inequalities, inspired
by Ville's proof

von Neumann
explains to
Dantzig about
strong duality of
linear programs

1947



*George
Dantzig*

1951

**Gale-Kuhn-
Tucker**
proof of LP duality,
Dantzig
proof* of
equivalence to
zero-sum games,
both in
Koopmans' book
*Activity Analysis
of Production and
Allocation*

Computing minimax strategies

- maximize v_R **Row utility**
subject to
for all c , $\sum_r p_r u_R(r, c) \geq v_R$ **Column optimality**
 $\sum_r p_r = 1$ **distributional constraint**

Equilibrium notions for general-sum games

General-sum games

- You could still play a minimax strategy in general-sum games
 - I.e., pretend that the opponent is only trying to hurt you

- But this is not rational:

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

- If Column was trying to hurt Row, Column would play Left, so Row should play Down
- In reality, Column will play Right (strictly dominant), so Row should play Up
- Is there a better generalization of minimax strategies in zero-sum games to general-sum games?

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%

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

The presentation game



Put effort into presentation (E)

Do not put effort into presentation (NE)

	<i>Pay attention (A)</i>	<i>Do not pay attention (NA)</i>
<i>Put effort into presentation (E)</i>	2, 2	-1, 0
<i>Do not put effort into presentation (NE)</i>	-7, -8	0, 0

- Pure-strategy Nash equilibria: (E, A), (NE, NA)
- Mixed-strategy Nash equilibrium:
($(\frac{4}{5} E, \frac{1}{5} NE), (\frac{1}{10} A, \frac{9}{10} NA)$)
 - Utility $-\frac{7}{10}$ for presenter, 0 for audience

The “equilibrium selection problem”

- You are about to play a game **that you have never played before** with a person **that you have never met**
- According to which equilibrium should you play?
- Possible answers:
 - Equilibrium that maximizes the sum of utilities (**social welfare**)
 - Or, at least not a Pareto-dominated equilibrium
 - So-called **focal** equilibria
 - “Meet in Paris” game: *You and a friend were supposed to meet in Paris at noon on Sunday, but you forgot to discuss where and you cannot communicate. All you care about is meeting your friend. Where will you go?*
 - Equilibrium that is the convergence point of some **learning process**
 - An equilibrium that is **easy to compute**
 - ...
- **Equilibrium selection is a difficult problem**

Computing a single Nash equilibrium



“Together with factoring, the complexity of finding a Nash equilibrium is in my opinion the most important concrete open question on the boundary of P today.”

Christos Papadimitriou,

STOC'01

[’91]

- **PPAD-complete** to compute one Nash equilibrium in a two-player game [Daskalakis, Goldberg, Papadimitriou STOC’06 / SIAM J. Comp. ’09; Chen & Deng FOCS’06 / Chen, Deng, Teng JACM’09]
- Is one Nash equilibrium all we need to know?

A useful reduction (SAT \rightarrow game)

[C. & Sandholm IJCAI'03, Games and Economic Behavior '08]

(Earlier reduction with weaker implications: Gilboa & Zemel GEB '89)

Formula: $(x_1 \text{ or } -x_2) \text{ and } (-x_1 \text{ or } x_2)$

Solutions: $x_1=\text{true}, x_2=\text{true}$
 $x_1=\text{false}, x_2=\text{false}$

Game:

	x_1	x_2	$+x_1$	$-x_1$	$+x_2$	$-x_2$	$(x_1 \text{ or } -x_2)$	$(-x_1 \text{ or } x_2)$	default
x_1	-2,-2	-2,-2	0,-2	0,-2	2,-2	2,-2	-2,-2	-2,-2	0,1
x_2	-2,-2	-2,-2	2,-2	2,-2	0,-2	0,-2	-2,-2	-2,-2	0,1
$+x_1$	-2,0	-2,2	1,1	-2,-2	1,1	1,1	-2,0	-2,2	0,1
$-x_1$	-2,0	-2,2	-2,-2	1,1	1,1	1,1	-2,2	-2,0	0,1
$+x_2$	-2,2	-2,0	1,1	1,1	1,1	-2,-2	-2,2	-2,0	0,1
$-x_2$	-2,2	-2,0	1,1	1,1	-2,-2	1,1	-2,0	-2,2	0,1
$(x_1 \text{ or } -x_2)$	-2,-2	-2,-2	0,-2	2,-2	2,-2	0,-2	-2,-2	-2,-2	0,1
$(-x_1 \text{ or } x_2)$	-2,-2	-2,-2	2,-2	0,-2	0,-2	2,-2	-2,-2	-2,-2	0,1
default	1,0	1,0	1,0	1,0	1,0	1,0	1,0	1,0	ϵ, ϵ

- Every satisfying assignment (if there are any) corresponds to an equilibrium with utilities 1, 1; exactly one additional equilibrium with utilities ϵ, ϵ that always exists
- Evolutionarily stable strategies Σ_2^P -complete [C. WINE 2013]

Some algorithm families for computing Nash equilibria of 2-player normal-form games






image from von Stengel

Lemke-Howson [J. SIAM '64]

Exponential time due to Savani & von Stengel [FOCS'04 / Econometrica'06]



Special cases / subroutines

[C. & Sandholm AAI'05, AAMAS'06; Benisch, Davis, Sandholm AAI'06 / JAIR'10; Kontogiannis & Spirakis APPROX'11; Adsul, Garg, Mehta, Sohoni STOC'11; ...]

- for both i , for any $s_i \in S_i - X_i$, $p_i(s_i) = 0$
- for both i , for any $s_i \in X_i$, $\sum p_{-i}(s_{-i})u_i(s_i, s_{-i}) = u_i$
- for both i , for any $s_i \in S_i - X_i$, $\sum p_{-i}(s_{-i})u_i(s_i, s_{-i}) \leq u_i$

Search over supports / MIP

[Dickhaut & Kaplan, Mathematica J. '91]

[Porter, Nudelman, Shoham AAI'04 / GEB'08]

[Sandholm, Gilpin, C. AAI'05]

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

Approximate equilibria

[Brown '51 / C. '09 / Goldberg, Savani, Sørensen, Ventre '11; Althöfer '94, Lipton, Markakis, Mehta '03, Daskalakis, Mehta, Papadimitriou '06, '07, Feder, Nazerzadeh, Saberi '07, Tsaknakis & Spirakis '07, Spirakis '08, Bosse, Byrka, Markakis '07, ...]

Search-based approaches (for 2 players)

- Suppose we know the **support** X_i of each player i 's mixed strategy in equilibrium
 - That is, which pure strategies receive positive probability
- Then, we have a linear feasibility problem:
 - for both i , for any $s_i \in S_i - X_i$, $p_i(s_i) = 0$
 - for both i , for any $s_i \in X_i$, $\sum p_{-i}(s_{-i}) u_i(s_i, s_{-i}) = u_i$
 - for both i , for any $s_i \in S_i - X_i$, $\sum p_{-i}(s_{-i}) u_i(s_i, s_{-i}) \leq u_i$
- Thus, we can search over possible supports
 - This is the basic idea underlying methods in [\[Dickhaut & Kaplan 91; Porter, Nudelman, Shoham AAI04/GEB08\]](#)
- Dominated strategies can be eliminated

Solving for a Nash equilibrium using MIP (2 players)

[Sandholm, Gilpin, C. AAAI'05]

- maximize *whatever you like (e.g., social welfare)*
- subject to
 - for both i , for any s_i , $\sum_{s_{-i}} \mathbf{p}_{s_{-i}} u_i(s_i, s_{-i}) = \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i \geq \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{p}_{s_i} \leq \mathbf{b}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i - \mathbf{u}_{s_i} \leq M(1 - \mathbf{b}_{s_i})$
 - for both i , $\sum_{s_i} \mathbf{p}_{s_i} = 1$
- \mathbf{b}_{s_i} is a binary variable indicating whether s_i is in the support, M is a large number


Lemke-Howson algorithm (1-slide sketch!)

	GREEN	ORANGE
RED	1, 0	0, 1
BLUE	0, 2	1, 0

player 2's utility as function of 1's mixed strategy

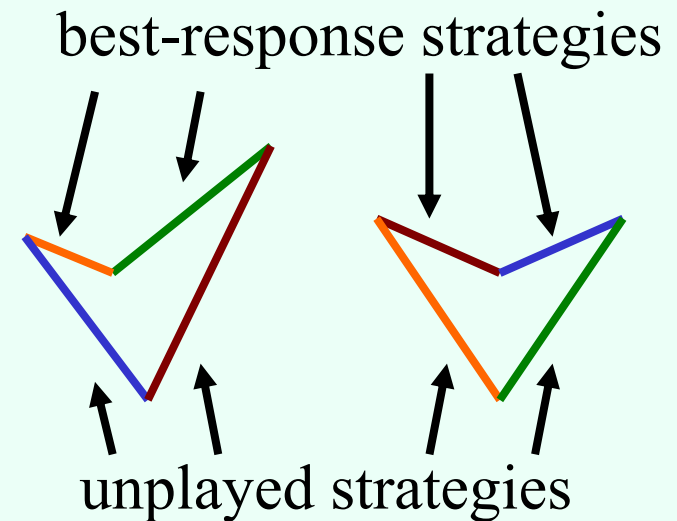


player 1's utility as function of 2's mixed strategy




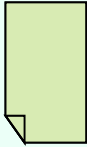
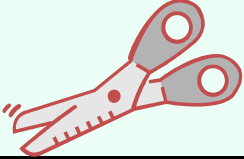

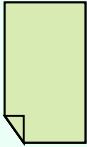
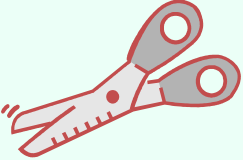
RED BLUE GREEN ORANGE

redraw both



- Strategy profile = pair of points
- Profile is an equilibrium iff every pure strategy is either a best response or unplayed
- I.e. equilibrium = pair of points that includes all the colors
 - ... except, pair of bottom points doesn't count (the "artificial equilibrium")
- Walk in some direction from the artificial equilibrium; at each step, throw out the color used twice

Correlated equilibrium [Aumann '74]

			
	0, 0 0	0, 1 1/6	1, 0 1/6
	1, 0 1/6	0, 0 0	0, 1 1/6
	0, 1 1/6	1, 0 1/6	0, 0 0

Correlated equilibrium LP

maximize *whatever*

subject to

$$\text{for all } r \text{ and } r', \quad \sum_c p_{r,c} u_R(r, c) \geq \sum_c p_{r,c} u_R(r', c)$$

Row incentive constraint

$$\text{for all } c \text{ and } c', \quad \sum_r p_{r,c} u_C(r, c) \geq \sum_r p_{r,c} u_C(r, c')$$


Column incentive constraint

$$\sum_{r,c} p_{r,c} = 1 \quad \text{distributional constraint}$$

Recent developments

Questions raised by security games

- Equilibrium **selection**?



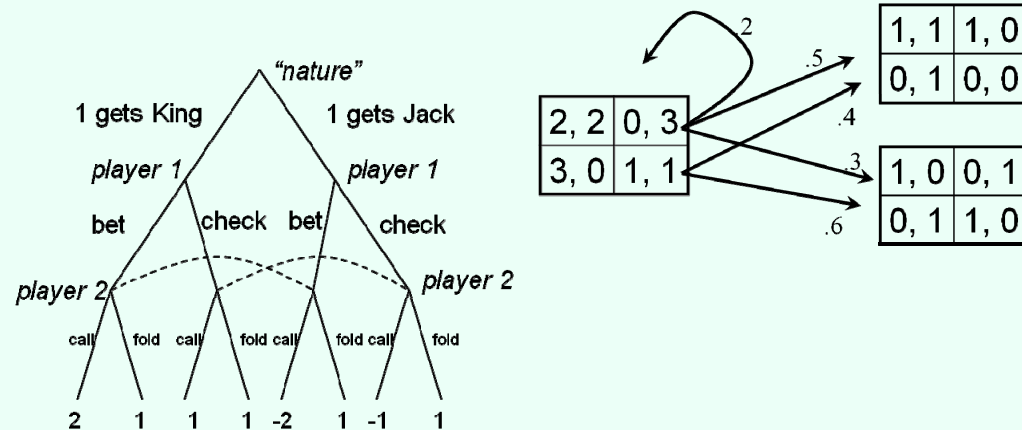
	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- How should we model **temporal / information structure**?

2, 2	-1, 0
-7, -8	0, 0

	L	R		L	R		
row player	U	4	6	column player	U	4	6
type 1 (prob. 0.5)	D	2	4	type 1 (prob. 0.5)	D	4	6





	L	R		L	R		
row player	U	2	4	column player	U	2	2
type 2 (prob. 0.5)	D	4	2	type 2 (prob. 0.5)	D	4	2







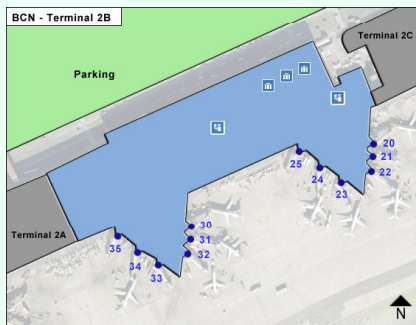
- What structure should **utility functions** have?
- Do our algorithms **scale**?

Observing the defender's distribution in security






Terminal A



Terminal B



observe



Mo

Tu

We

Th

Fr

Sa

This model is not uncontroversial... [Pita, Jain, Tambe, Ordóñez, Kraus AIJ'10; Korzhyk, Yin, Kiekintveld, C., Tambe JAIR'11; Korzhyk, C., Parr AAMAS'11]

Commitment (Stackelberg strategies)

Commitment

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

Unique Nash
equilibrium (iterated
strict dominance
solution)




von Stackelberg

- 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

	0	1
.49	1, 1	3, 0
.51	0, 0	2, 1

Sometimes also called a **Stackelberg (mixed) 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

[C. & Sandholm EC'06, von Stengel & Zamir GEB'10]

- Separate LP for every column c^* :

maximize $\sum_r p_r u_R(r, c^*)$ Row utility

subject to

for all c , $\sum_r p_r u_C(r, c^*) \geq \sum_r p_r u_C(r, c)$ Column optimality

$\sum_r p_r = 1$ distributional constraint

On the game we saw before

x	1, 1	3, 0
y	0, 0	2, 1

maximize $1x + 0y$

subject to

$$1x + 0y \geq 0x + 1y$$

$$x + y = 1$$

$$x \geq 0$$

$$y \geq 0$$

maximize $3x + 2y$

subject to

$$0x + 1y \geq 1x + 0y$$


$$x + y = 1$$

$$x \geq 0$$

$$y \geq 0$$

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



Generalizing beyond zero-sum games


Minimax, **Nash**, Stackelberg all agree in zero-sum games




minimax strategies



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



Nash equilibrium



Stackelberg mixed strategies

Other nice properties of commitment to mixed strategies

- No equilibrium selection problem



0, 0	-1, 1
1, -1	-5, -5

- Leader's payoff at least as good as any Nash eq. or even correlated eq. (von Stengel & Zamir [GEB '10]; see also C. & Korzhyk [AAAI '11], Letchford, Korzhyk, C. [JAAMAS'14])




IV



Some other work on commitment in unrestricted games

2, 2	-1, 0
-7, -8	0, 0

normal-form games



extensive-form games


[Letchford & C., EC'10]

- learning to commit [Letchford, C., Munagala SAGT'09]
- correlated strategies [C. & Korzhyk AAI'11]
- uncertain observability [Korzhyk, C., Parr AAMAS'11]

		L	R		L	R	
row player	U	4	6	column player	U	4	6
type 1 (prob. 0.5)	D	2	4	type 1 (prob. 0.5)	D	4	6
		L	R		L	R	
row player	U	2	4	column player	U	2	2
type 2 (prob. 0.5)	D	4	2	type 2 (prob. 0.5)	D	4	2

commitment in Bayesian games

- [C. & Sandholm EC'06; Paruchuri, Pearce, Marecki, Tambe, Ordóñez, Kraus AAMAS'08; Letchford, C., Munagala SAGT'09; Pita, Jain, Tambe, Ordóñez, Kraus AIJ'10; Jain, Kiekintveld, Tambe AAMAS'11; ...]



stochastic games

[Letchford, MacDermed, C., Parr, Isbell, AAI'12]

Security games

Example security game


- 3 airport terminals to defend (A, B, C)
- Defender can place checkpoints at 2 of them
- Attacker can attack any 1 terminal

	A	B	C
{A, B}	0, -1	0, -1	-2, 3
{A, C}	0, -1	-1, 1	0, 0
{B, C}	-1, 1	0, -1	0, 0

Security resource allocation games


[Kiekintveld, Jain, Tsai, Pita, Ordóñez, Tambe AAMAS'09]

- Set of targets T
- Set of security resources Ω available to the defender (leader)
- Set of schedules $\mathcal{S} \subseteq 2^T$
- Resource ω can be assigned to one of the schedules in $A(\omega) \subseteq \mathcal{S}$
- 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
- $U_d^c(t) \geq U_d^u(t); U_a^c(t) \leq U_a^u(t)$




Game-theoretic properties of security resource allocation games [Korzhyk, Yin, Kiekintveld, C., Tambe JAIR'11]

- For the defender:
Stackelberg strategies are also Nash strategies
 - minor assumption needed
 - not true with multiple attacks
- Interchangeability property for Nash equilibria (“solvable”)
 - no equilibrium selection problem
 - still true with multiple attacks[Korzhyk, C., Parr IJCAI'11]





1, 2	1, 0	2, 2
1, 1	1, 0	2, 1
0, 1	0, 0	0, 1

Scalability in security games



basic model

[Kiekintveld, Jain, Tsai, Pita, Ordóñez, Tambe AAMAS'09; Korzhyk, C., Parr, AAI'10; Jain, Kardeş, Kiekintveld, Ordóñez, Tambe AAI'10; Korzhyk, C., Parr, IJCAI'11]



*games on graphs
(usually zero-sum)*

[Halvorson, C., Parr IJCAI'09; Tsai, Yin, Kwak, Kempe, Kiekintveld, Tambe AAI'10; Jain, Korzhyk, Vaněk, C., Pěchouček, Tambe AAMAS'11; Jain, C., Tambe AAMAS'13; Xu, Fang, Jiang, C., Dughmi, Tambe AAI'14]

Techniques:

compact linear/integer programs

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

Subject to $\forall \omega: \sum_s c_{\omega,s} \leq 1$

$\forall t: \sum_{\omega} \sum_{s \in S} c_{\omega,s} \leq 1$ Marginal probability of t^* being defended (?)

Distributional constraints

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

min subject to



strategy generation



$$\begin{aligned}
 & \sigma_h(s_{h_0}) + \dots \sigma_h(s_{h_k}) \\
 & u \geq \sigma_h(s_{h_0}) \cdot u(s_{s_0}, s_{h_0}) + \dots \sigma_h(s_{h_2}) \cdot u(s_{s_0}, s_{h_2}) \\
 & u \geq \sigma_h(s_{h_0}) \cdot u(s_{s_1}, s_{h_0}) + \dots \sigma_h(s_{h_2}) \cdot u(s_{s_1}, s_{h_2}) \\
 & \vdots \\
 & u \geq \sigma_h(s_{h_0}) \cdot u(s_{s_k}, s_{h_0}) + \dots \sigma_h(s_{h_k}) \cdot u(s_{s_k}, s_{h_k})
 \end{aligned} = 1$$

Compact LP

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

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

Subject to

$$\forall \omega : \sum_s c_{\omega,s} \leq 1$$

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

Marginal probability
of t^* being defended (?)


Distributional constraints

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

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

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

Birkhoff-von Neumann theorem


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

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

$$= .1 \begin{array}{|c|c|c|} \hline 1 & 0 & 0 \\ \hline 0 & 0 & 1 \\ \hline 0 & 1 & 0 \\ \hline \end{array} + .1 \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 0 & 0 & 1 \\ \hline 1 & 0 & 0 \\ \hline \end{array} + .5 \begin{array}{|c|c|c|} \hline 0 & 0 & 1 \\ \hline 0 & 1 & 0 \\ \hline 1 & 0 & 0 \\ \hline \end{array} + .3 \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 1 & 0 & 0 \\ \hline 0 & 0 & 1 \\ \hline \end{array}$$

- 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

Schedules of size 1 using BvN



	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



Algorithms & complexity

[Korzhyk, C., Parr AAI'10]

		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

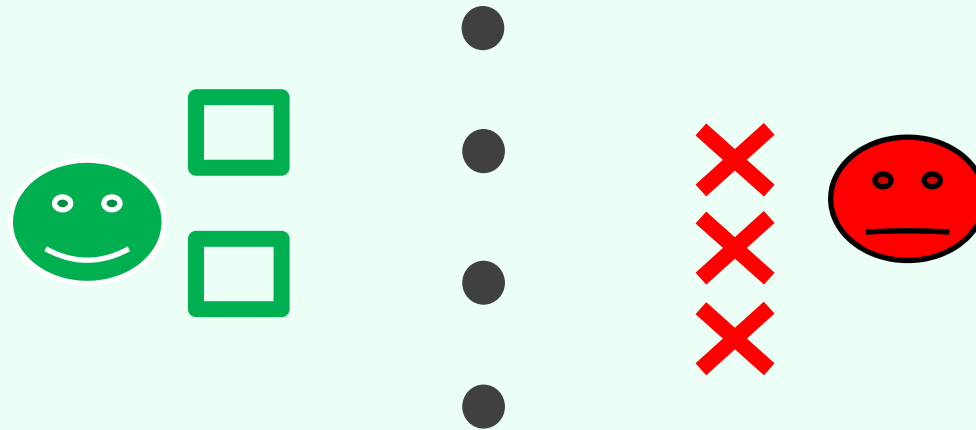
Also: security games on graphs

[Letchford, C. AAI'13]

Security games with multiple attacks

[Korzhyk, Yin, Kiekintveld, C., Tambe JAIR'11]

- The attacker can choose multiple targets to attack




- The utilities are added over all attacked targets
- Stackelberg NP-hard; Nash polytime-solvable and interchangeable [Korzhyk, C., Parr IJCAI'11]
 - Algorithm generalizes ORIGAMI algorithm for single attack [Kiekintveld, Jain, Tsai, Pita, Ordóñez, Tambe AAMAS'09]

Actual Security Schedules: Before vs. After


Boston, Coast Guard – “PROTECT” algorithm

slide courtesy of Milind Tambe

Before PROTECT



After PROTECT



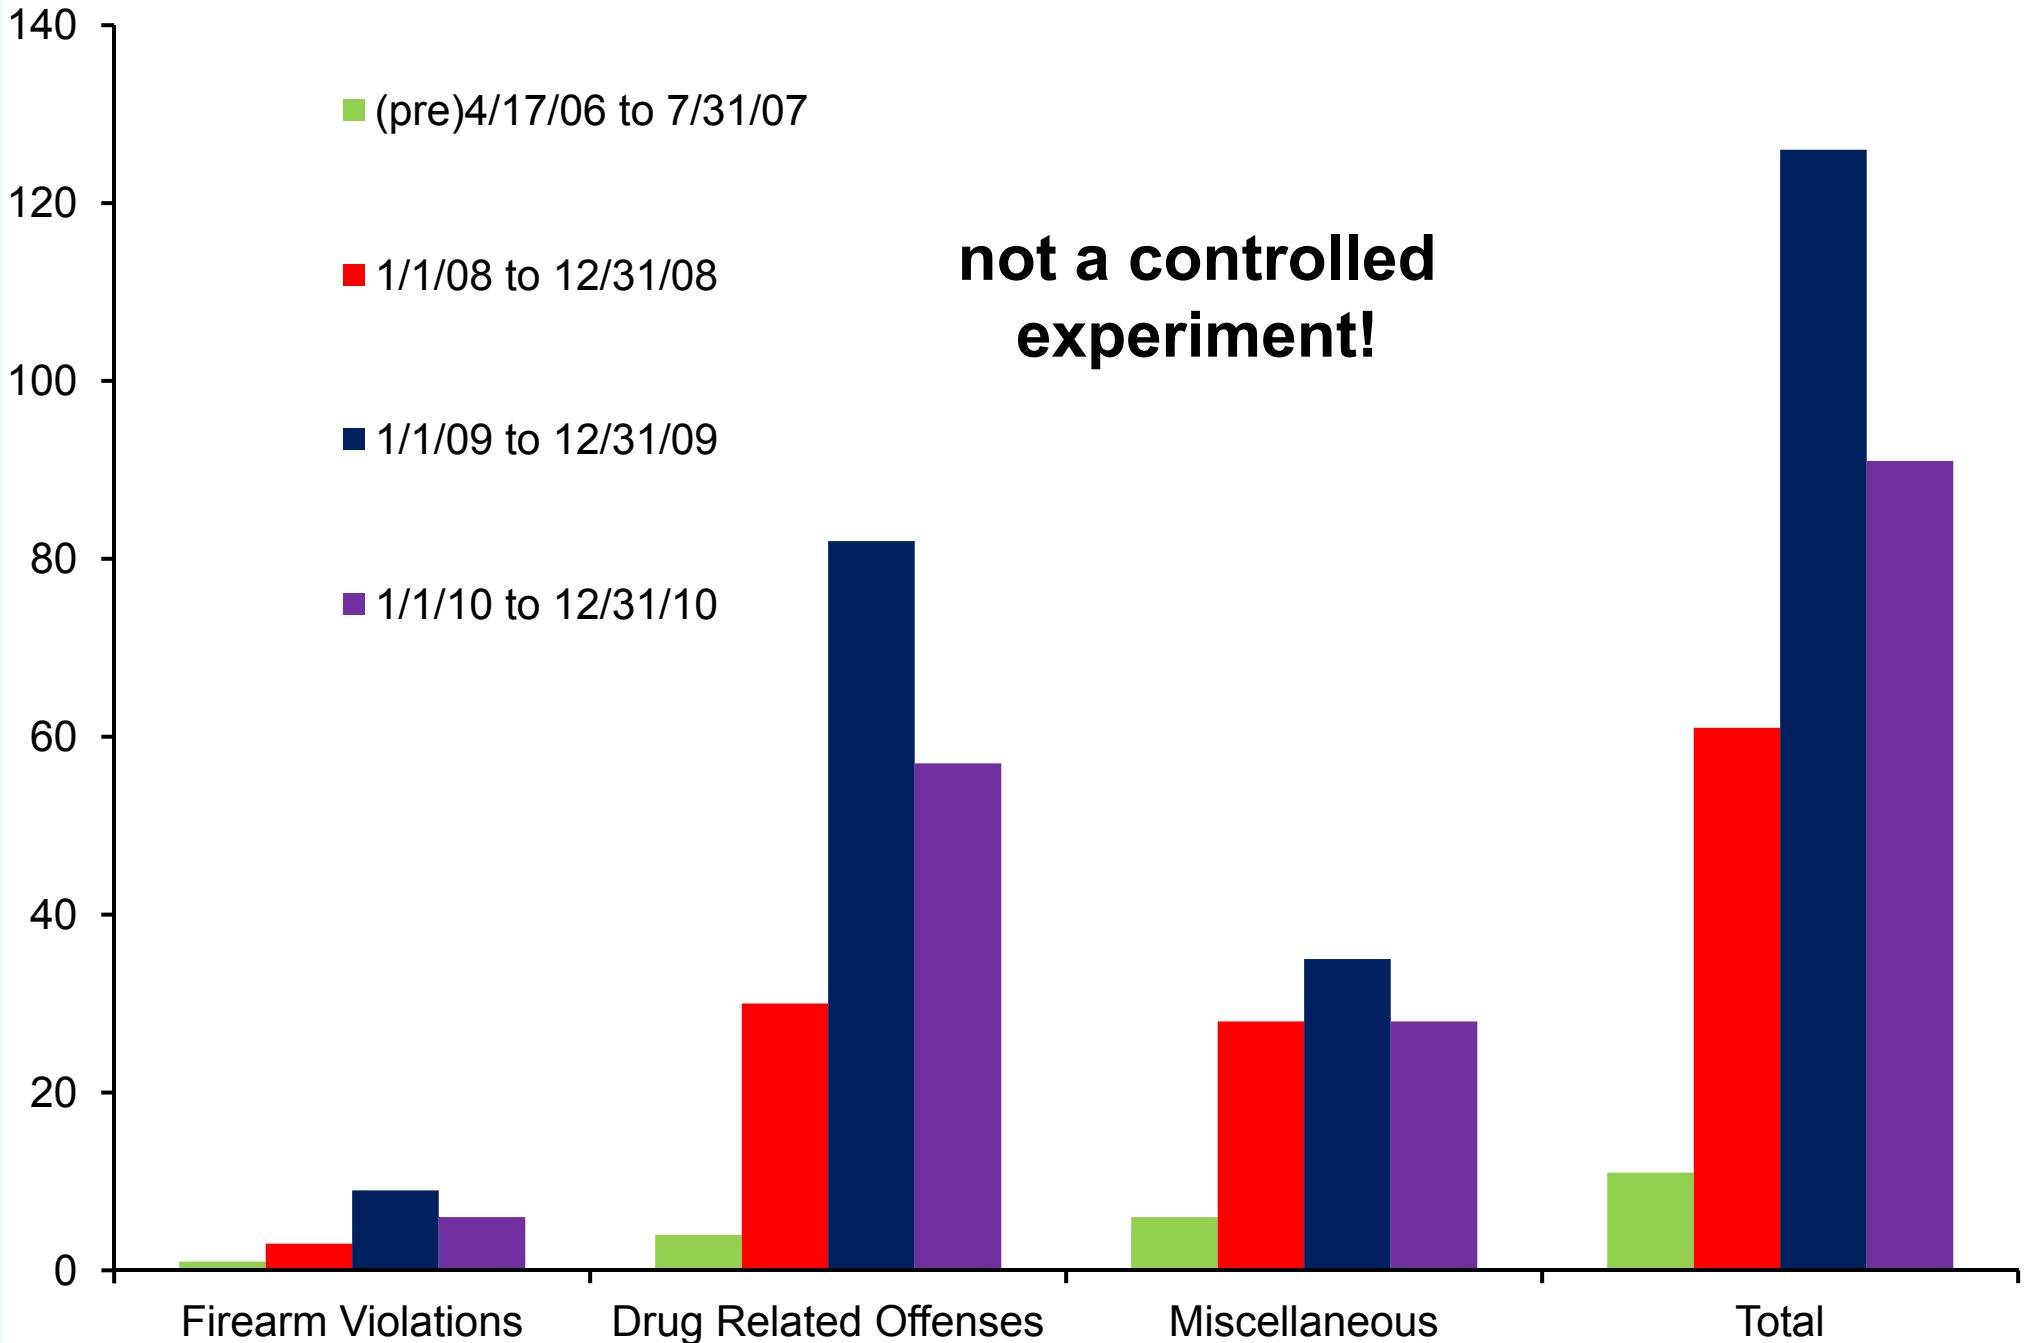
Industry port partners comment:

“The Coast Guard seems to be everywhere, all the time.”

Data from LAX checkpoints


*slide courtesy of
Milind Tambe*

before and after “ARMOR” algorithm



Placing checkpoints in a city

[Tsai, Yin, Kwak, Kempe, Kiekintveld, Tambe AAI'10; Jain, Korzhyk, Vaněk, C., Pěchouček, Tambe AAMAS'11; Jain, C., Tambe AAMAS'13]



Learning in games

Learning in (normal-form) games

- Learn how to play a game by
 - playing it many times, and
 - updating your strategy based on experience
- Why?
 - Some of the game's utilities (especially the other players') may be unknown to you
 - The other players may not be playing an equilibrium strategy
 - Computing an optimal strategy can be hard
 - Learning is what humans typically do
 - ...
- Does learning converge to equilibrium?

Iterated best response

- In the first round, play something arbitrary
- In each following round, play a best response against what the other players played in the **previous** round
- If all players play this, it can converge (i.e., we reach an equilibrium) or cycle

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

rock-paper-scissors

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

a simple congestion game

- **Alternating best response**: players alternately change strategies: one player best-responds each odd round, the other best-responds each even round

Fictitious play [Brown 1951]

- In the first round, play something arbitrary
- In each following round, play a best response against the **empirical distribution** of the other players' play
 - I.e., as if other player randomly selects from his past actions
- Again, if this converges, we have a Nash equilibrium
- Can still fail to converge...


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



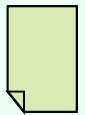

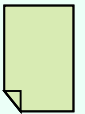
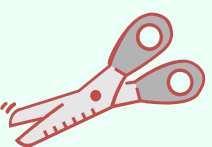
rock-paper-scissors

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






a simple congestion game

Fictitious play on rock-paper- scissors





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

Row



Column



30% R, 50% P, 20% S

30% R, 20% P, 50% S

Does the empirical distribution of play converge to equilibrium?

- ... for iterated best response?
- ... for fictitious play?

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

Fictitious play is guaranteed to converge in...

- Two-player zero-sum games [Robinson 1951]
- Generic 2x2 games [Miyasawa 1961]
- Games solvable by iterated strict dominance [Nachbar 1990]
- Weighted potential games [Monderer & Shapley 1996]
- **Not** in general [Shapley 1964]
- But, fictitious play always converges to the set of $\frac{1}{2}$ -approximate equilibria [C. 2009; more detailed analysis by Goldberg, Savani, Sørensen, Ventre 2011]

Shapley's game on which fictitious play does not converge

starting with (U, C):

0, 0	0, 1	1, 0
1, 0	0, 0	0, 1
0, 1	1, 0	0, 0

“Teaching”

- Suppose you are playing against a player that uses one of these learning strategies
 - Fictitious play, anything with no regret, ...
- Also suppose you are **very patient**, i.e., you only care about what happens in the long run
- How will you (the row player) play in the following repeated games?
 - Hint: the other player will **eventually best-respond** to whatever you do

4, 4	3, 5
5, 3	0, 0

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

- Note relationship to optimal strategies to commit to
- There is some work on learning strategies that are in **equilibrium** with each other [Brafman & Tennenholtz AIJ04]

Hawk-Dove Game

[Price and Smith, 1973]

	Dove	Hawk
Dove	1, 1	0, 2
Hawk	2, 0	-1, -1

- Unique *symmetric* equilibrium:
50% Dove, 50% Hawk

Evolutionary game theory

- Given: a symmetric 2-player game

	Dove	Hawk
Dove	1, 1	0, 2
Hawk	2, 0	-1, -1

- Population of players; players randomly matched to play game
- Each player plays a pure strategy
 - p_s = fraction of players playing strategy s
 - p = vector of all fractions p_s (the state)
- Utility for playing s is $u(s, p) = \sum_{s'} p_{s'} u(s, s')$
- Players reproduce at rate proportional to their utility; their offspring play the same strategy
 - $$dp_s(t)/dt = p_s(t)(u(s, p(t)) - \sum_{s'} p_{s'} u(s', p(t)))$$
 - Replicator dynamic
- What are the steady states?

Stability

	Dove	Hawk
Dove	1, 1	0, 2
Hawk	2, 0	-1, -1

- A steady state is stable if slightly perturbing the state will not cause us to move far away from the state
- **Proposition:** every stable steady state is a Nash equilibrium of the symmetric game
- Slightly stronger criterion: a state is **asymptotically stable** if it is stable, and after slightly perturbing this state, we will (in the limit) return to this state

Evolutionarily stable strategies

[Price and Smith, 1973]

- Now suppose players play **mixed** strategies
- A (single) mixed strategy σ is **evolutionarily stable** if the following is true:
 - Suppose all players play σ
 - Then, whenever a very small number of **invaders** enters that play a different strategy σ' ,

the players playing σ must get strictly **higher** utility than those playing σ' (i.e., σ must be able to **repel invaders**)

Properties of ESS

- **Proposition.** A strategy σ is evolutionarily stable if and only if the following conditions both hold:

(1) For all σ' , we have $u(\sigma, \sigma) \geq u(\sigma', \sigma)$ (i.e., **symmetric Nash equilibrium**)

(2) For all $\sigma' (\neq \sigma)$ with $u(\sigma, \sigma) = u(\sigma', \sigma)$, we have $u(\sigma, \sigma') > u(\sigma', \sigma')$

- **Theorem** [Taylor and Jonker 1978, Hofbauer et al. 1979, Zeeman 1980].

Every ESS is asymptotically stable under the

replicator dynamic. (Converse does not hold [van Damme 1987].)

Invasion (1/2)

	Dove	Hawk
Dove	1, 1	0, 2
Hawk	2, 0	-1, -1




- Given: population P_1 that plays $\sigma = 40\%$ Dove, 60% Hawk
- Tiny population P_2 that plays $\sigma' = 70\%$ Dove, 30% Hawk **invades**
- $u(\sigma, \sigma) = .16*1 + .24*2 + .36*(-1) = .28$ but
 $u(\sigma', \sigma) = .28*1 + .12*2 + .18*(-1) = .34$
- σ' (initially) grows in the population; invasion is **successful**

Invasion (2/2)

	Dove	Hawk
Dove	1, 1	0, 2
Hawk	2, 0	-1, -1

- Now P_1 plays $\sigma = 50\%$ Dove, 50% Hawk
- Tiny population P_2 that plays $\sigma' = 70\%$ Dove, 30% Hawk **invades**
- $u(\sigma, \sigma) = u(\sigma', \sigma) = .5$, so second-order effect:
- $u(\sigma, \sigma') = .35*1 + .35*2 + .15*(-1) = .9$ but
 $u(\sigma', \sigma') = .49*1 + .21*2 + .09*(-1) = .82$
- σ' shrinks in the population; invasion is **repelled**

Rock- Paper- Scissors



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

- Only one Nash equilibrium (Uniform)
- $u(\text{Uniform}, \text{Rock}) = u(\text{Rock}, \text{Rock})$
- No ESS

“Safe-Left-Right”

	Safe	Left	Right
Safe	1, 1	1, 1	1, 1
Left	1, 1	0, 0	2, 2
Right	1, 1	2, 2	0, 0

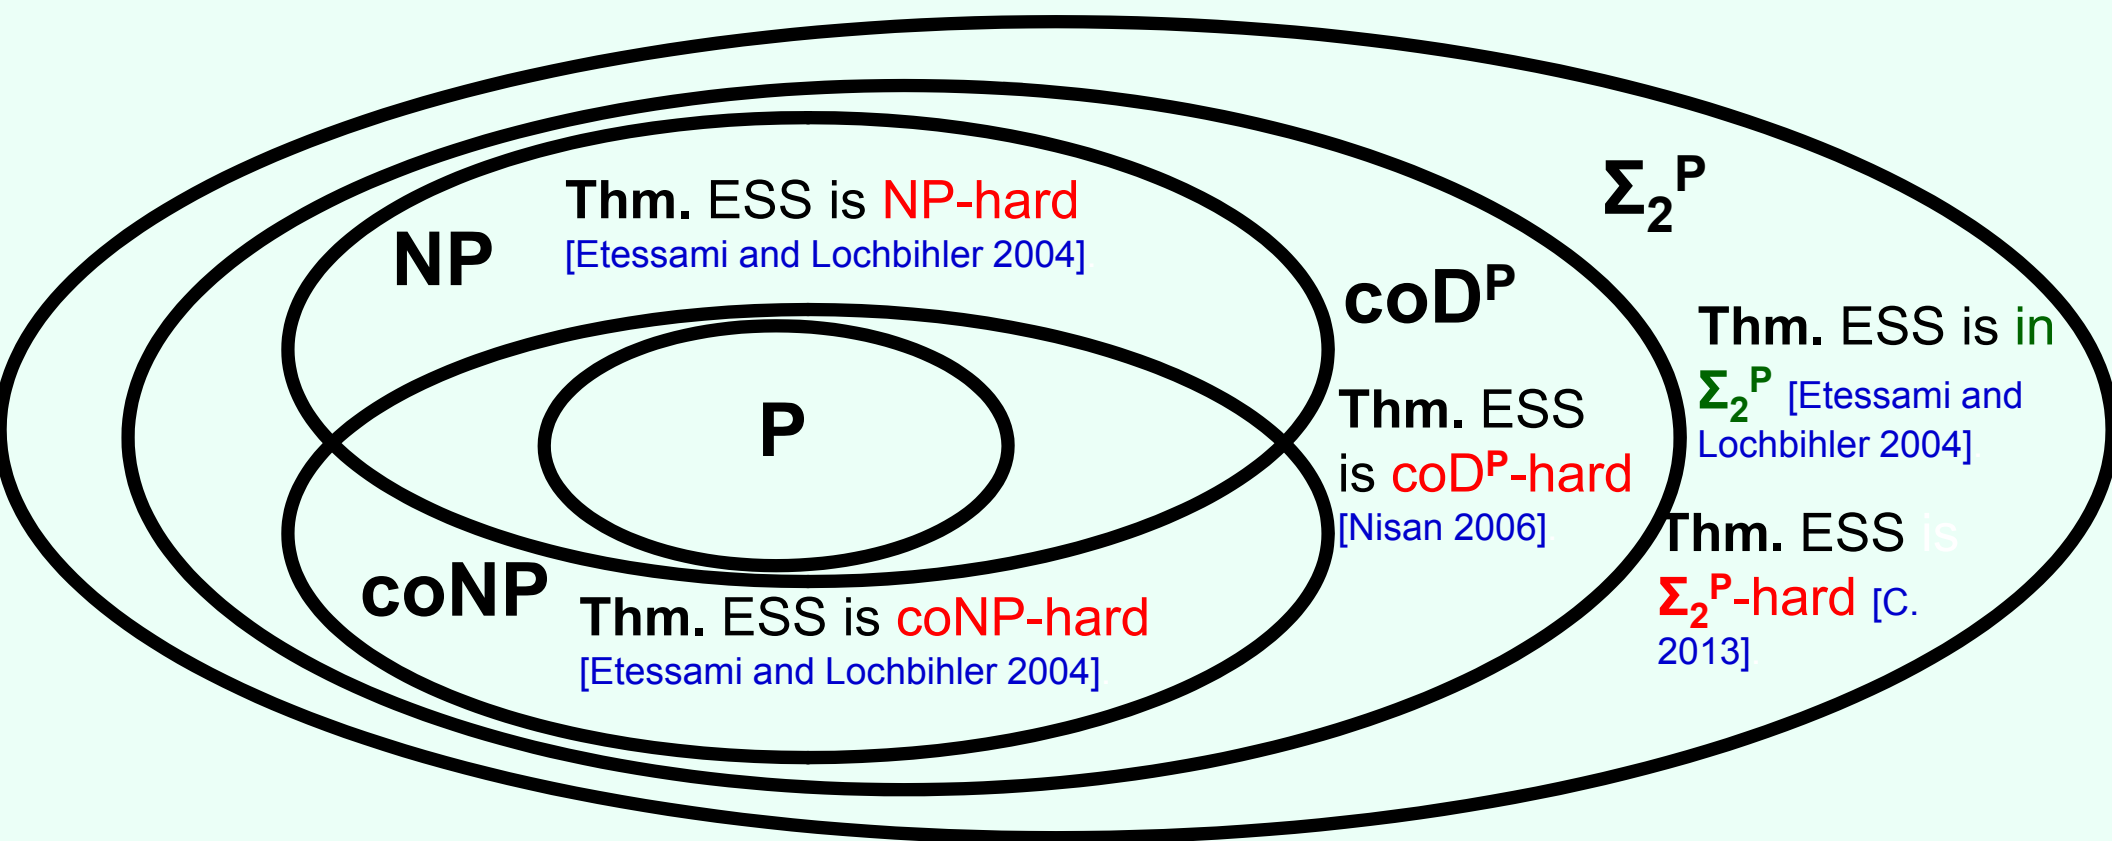
- Can 100% Safe be invaded?
- Is there an ESS?

The ESS problem

Input: symmetric 2-player normal-form game.

Q: Does it have an evolutionarily stable strategy?

(Hawk-Dove: yes. Rock-Paper-Scissors: no. Safe-Left-Right: no.)



The standard Σ_2^P -complete problem

Input: Boolean formula f over variables X_1 and X_2

Q: Does there exist an assignment of values to X_1 such that for every assignment of values to X_2 f is true?

Discussion of implications

- Many of the techniques for finding (optimal) Nash equilibria **will not extend to ESS**
- Evolutionary game theory gives a possible **explanation of how equilibria are reached...**
... for this purpose it would be good if its solution concepts aren't (very) hard to compute!

Learning in Stackelberg games

[Letchford, C., Munagala SAGT'09]


See also here at NIPS'14: Blum, Haghtalab, Procaccia [Th54]

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

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

Learning in Stackelberg games...

[Letchford, C., Munagala SAGT'09]



Theorem. Finding 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 # of follower actions
- d is # of leader actions

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

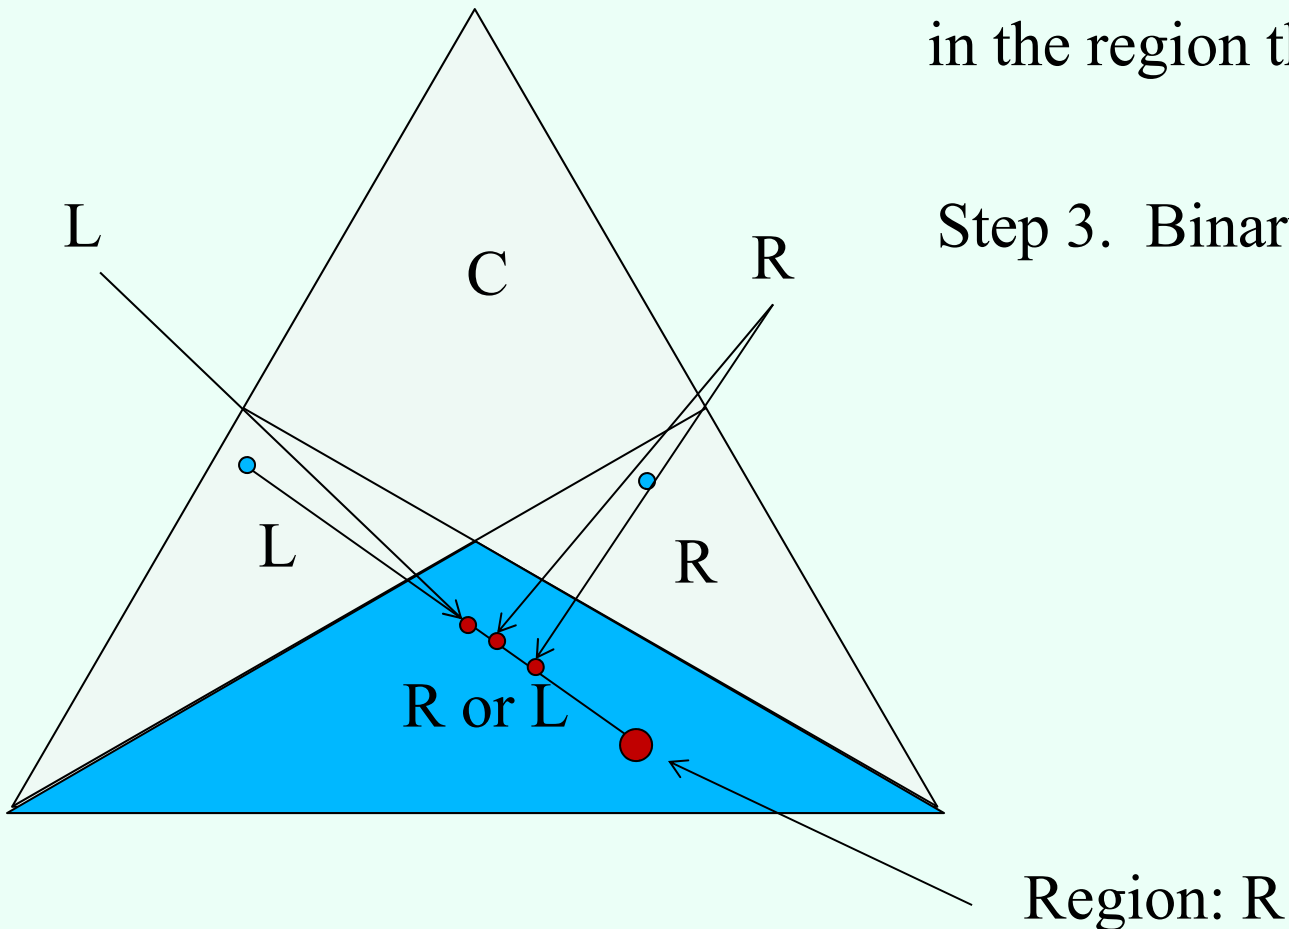
Finding a point on an unknown hyperplane

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

Intermediate state



Determining the hyperplane

Intermediate state



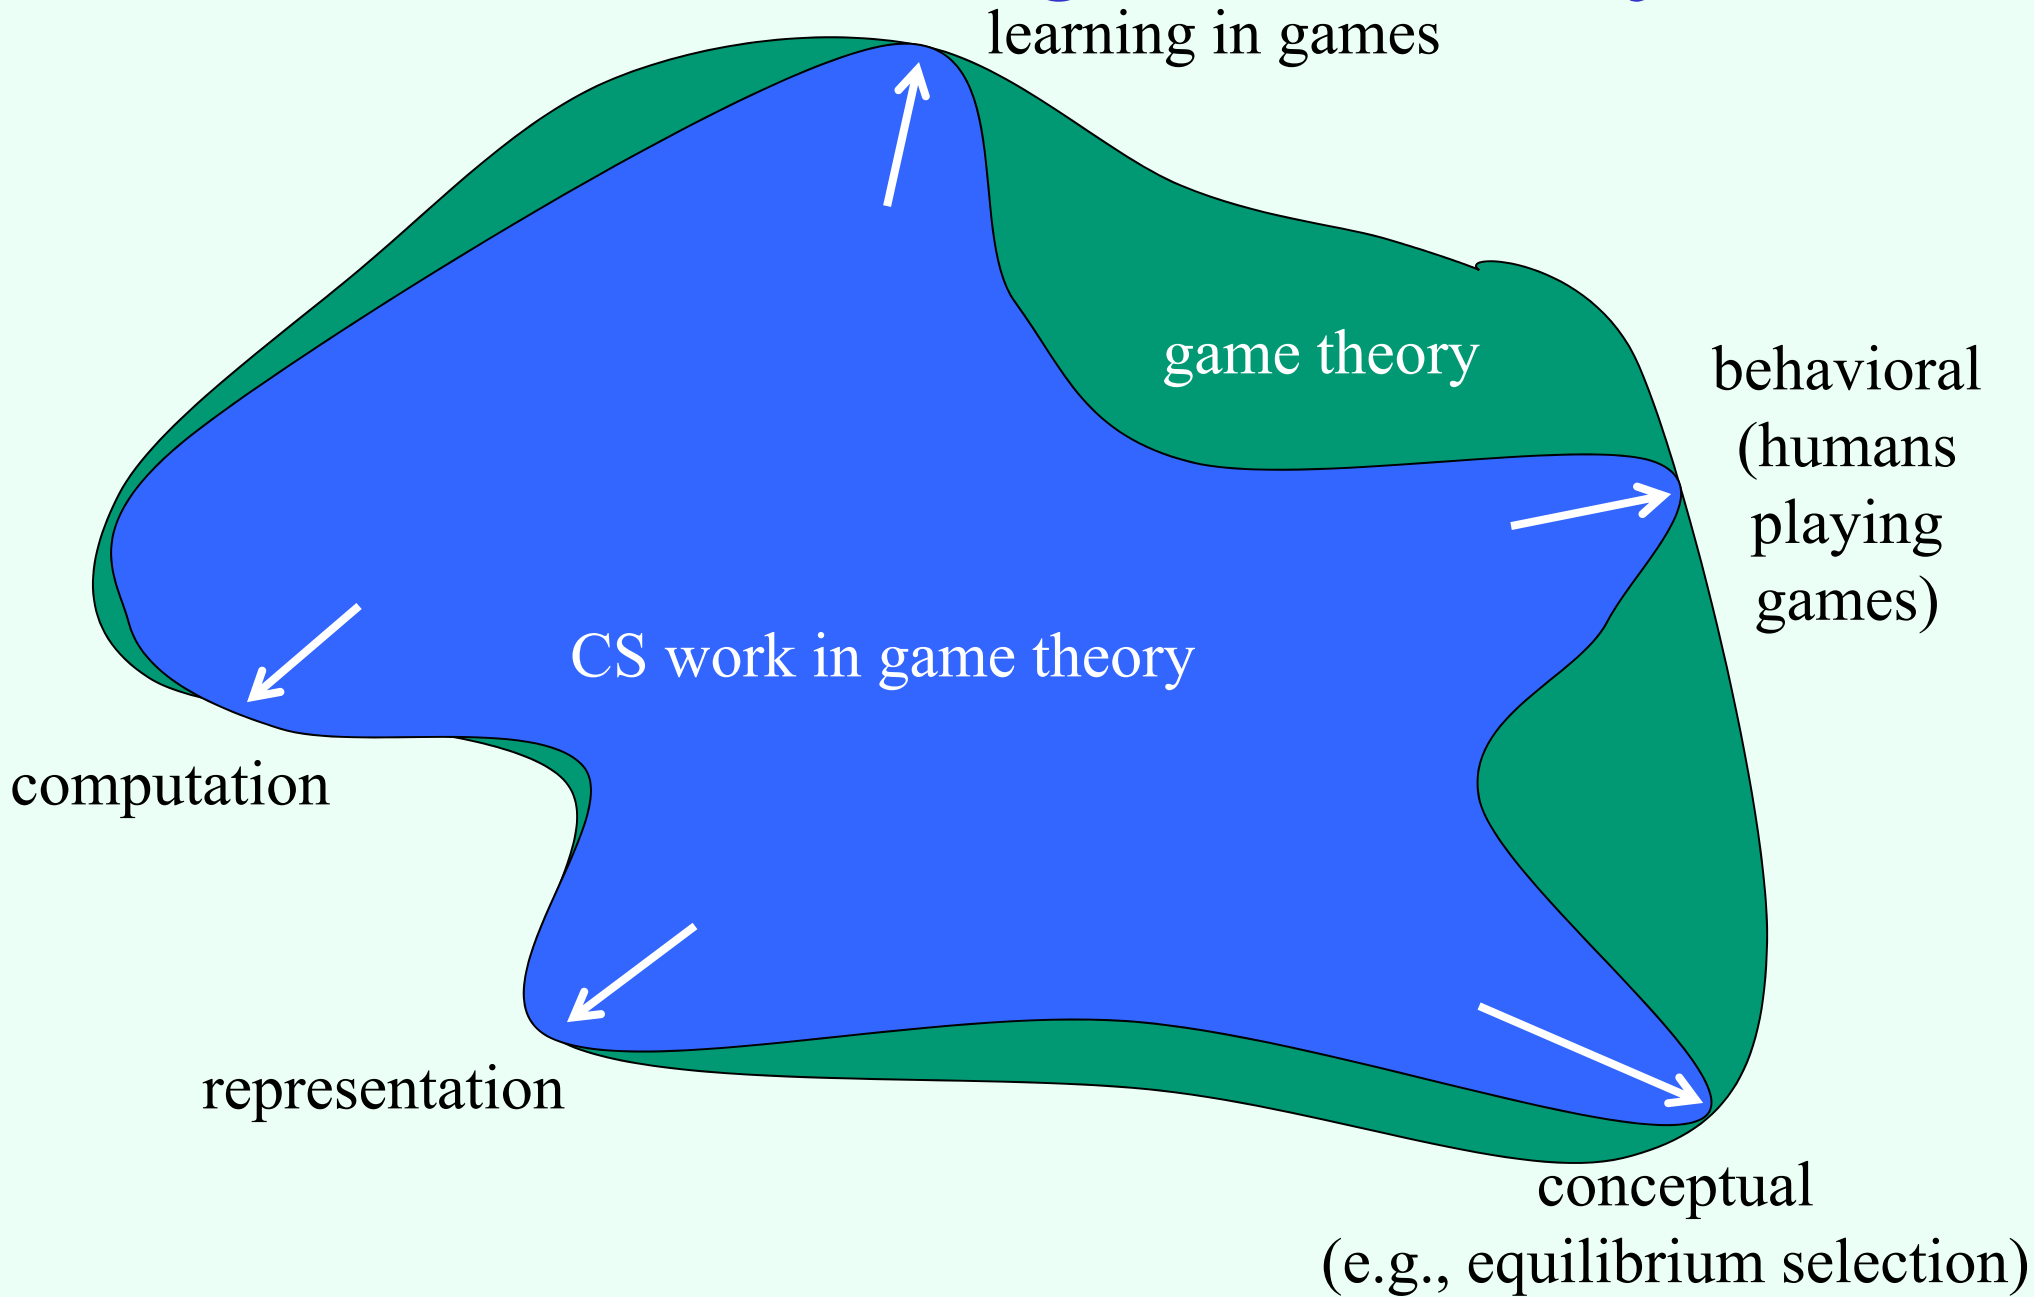
Step 1. Sample a regular d-simplex centered at the point

Step 2. Connect d lines between points on opposing sides

Step 3. Binary search along these lines

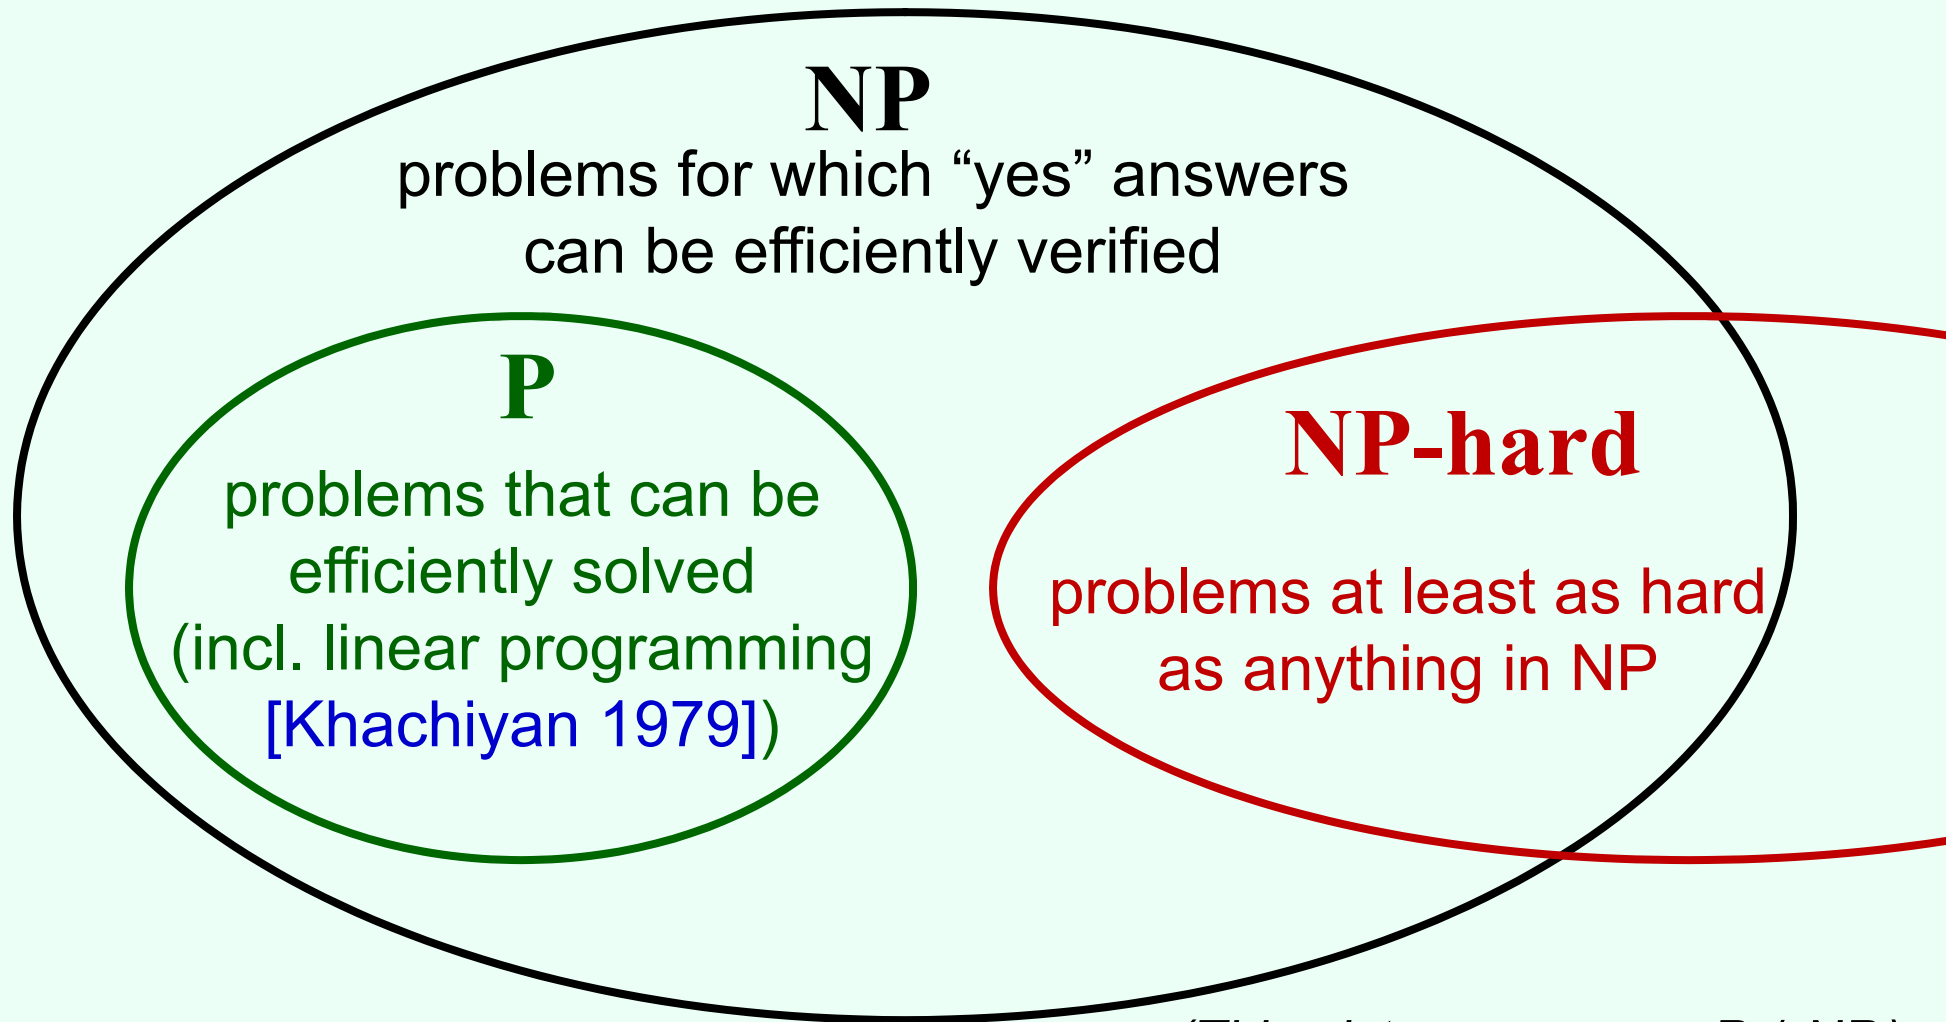
Step 4. Determine hyperplane (and update the region estimates with this information)

In summary: CS pushing at some of the boundaries of game theory



Backup slides

Computational complexity theory



(This picture assumes $P \neq NP$.)

- Is $P = NP$? [Cook 1971, Karp 1972, Levin 1973, ...]

Two computational questions for iterated dominance

- 1. Can a **given strategy** be eliminated using iterated dominance?
- 2. Is there some path of elimination by iterated dominance such that only **one strategy per player remains**?

- For strict dominance (with or without dominance by mixed strategies), both can be solved in polynomial time due to path-independence:
 - Check if any strategy is dominated, remove it, repeat
- For weak dominance, both questions are NP-hard (even when all utilities are 0 or 1), with or without dominance by mixed strategies [C., Sandholm 05]
 - Weaker version proved by [Gilboa, Kalai, Zemel 93]

Matching pennies with a sensitive target

		<i>Them</i>	
		L	R
<i>Us</i>	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

		<i>Them</i>	
		L	R
<i>Us</i>	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

		<i>Them</i>	
		L	R
<i>Us</i>	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)

Practice games

20, -20	0, 0
0, 0	10, -10

20, -20	0, 0	10, -10
0, 0	10, -10	8, -8

Correlated equilibrium as Bayes-Nash equilibrium

	$\theta_2=1$	$\theta_2=2$	$\theta_2=3$																											
$\theta_1=1$	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">0</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
$\theta_1=2$	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">0</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
$\theta_1=3$	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">1/6</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0	<table border="1"> <tr><td>0, 0</td><td>0, 1</td><td>1, 0</td></tr> <tr><td>1, 0</td><td>0, 0</td><td>0, 1</td></tr> <tr><td>0, 1</td><td>1, 0</td><td>0, 0</td></tr> </table> <p style="text-align: center;">0</p>	0, 0	0, 1	1, 0	1, 0	0, 0	0, 1	0, 1	1, 0	0, 0
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												
0, 0	0, 1	1, 0																												
1, 0	0, 0	0, 1																												
0, 1	1, 0	0, 0																												

Stackelberg mixed strategies deserve recognition as a separate solution concept!

- Seeing it only as a solution of a modified (extensive-form) game makes it hard to see...
 - when it **coincides** with other solution concepts
 - how **utilities compare** to other solution concepts
 - how to **compute** solutions
 - ...
- Does not mean it's not **also** useful to think of it as a backward induction solution
- Similar story for correlated equilibrium



Committing to a correlated strategy

[C. & Korzhyk AAI'11]

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

LP for optimal correlated strategy to commit to

maximize $\sum_{r,c} p_{r,c} u_C(r, c)$ leader utility

subject to

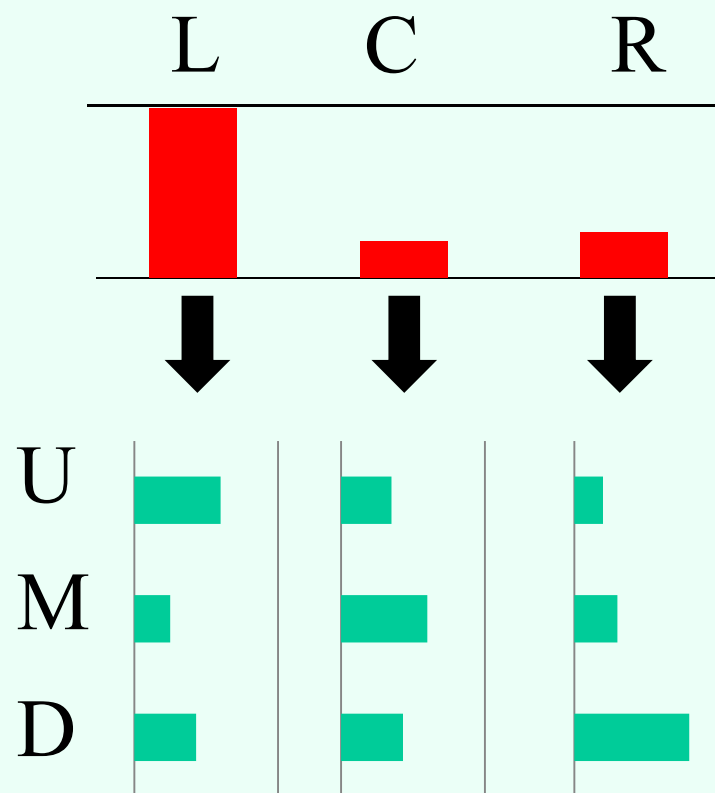
for all c and c' , $\sum_r p_{r,c} u_C(r, c) \geq \sum_r p_{r,c'} u_C(r, c')$

Column incentive constraint

$\sum_{r,c} p_{r,c} = 1$ distributional constraint


Equivalence to Stackelberg

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




3-player example

Leader



2 Utilities 1



Unique optimal correlated strategy to commit to:

50%



50%



Different from Stackelberg / CE

The Polynomial Hierarchy

$$\exists^p L = \{ \mathbf{x} \text{ in } \{0,1\}^* \mid (\exists \mathbf{w} \text{ in } \{0,1\}^{\leq p(|x|)}) (\mathbf{x}, \mathbf{w}) \text{ in } L \}$$

$$\forall^p L = \{ \mathbf{x} \text{ in } \{0,1\}^* \mid (\forall \mathbf{w} \text{ in } \{0,1\}^{\leq p(|x|)}) (\mathbf{x}, \mathbf{w}) \text{ in } L \}$$

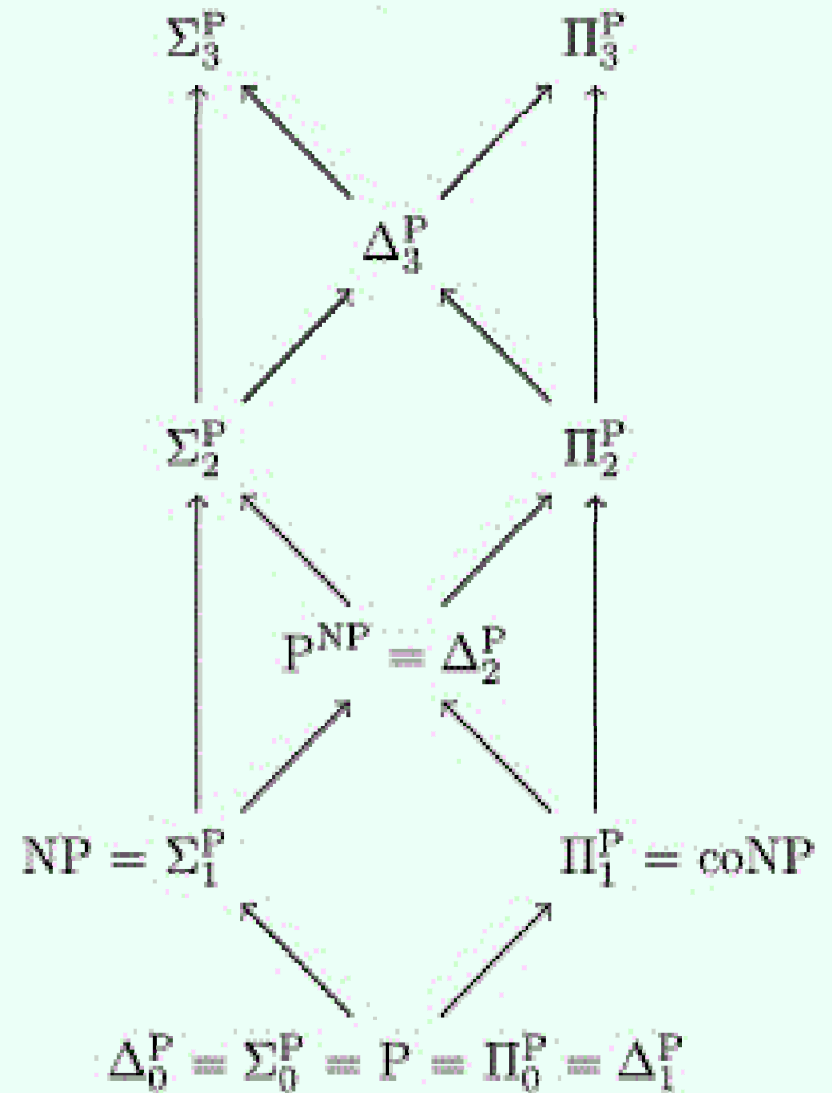
$$\exists^p C = \{ \exists^p L \mid p \text{ is a polynomial} \\ \text{and } L \text{ in } C \}$$

$$\forall^p C = \{ \forall^p L \mid p \text{ is a polynomial} \\ \text{and } L \text{ in } C \}$$

$$\Sigma_0^P = \Pi_0^P = P$$

$$\Sigma_{i+1}^P = \exists^P \Pi_i^P$$

$$\Pi_{i+1}^P = \forall^P \Sigma_i^P$$



The ESS-RESTRICTED-SUPPORT problem

Input: symmetric 2-player normal-form game, subset T of the strategies S

Q: Does the game have an evolutionarily stable strategy whose support is restricted to (a subset of) T ?

MINMAX-CLIQUE

proved $\Pi_2^P (= \text{co}\Sigma_2^P)$ -complete by Ko and Lin [1995]

Input: graph $G = (V, E)$, sets I and J , partition of V into subsets V_{ij} (for i in I and j in J), number k

Q: Is it the case that for every function $t : I \rightarrow J$, $\bigcup_i V_{i,t(i)}$ has a clique of size k ?

Thank you, compendium by Schaefer and Umans!





Illustration of reduction



T

	s_{11}	s_{12}	s_{21}	s_{22}	$s_{v_{11}}$	$s_{v_{12}}$	$s_{v_{21}}$	$s_{v_{22}}$	s_0
s_{11}	1	0	2	2	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{12}	0	1	2	2	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{21}	2	2	1	0	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
s_{22}	2	2	0	1	$3/2$	$3/2$	$3/2$	$3/2$	$3/2$
$s_{v_{11}}$	$3/2$	0	$3/2$	$3/2$	0	0	3	3	0
$s_{v_{12}}$	0	$3/2$	$3/2$	$3/2$	0	0	0	3	0
$s_{v_{21}}$	$3/2$	$3/2$	$3/2$	0	3	0	0	0	0
$s_{v_{22}}$	$3/2$	$3/2$	0	$3/2$	3	3	0	0	0
s_0	$3/2$	$3/2$	$3/2$	$3/2$	0	0	0	0	0

Unrestricted support?

- Just duplicate all the strategies outside T ...
- (Appendix: result still holds in games in which every pure strategy is the unique best response to some mixed strategy)

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