One Equilibrium Is Not Enough: Computing Game-Theoretic Solutions to Act Strategically

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

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

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

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

Vincent Conitzer
Duke University

My wonderful co-authors (alphabetically):

Krzysztof Apt, CWI Amsterdam. Sayan Bhattacharya, Duke. Craig Boutilier, U. Toronto. Andrew Davenport, IBM Research. Jonathan Derryberry, CMU. Bruce Donald, Duke. Joseph Farfel, Duke. Nikesh Garera, Johns Hopkins. Andrew Gilpin, CMU. Mingyu Guo, Liverpool. Erik Halvorson, Duke. Paul Harrenstein, TU Munich. Ryo Ichimura, Kyushu U. Nicole Immorlica, Northwestern. Atsushi Iwasaki, Kyushu U. Kamal Jain, MSR. Manish Jain, USC. Jayant Kalagnanam, IBM Research. Christopher Kiekintveld, UT El Paso. Dmytro Korzhyk, Duke. Jerome Lang, U. Paris-Dauphine. Joshua Letchford, Duke. Vangelis Markakis, Athens U. Econ. and Business. Kohki Maruono, Kyushu U. Kamesh Munagala, Duke. Yoshifusa Omori, Kyushu U. Naoki Ohta, Kyushu U. Ron Parr, Duke. Michal Pechoucek, Prague TU. Ariel Procaccia, Harvard U. Daniel Reeves, Yahoo! Research. Matthew Rognlie, MIT. Jeff Rosenschein, Hebrew U. Yuko Sakurai, Kyushu U. Tuomas Sandholm, CMU. Paolo Santi, IIT CNR. Yasufumi Satoh, Kyushu U. Peng Shi, MIT. Milind Tambe, USC. Taiki Todo, Kyushu U. Ondrej Vanek, Prague TU. Liad Wagman, Illinois Institute of Technology. Toby Walsh, NICTA and UNSW. Mathijs de Weerdt, TU Delft. Lirong Xia, Duke. Zhengyu Yin, USC. Makoto Yokoo, Kyushu U. Michael Zuckerman, Hebrew U.

Multiple entities with different interests

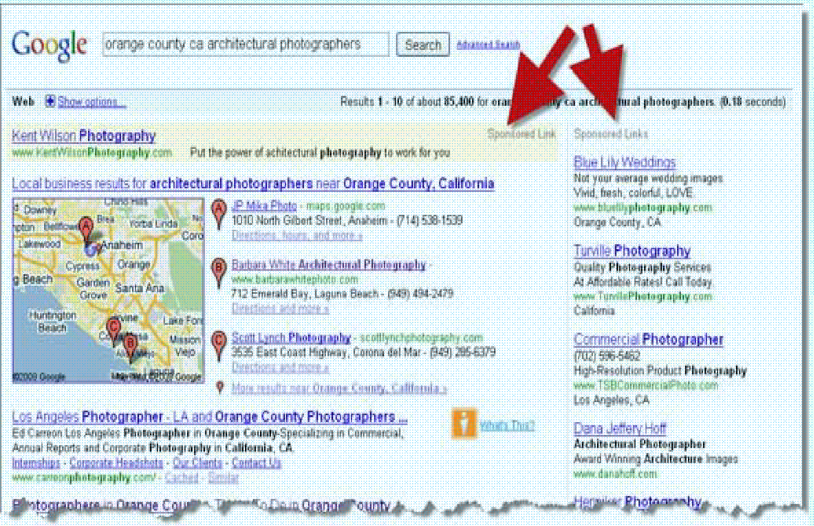

Auctions

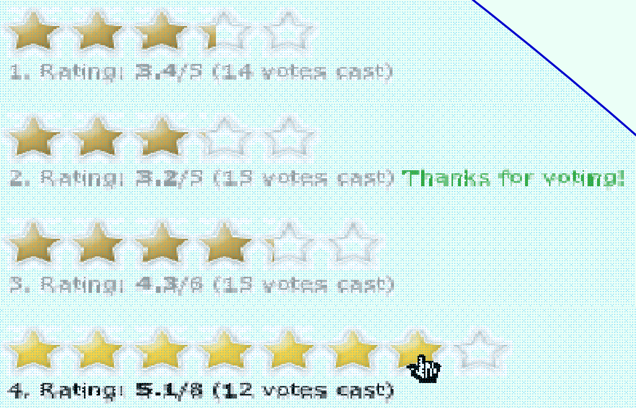

Rating/voting systems

Security

THIS TALK

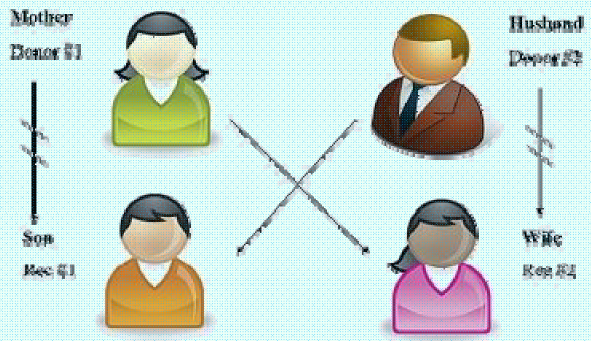

Kidney exchanges

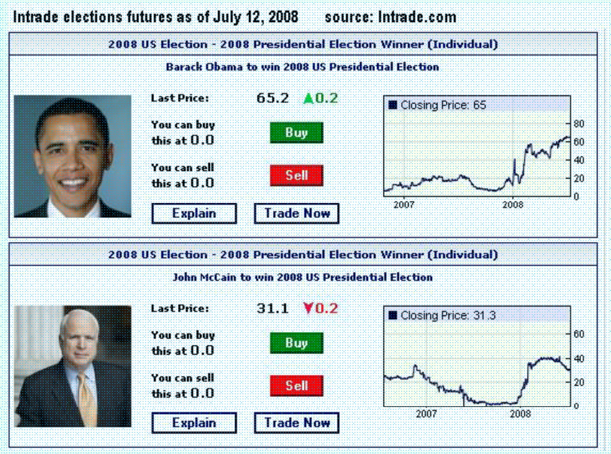

Prediction markets

overview: C., CACM March 2010

Donation matching

How can AI help?

Closer to home…

Game playing

Multiagent systems

Goal:
Blocked(Room0)

Goal:
Clean(Room0)

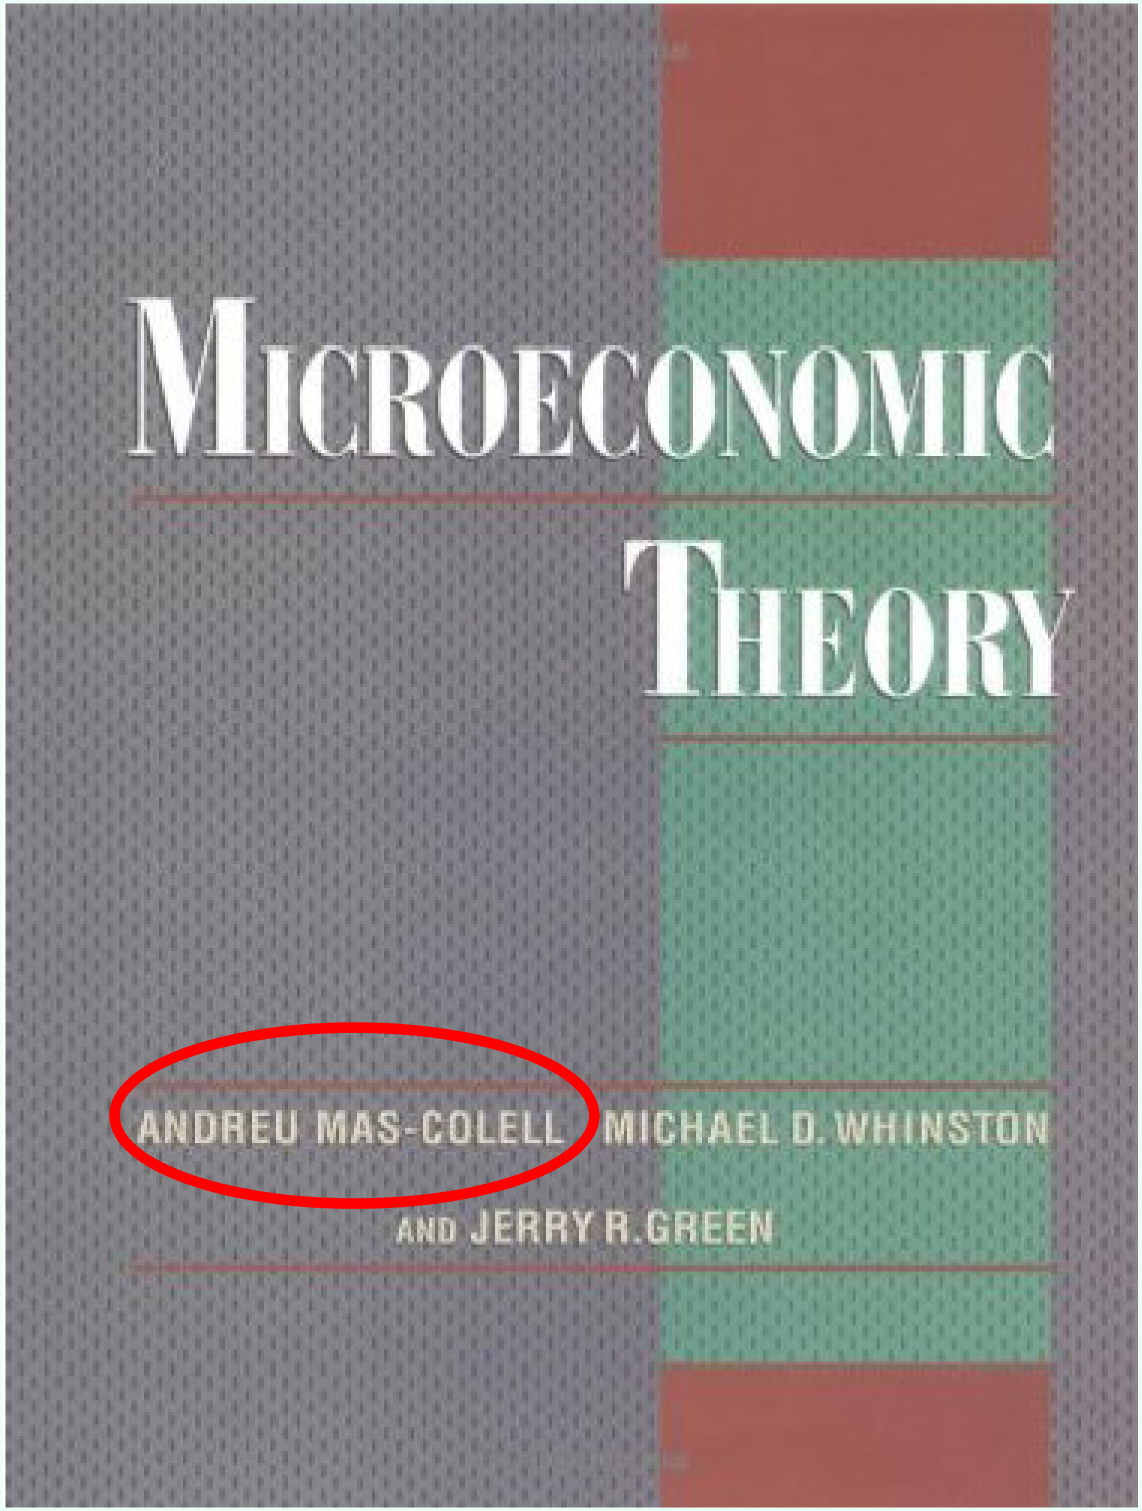
MICROECONOMIC
THEORY
ANDREU MAS-COLELL
MICHAEL D. WHINSTON
AND JERRY R. GREEN

Some microeconomic theory tools for AI

GAME THEORY

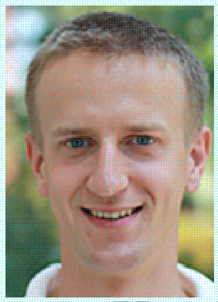

Dima Korzhyk

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

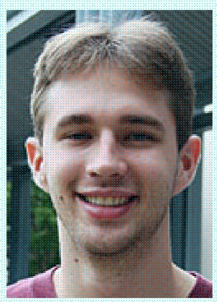

Josh Letchford

THIS TALK

SOCIAL CHOICE

A>B>C
B>A>C
C>B>A
↓
B wins

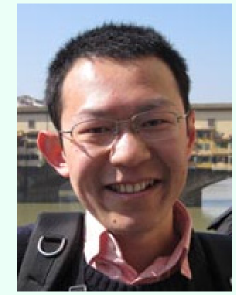

Lirong Xia

MECHANISM DESIGN

v_1=42
v_2=30
v_3=20
→ 1 wins, pays 30

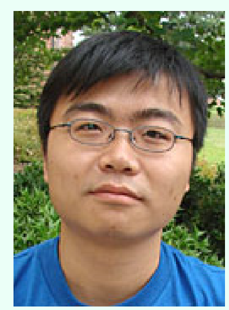

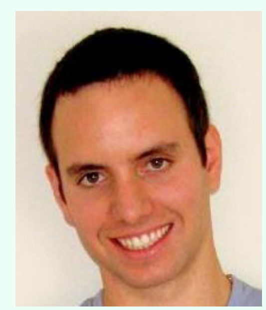

Mingyu Guo *Liad Wagman*

Penalty kick example

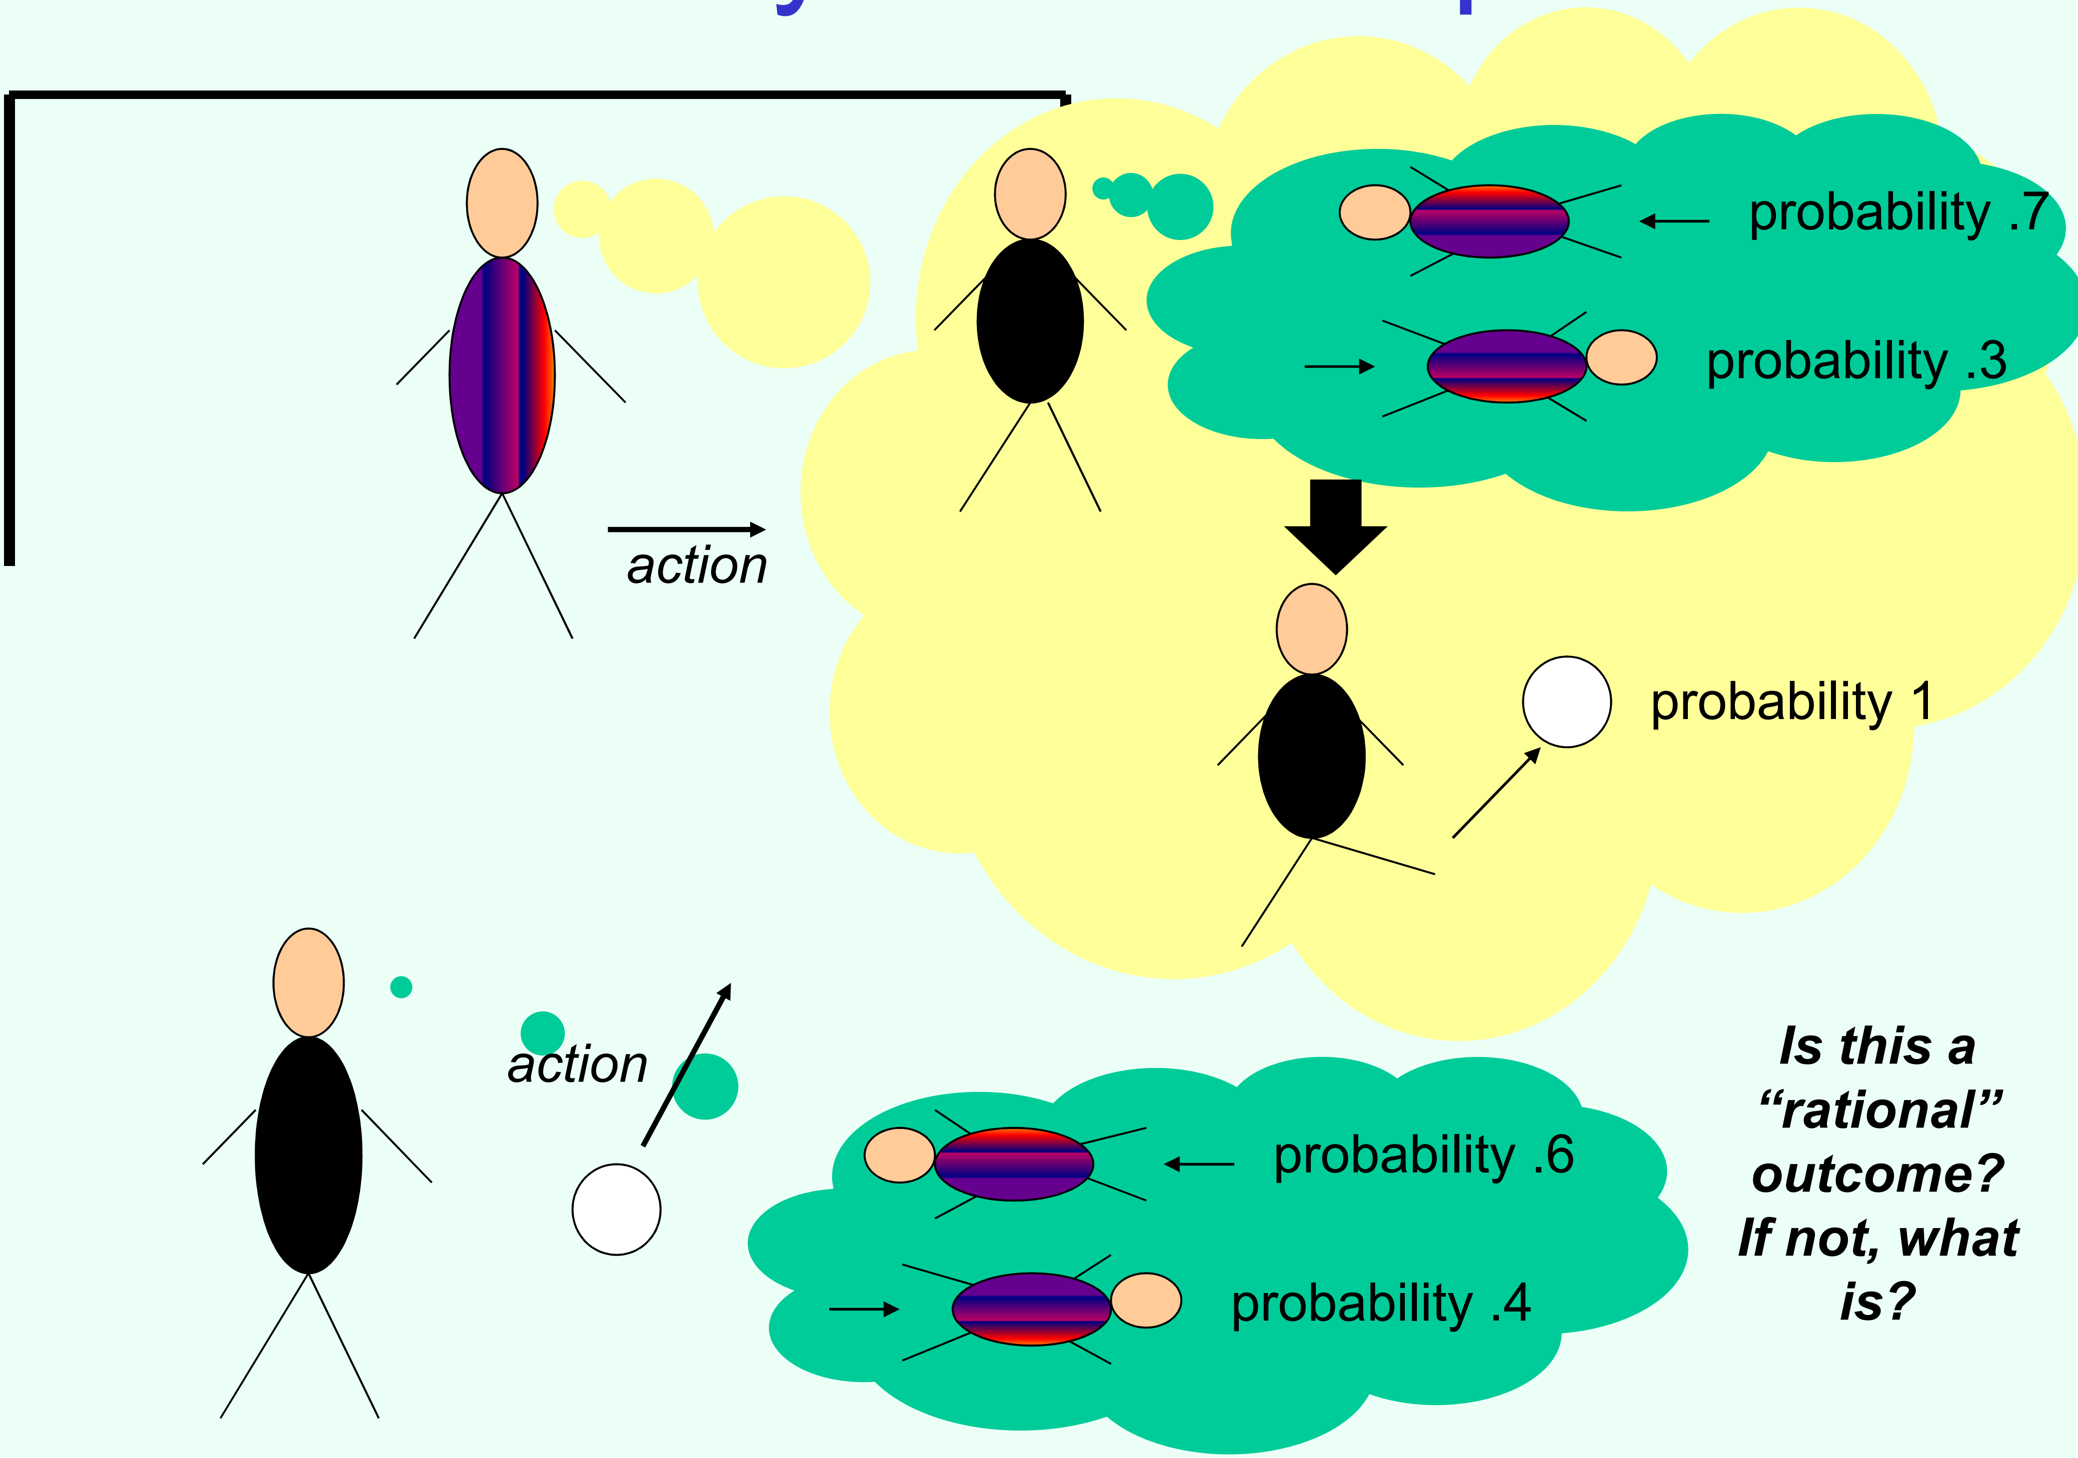

Is this a "rational" outcome? If not, what is?

Penalty kick

(also known as: matching pennies)

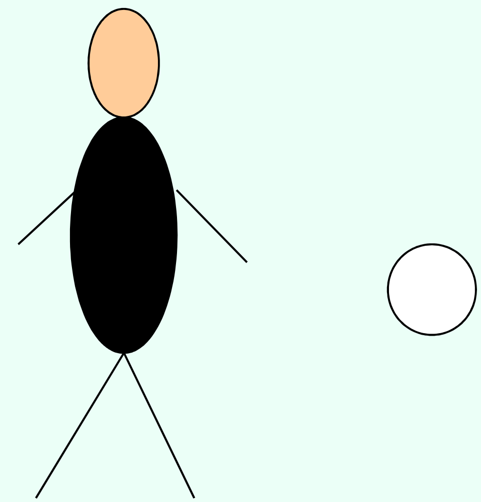

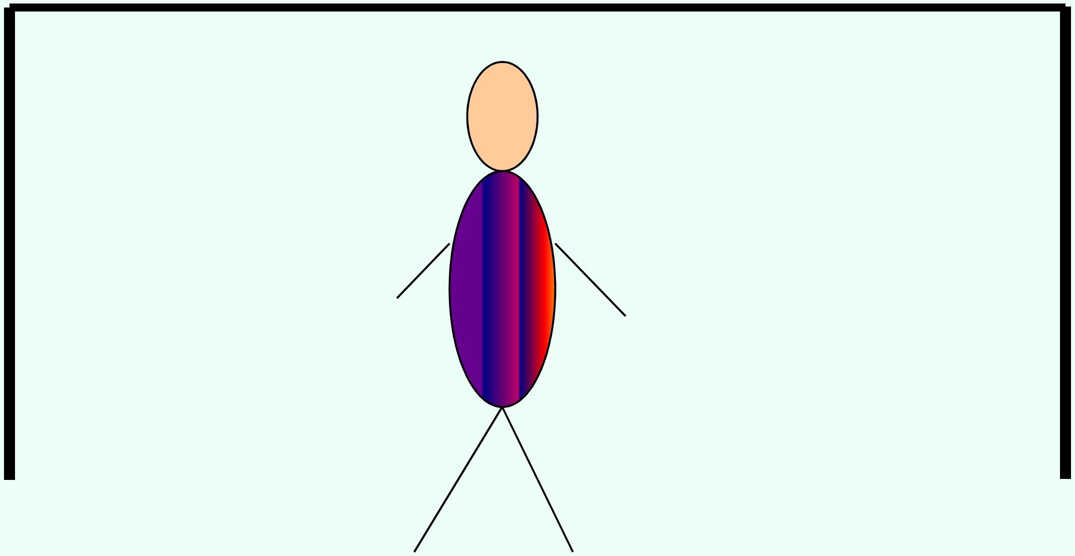

		.5 L	.5 R
.5	L	0, 0	-1, 1
.5	R	-1, 1	0, 0

Security example

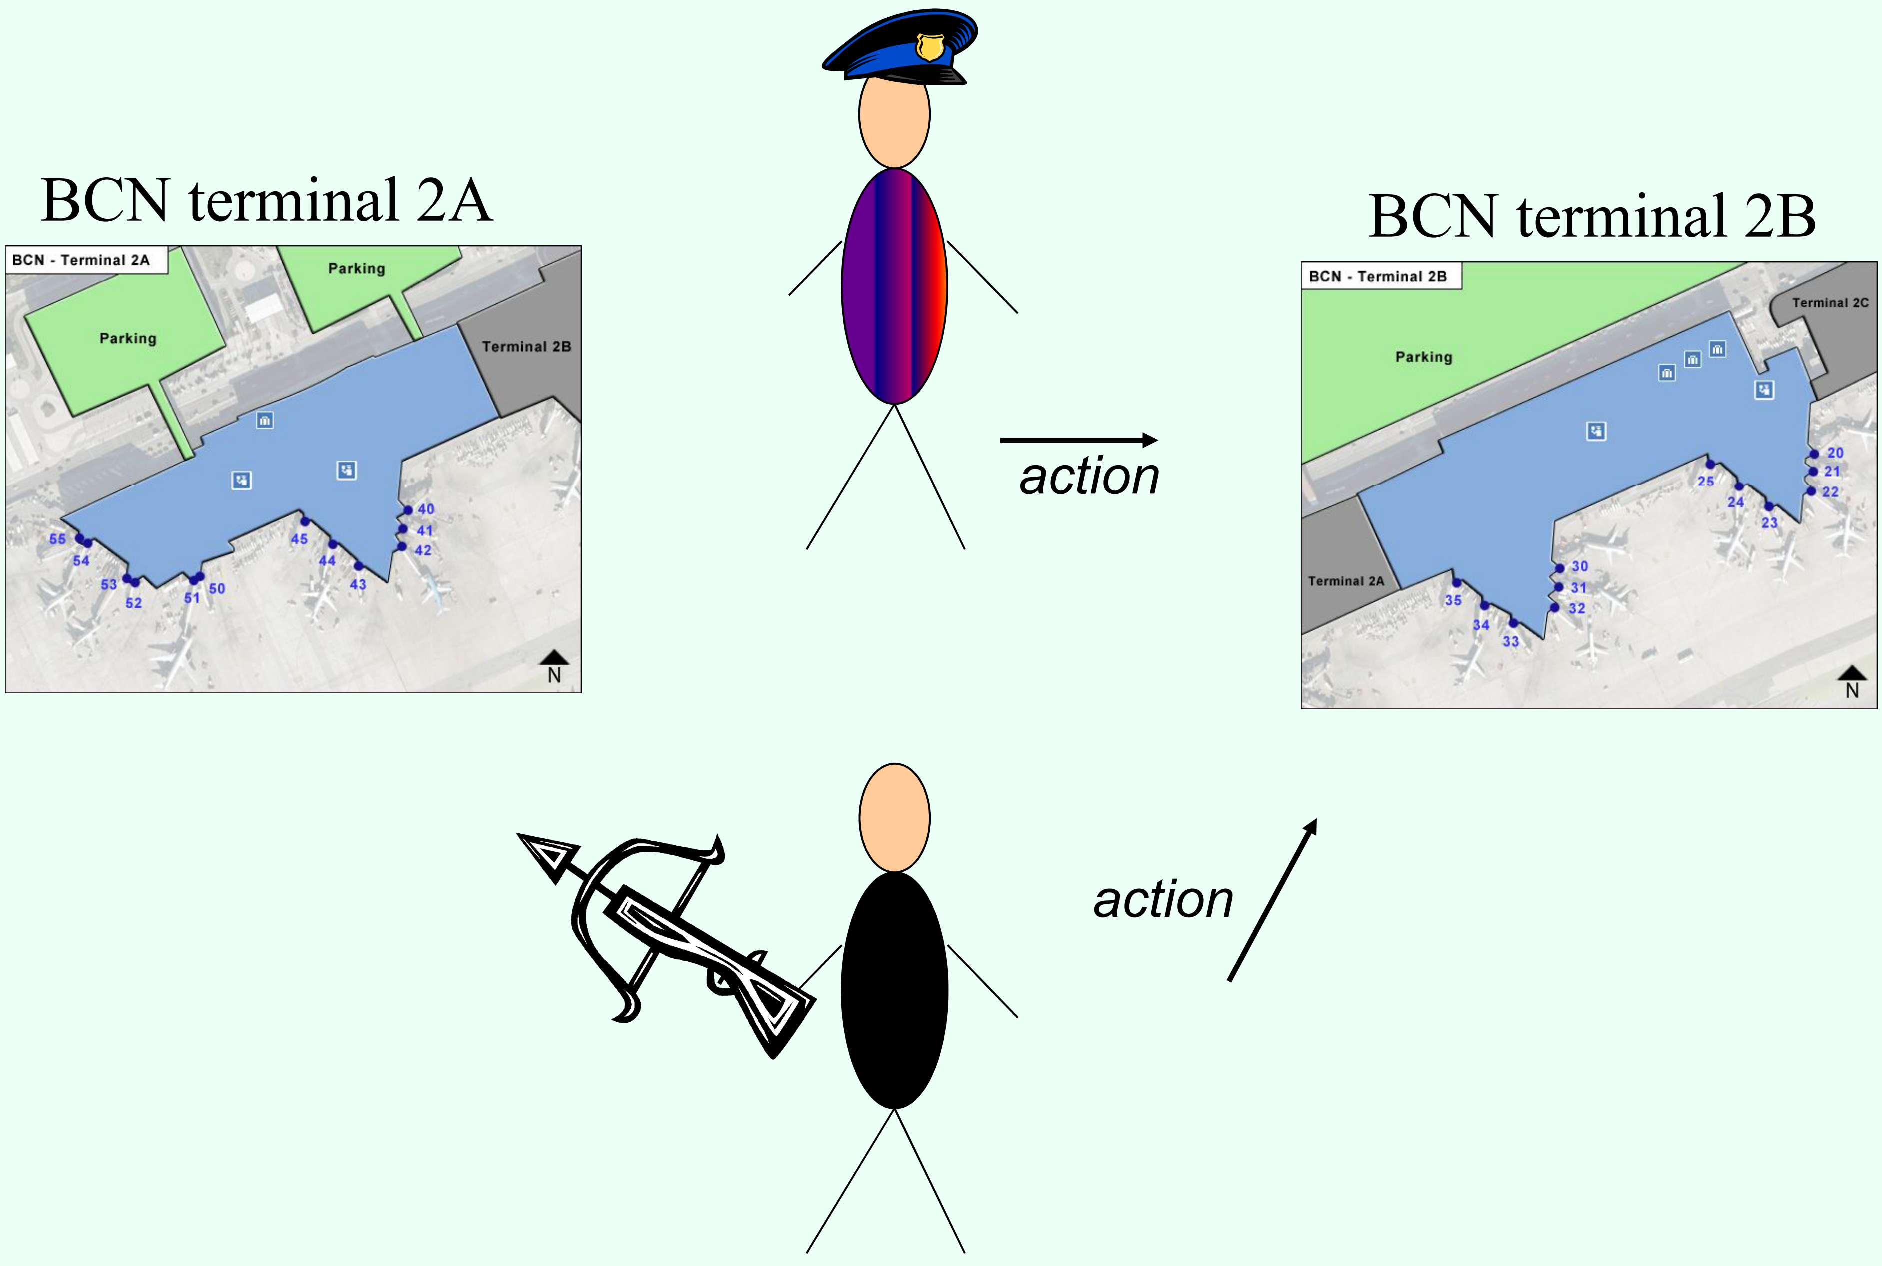

Security game

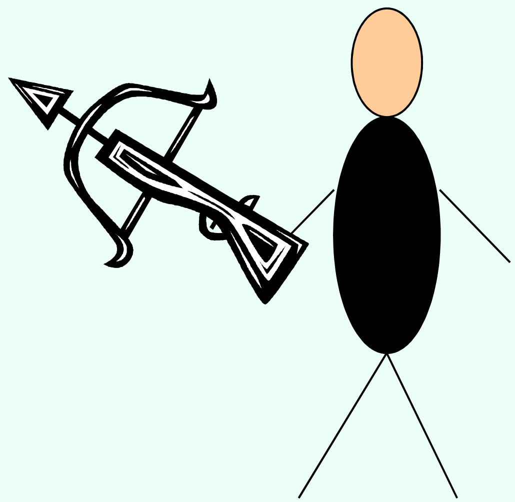

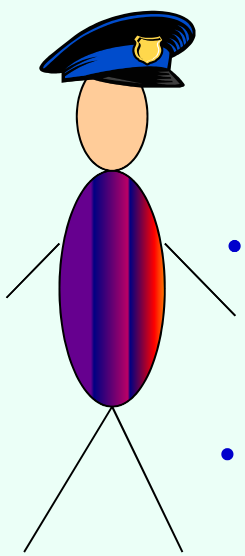

		2A	2B
.33	2A	0, 0	-1, 2
.67	2B	-1, 1	0, 0

Recent deployments in security

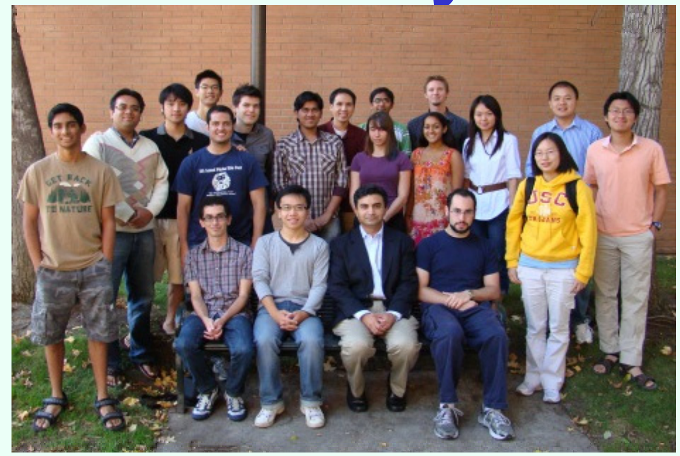

- Tambe's TEAMCORE group at USC
- Airport security
  - Where should checkpoints, canine units, etc. be deployed?
  - Deployed at LAX and another US airport, being evaluated for deployment at all US airports
- Federal Air Marshals
- Coast Guard
- …

“Should I buy an SUV?”

(also known as the Prisoner’s Dilemma)

purchasing + gas cost

cost: 5

cost: 3

accident cost

cost: 5 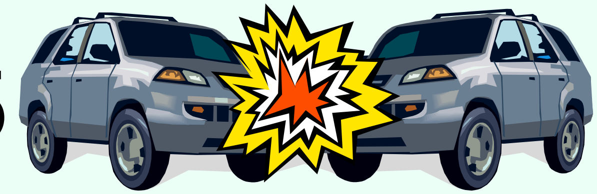 cost: 5

cost: 8  cost: 2

cost: 5  cost: 5

Computational aspects of dominance: Gilboa, Kalai, Zemel Math of OR ‘93; C. & Sandholm EC ’05, AAAI’05; Brandt, Brill, Fischer, Harrenstein TOCS ‘11

SUV	-10, -10	-7, -11
car	-11, -7	-8, -8

"Chicken"

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

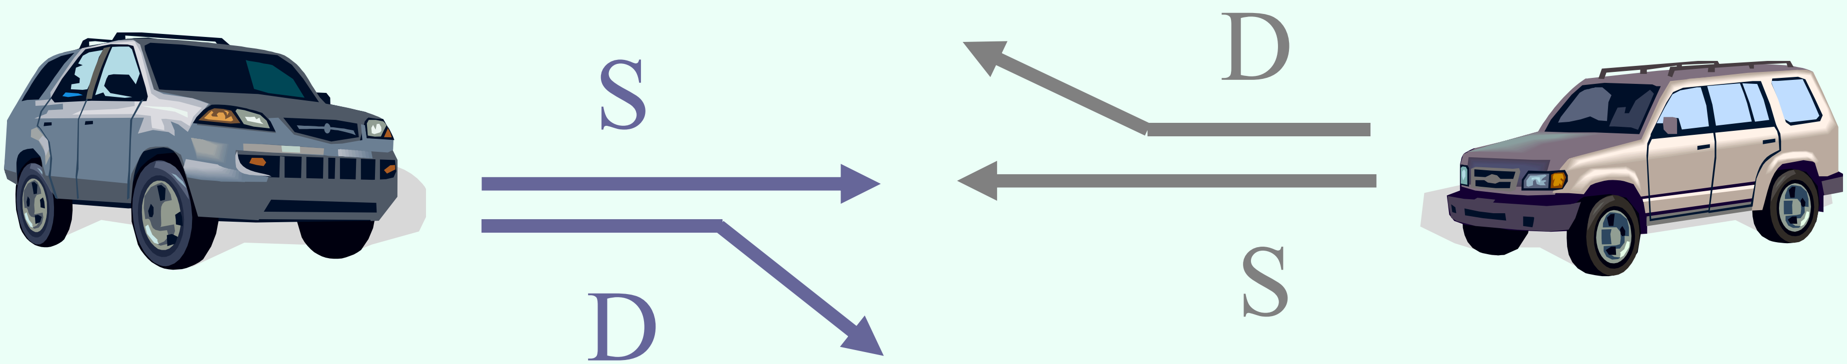

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

Nash equilibrium [Nash '50]

- 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

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

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

  ((4/5 E, 1/5 NE), (1/10 A, 9/10 NA))
  - Utility -7/10 for presenter, 0 for audience

Modeling and representing games

THIS TALK (unless specified otherwise)

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

normal-form games

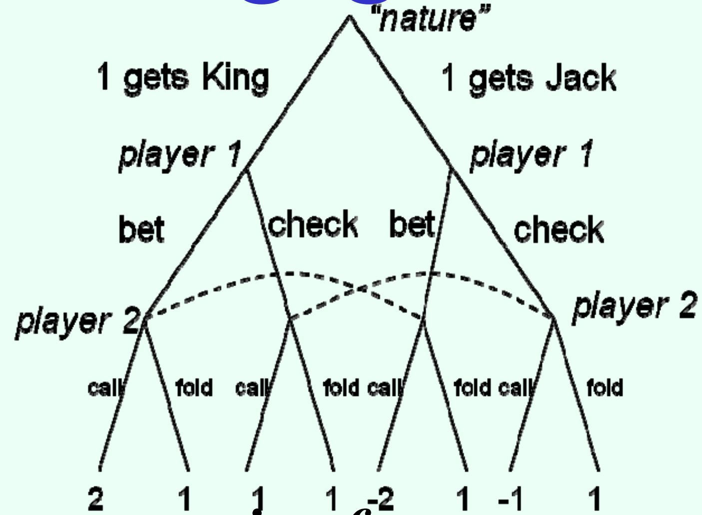

extensive-form games

row player type 1 (prob. 0.5)

	L	R
U	4	6
D	2	4

column player type 1 (prob. 0.5)

	L	R
U	4	6
D	4	6

row player type 2 (prob. 0.5)

	L	R
U	2	4
D	4	2

column player type 2 (prob. 0.5)

	L	R
U	2	2
D	4	2

Bayesian games

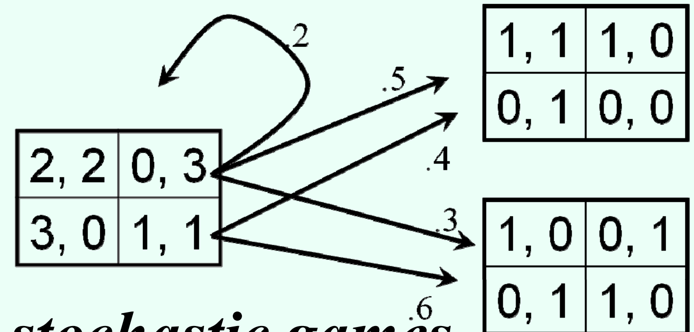

stochastic games

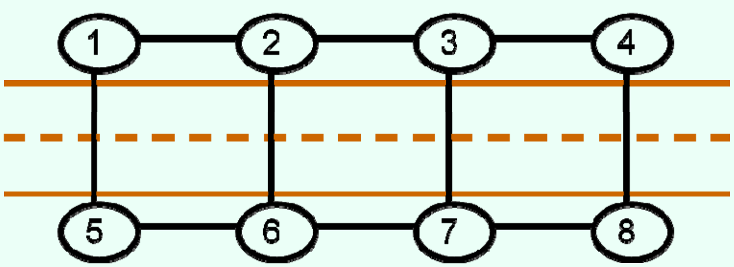

graphical games

[Kearns, Littman, Singh UAI'01]

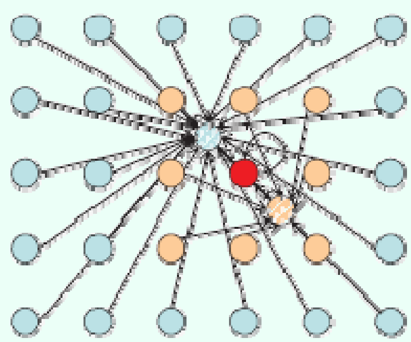

action-graph games

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

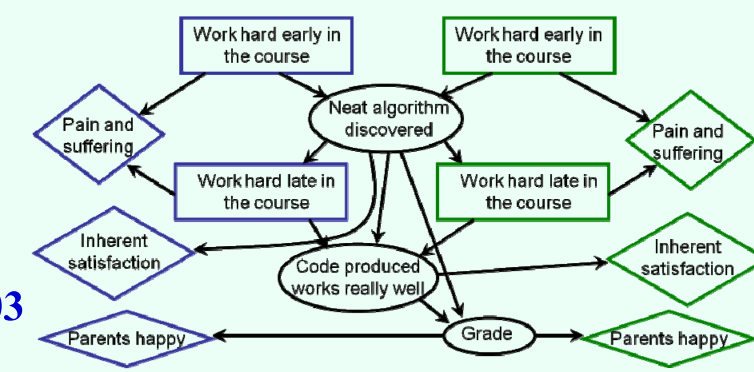

MAIDs

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

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, even in a two-player game [Daskalakis, Goldberg, Papadimitriou STOC'06; Chen & Deng FOCS'06]
  - still holds for FPTAS / smoothed poly [Chen, Deng, Teng FOCS'06]
- Is one Nash equilibrium all we need to know?

A useful reduction (SAT → game)

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

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

Formula: **(x_1 or -x_2) and (-x_1 or x_2)**

Solutions:
x_1=true,x_2=true
x_1=false,x_2=false

Game:

	x_1	**x_2**	**+x_1**	**-x_1**	**+x_2**	**-x_2**	**(x_1 or -x_2)**	**(-x_1 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 or -x_2)	-2,-2	-2,-2	0,-2	2,-2	2,-2	0,-2	-2,-2	-2,-2	0,1
(-x_1 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

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

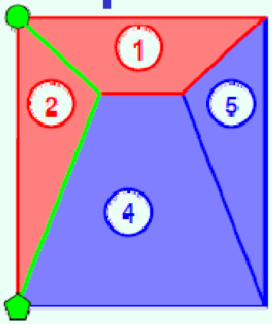

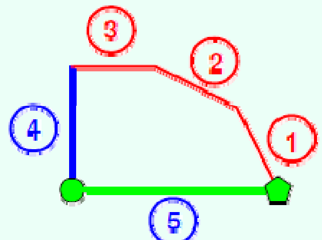

image from von Stengel

Lemke-Howson [J. SIAM '64]

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

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

Search over supports / MIP

[Dickhaut & Kaplan, Mathematica J. '91]
[Porter, Nudelman, Shoham AAAI'04 / GEB'08]
[Sandholm, Gilpin, C. AAAI'05]

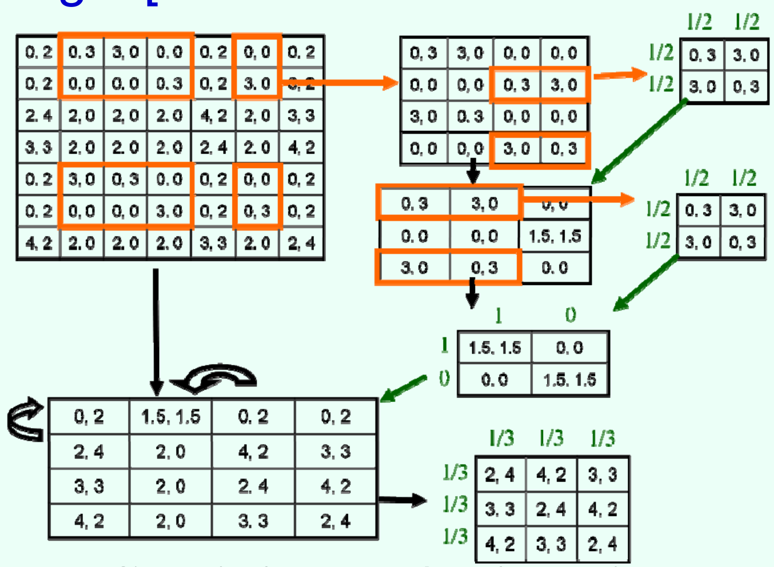

Special cases / subroutines

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

	0, 1	0, 1	½, ½	½, ½	½, ½
	1, 0	1, 0	0, 1	0, 1	0, 1
	1, 0	1, 0	0, 1	0, 1	0, 1
	½, ½	½, ½	1, 0	1, 0	1, 0
	½, ½	½, ½	1, 0	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, …]

Sidestepping the problems

(one *solution concept* is not enough…?)

Nash is not optimal if one player can commit

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

Unique Nash equilibrium

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 to mixed strategies

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

– Sometimes also called a Stackelberg (mixed) strategy

Observing the defender's distribution in security

BCN terminal 2A

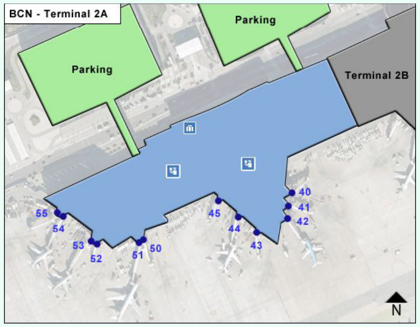

BCN terminal 2B

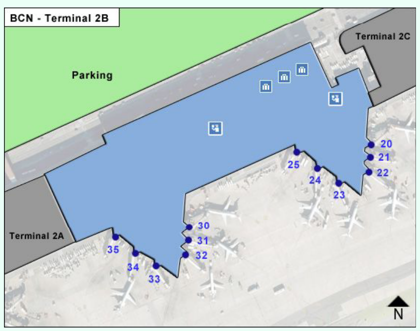

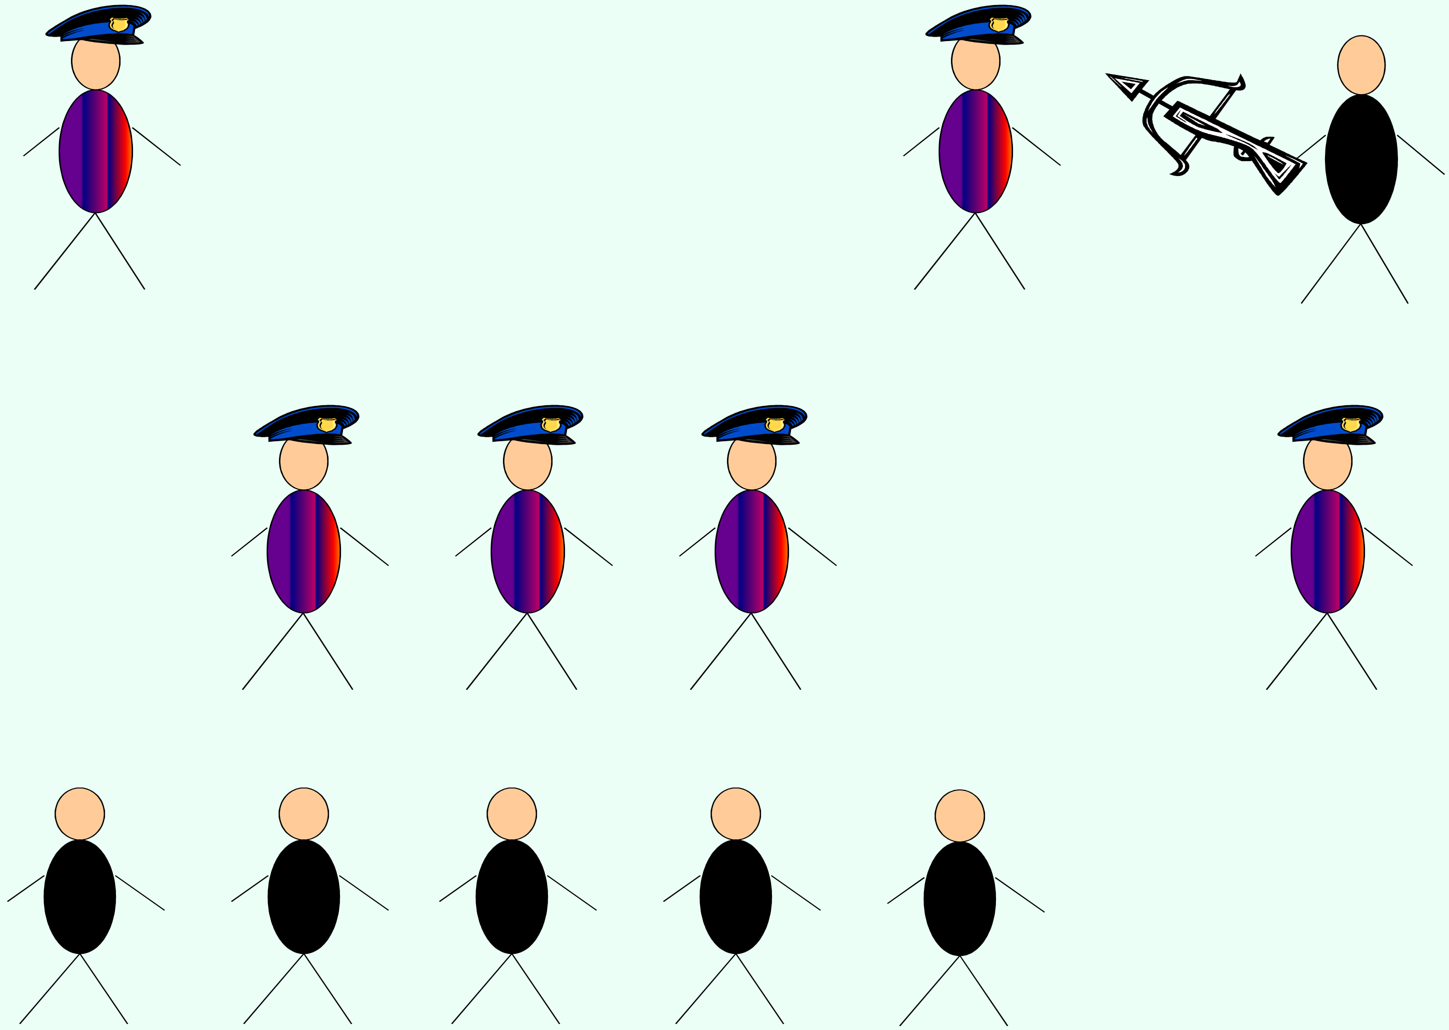

observe

Mo Tu We Th Fr Sa

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

Computing the optimal mixed strategy to commit to

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

- Separate LP for every column c^*:

maximize $\Sigma_r \, p_r \, u_R(r, c^*)$ leader utility

subject to

for all c, $\Sigma_r \, p_r \, u_C(r, c) \leq \Sigma_r \, p_r \, u_C(r, c^*)$ follower optimality

$\Sigma_r \, p_r = 1$ distributional constraint

Other nice properties of commitment to mixed strategies

- Agrees w. Nash in zero-sum games

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

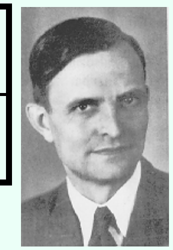

- Leader's payoff at least as good as any Nash eq. or even correlated eq. (von Stengel & Zamir [GEB '10]; see also C. & Korzhyk [AAAI '11], Letchford & C. [draft])

$\geq$

- No equilibrium selection problem

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

Some other work on commitment in unrestricted games

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

normal-form games

learning to commit [Letchford, C., Munagala SAGT'09]

uncertain observability [Korzhyk, C., Parr AAMAS'11]

correlated strategies [C. & Korzhyk, AAAI'11]

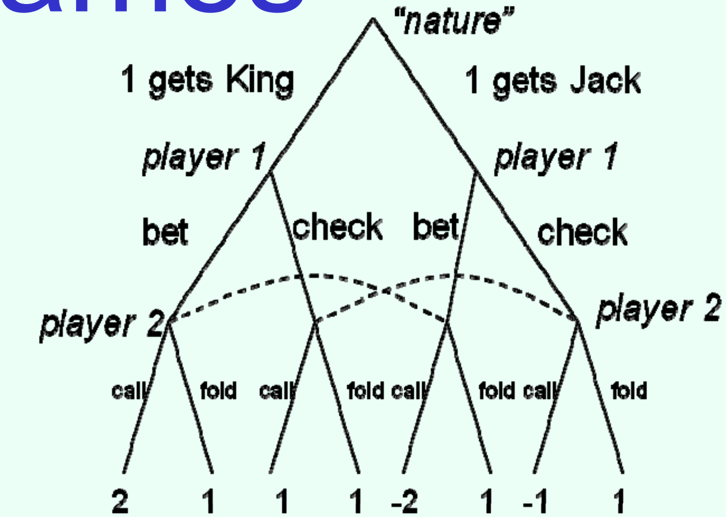

extensive-form games

[Letchford & C., EC'10]

row player type 1 (prob. 0.5)

	L	R
U	4	6
D	2	4

column player type 1 (prob. 0.5)

	L	R
U	4	6
D	4	6

row player type 2 (prob. 0.5)

	L	R
U	2	4
D	4	2

column player type 2 (prob. 0.5)

	L	R
U	2	2
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]

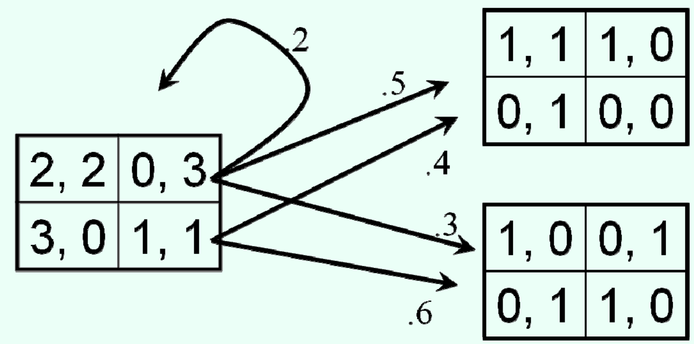

stochastic games

ongoing work with Korzhyk, Letchford, Parr

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 $S \subseteq 2^T$
- 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
- $U_d^c(t) \geq U_d^u(t); U_a^c(t) \leq U_a^u(t)$

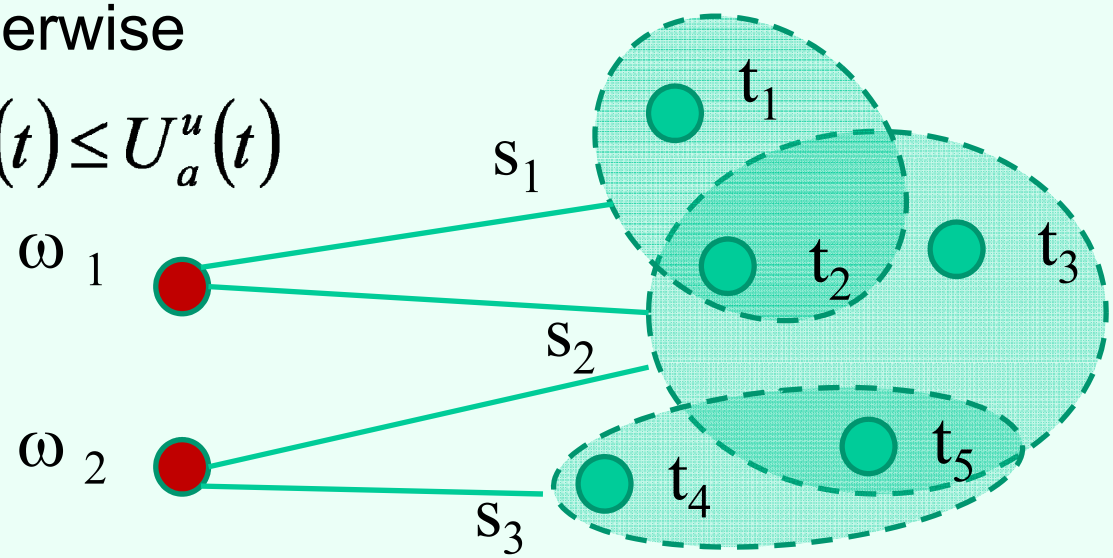

Game-theoretic properties of security resource allocation games [Korzhyk, Yin, Kiekintveld, C., Tambe JAIR'11]

- For the defender: Stackelberg strategies are also Nash strategies
  – minor assumption needed
  – not true with multiple attacks

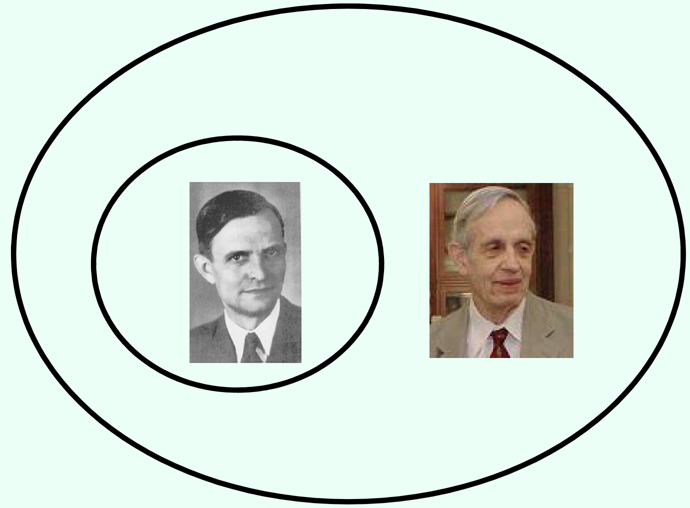

- Interchangeability property for Nash equilibria ("solvable")
  - no equilibrium selection problem
  - still true with multiple attacks [Korzhyk, C., Parr IJCAI'11 – poster W. 3:30pm, talk F. 10:30am]

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

Scalability in security games

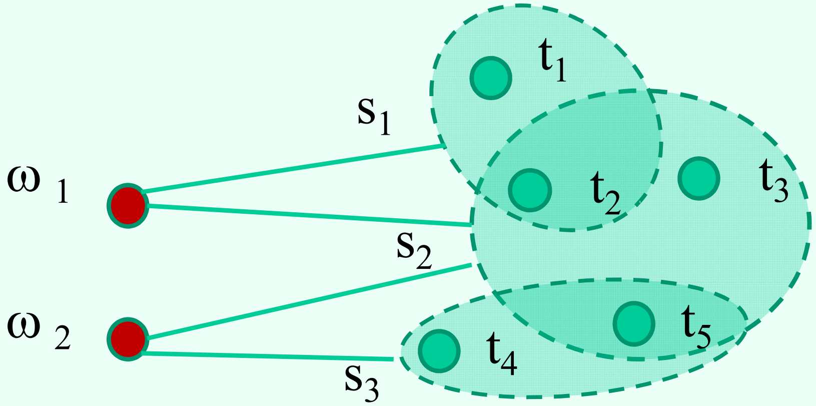

basic model

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

→

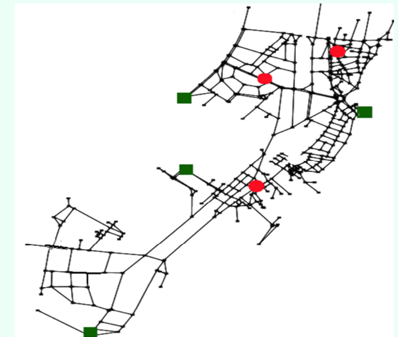

games on graphs (usually zero-sum)

[Halvorson, C., Parr IJCAI’09; Tsai, Yin, Kwak, Kempe, Kiekintveld, Tambe AAAI’10; Jain, Korzhyk, Vaněk, C., Pěchouček, Tambe AAMAS’11]; ongoing work with Letchford, Vorobeychik

Techniques:

compact linear/integer programs

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

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

strategy generation

$$\begin{array}{ll}
\min & u \\
\text{subject to} & \\
& \sigma_h(s_{h_0}) + \dots + \sigma_h(s_{h_k}) = 1 \\
& u \ge \sigma_h(s_{h_0}) \cdot u(s_{s_0}, s_{h_0}) + \dots + \sigma_h(s_{h_k}) \cdot u(s_{s_0}, s_{h_k}) \\
& u \ge \sigma_h(s_{h_0}) \cdot u(s_{s_1}, s_{h_0}) + \dots + \sigma_h(s_{h_k}) \cdot u(s_{s_1}, s_{h_k}) \\
& \vdots \qquad \ddots \qquad \vdots \\
& u \ge \sigma_h(s_{h_0}) \cdot u(s_{s_k}, s_{h_0}) + \dots + \sigma_h(s_{h_k}) \cdot u(s_{s_k}, s_{h_k})
\end{array}$$

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

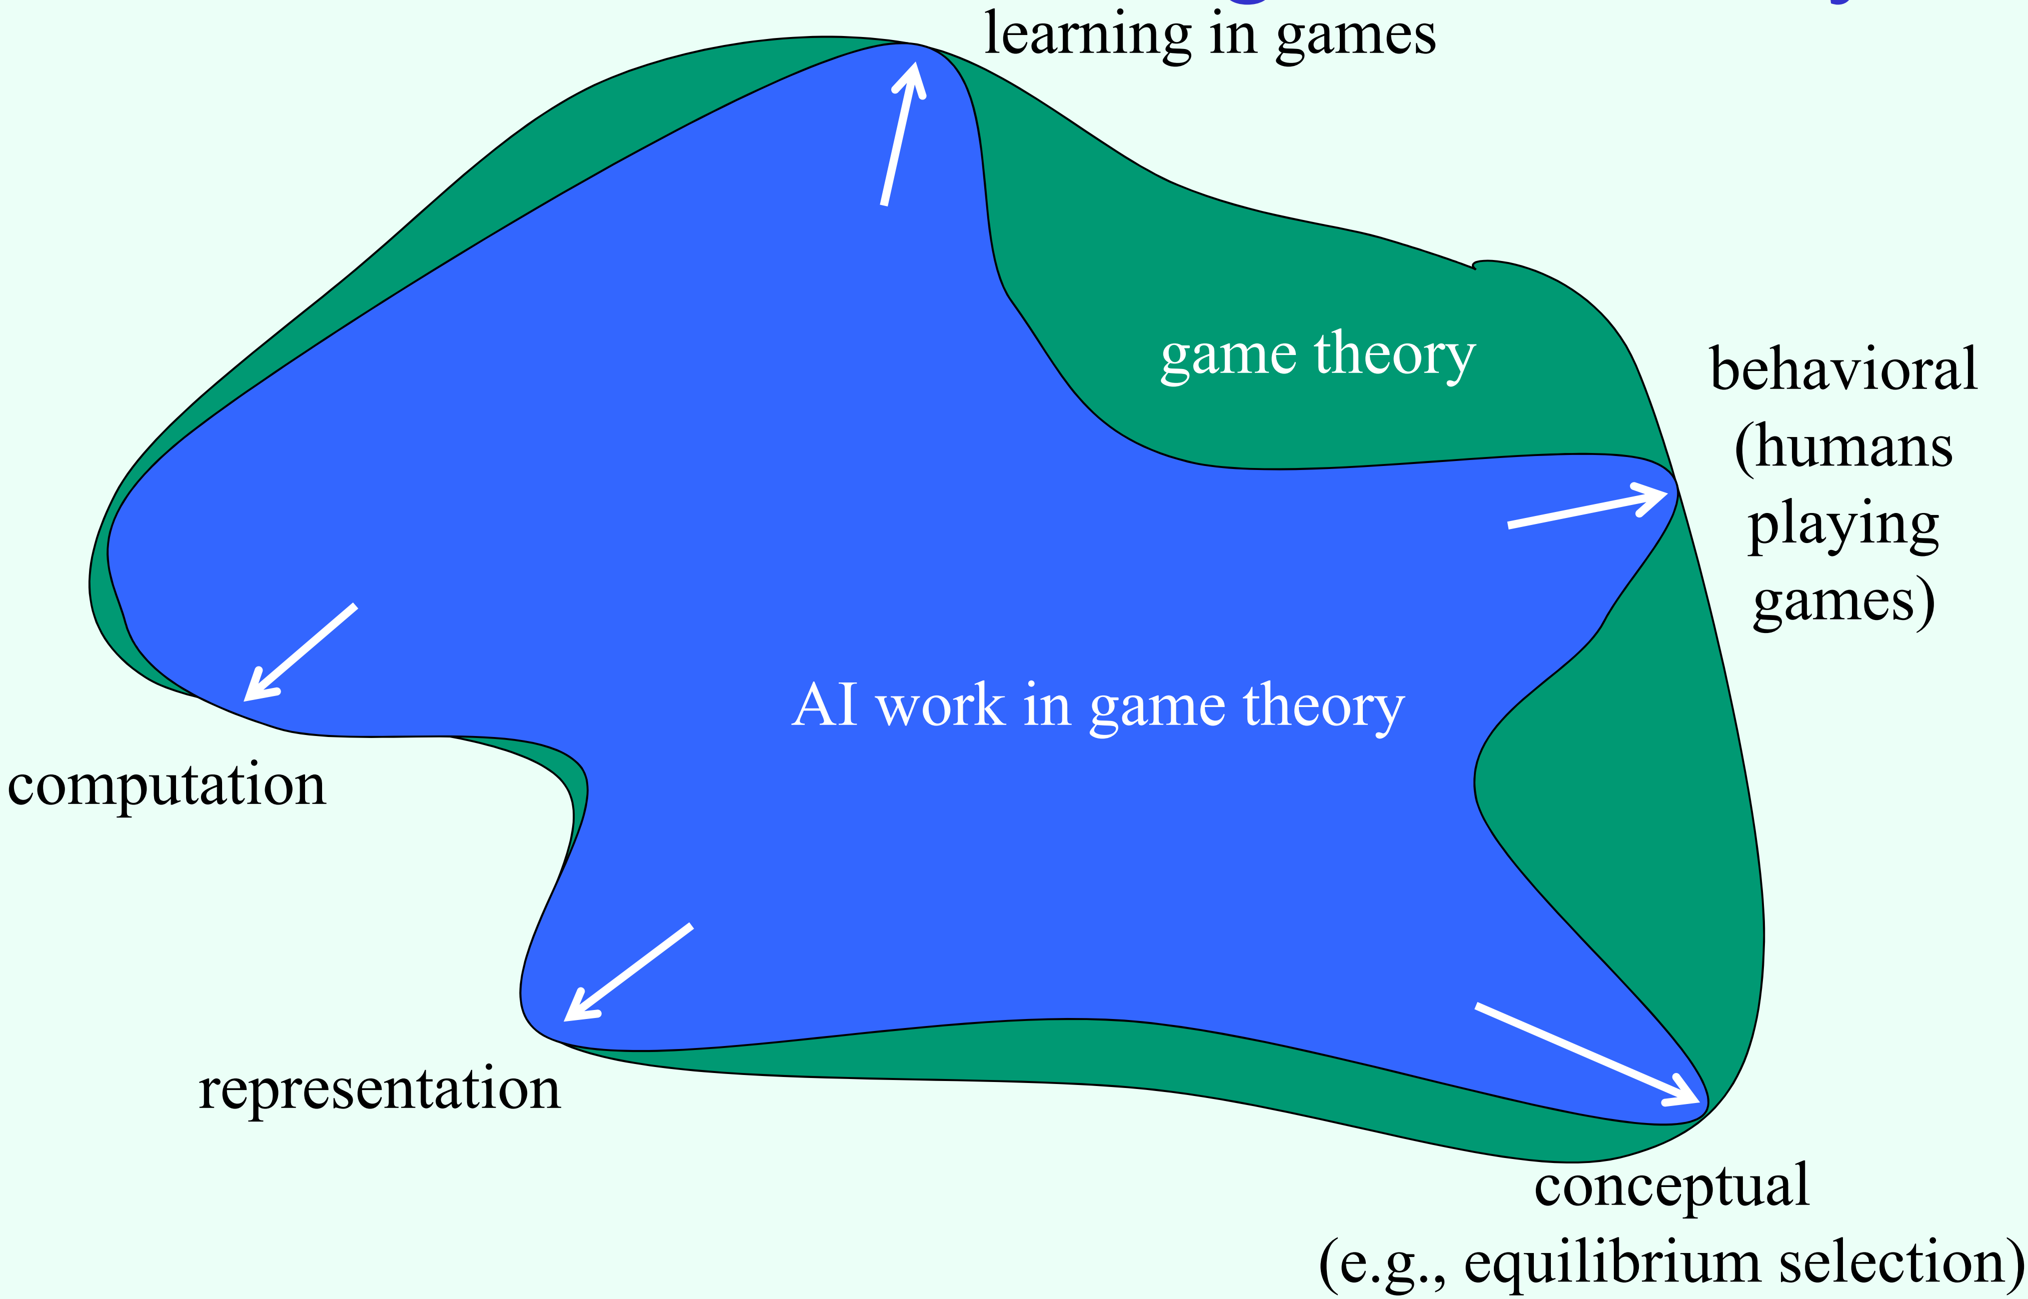

Funding

Any opinions, conclusions or recommendations are mine and do not necessarily reflect the views of the funding agencies

Academic family

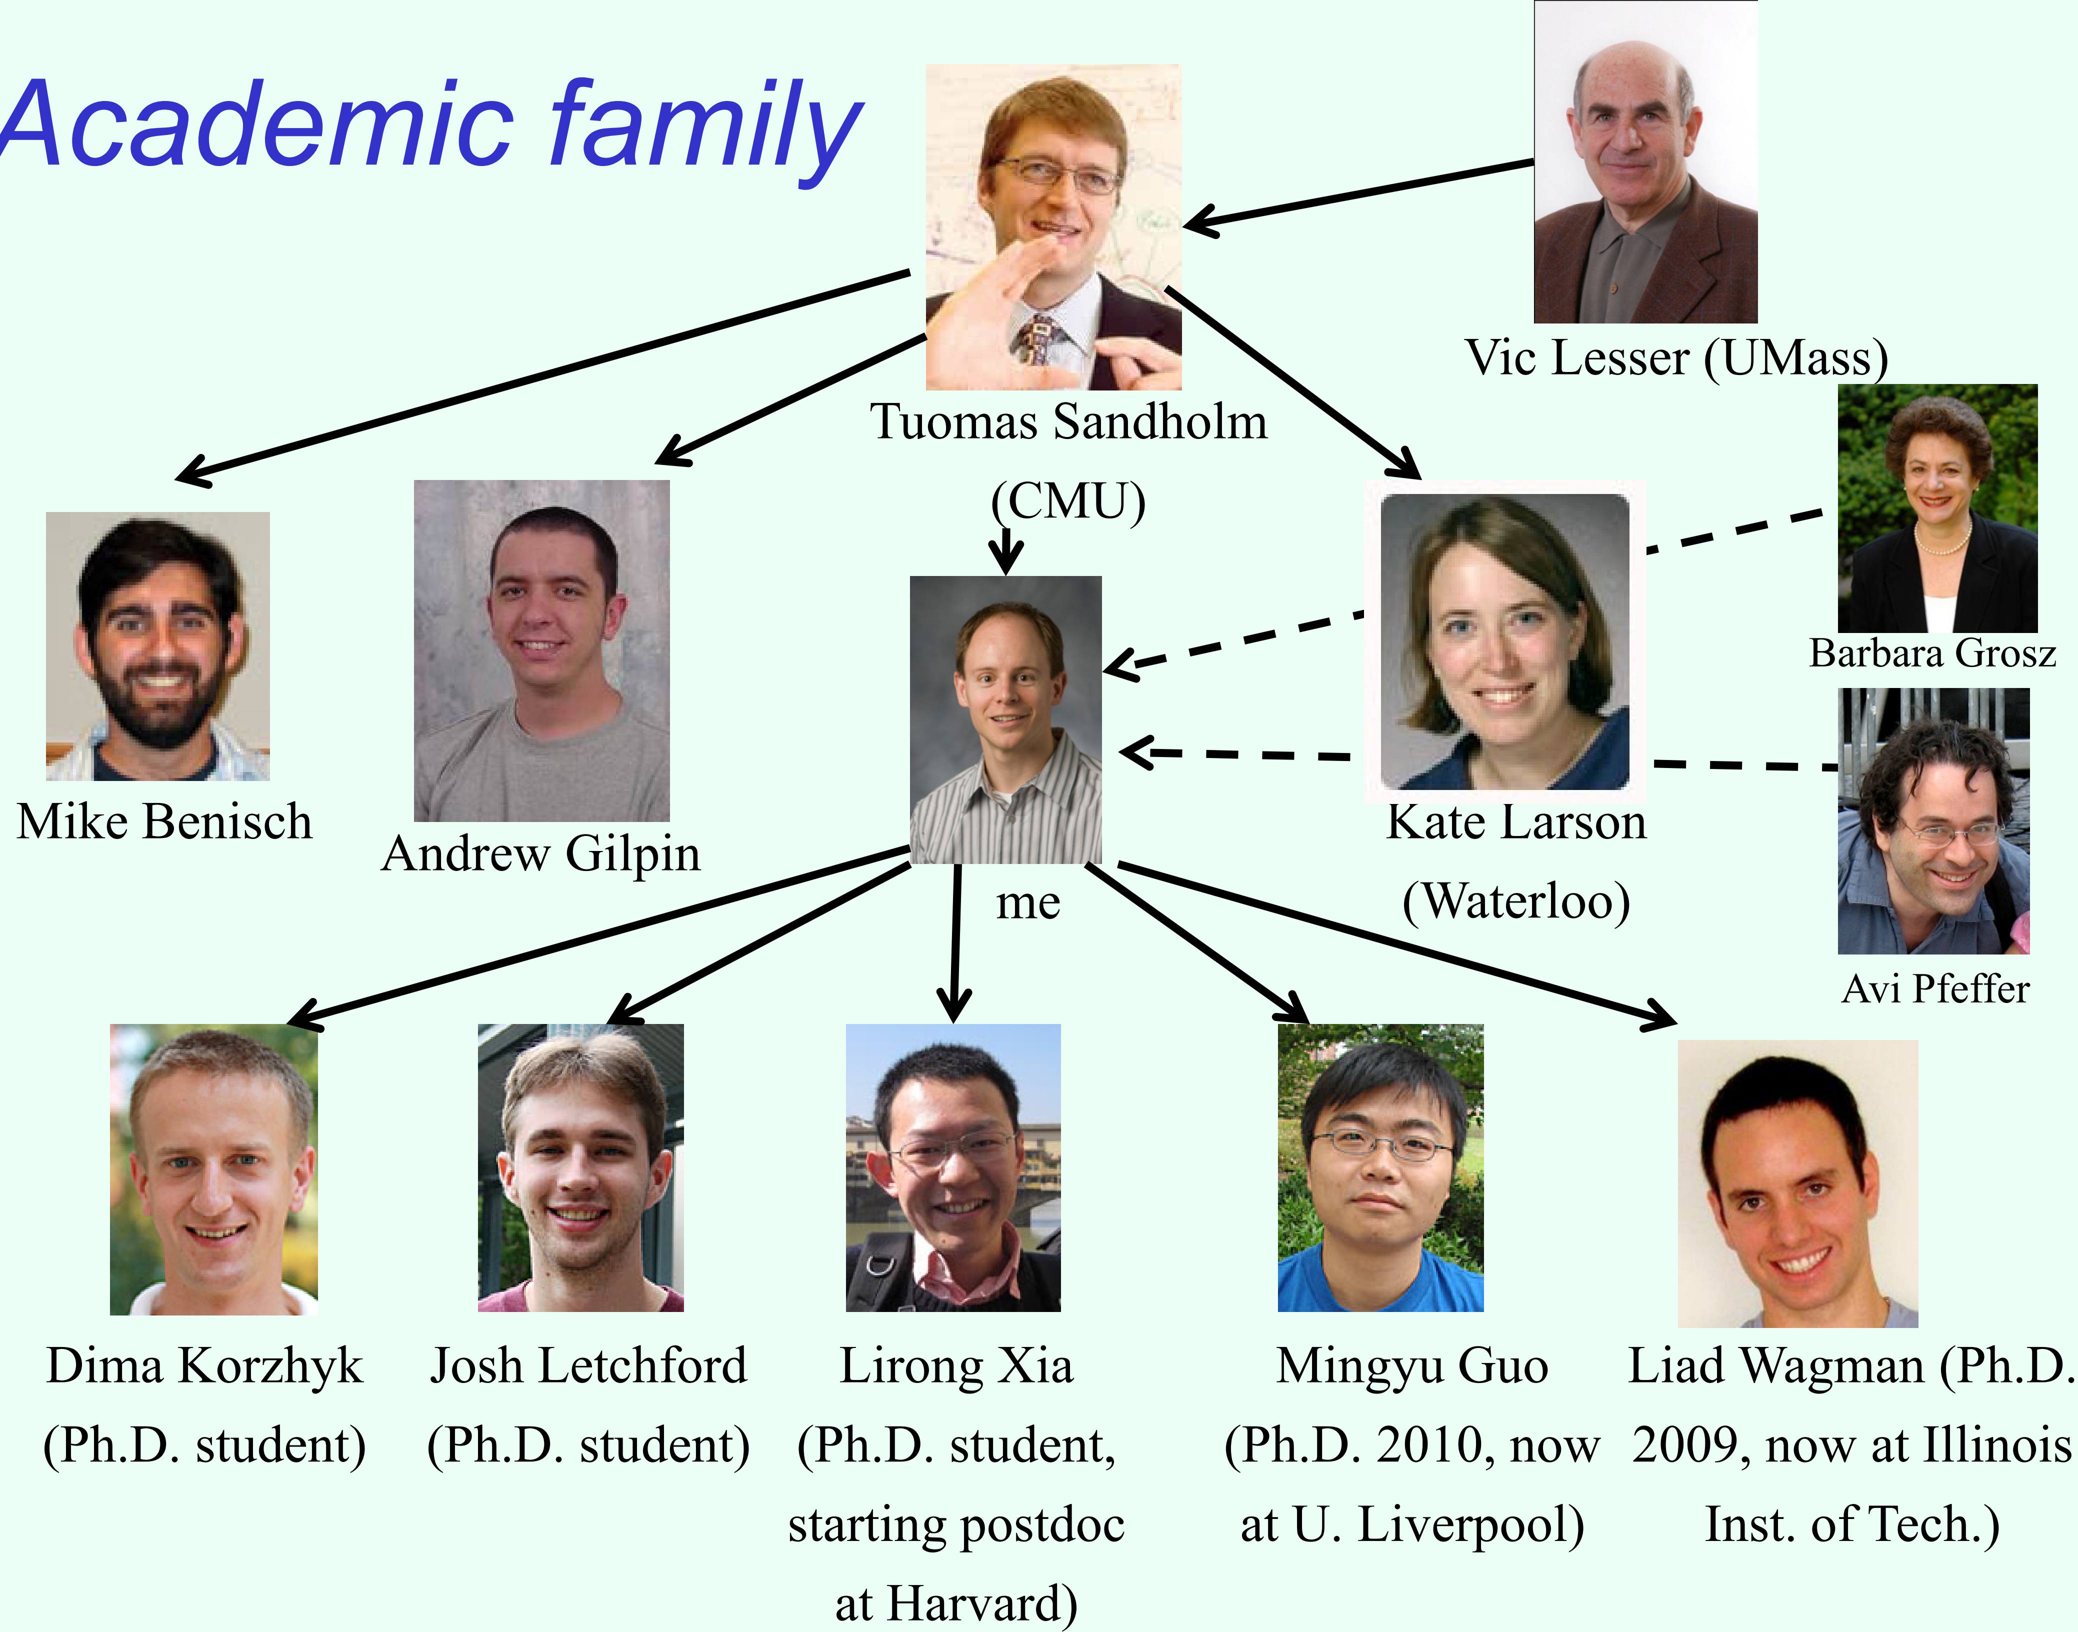

AI at Duke


Ron Parr


Carlo Tomasi


Bruce Donald


Alex Hartemink


me

SECONDARY CS FACULTY


Silvia Ferrari


Sayan Mukherjee


Uwe Ohler


Mauro Maggioni

C&T book(s?)

office!

D207


Don Loveland


Alan Biermann

PROFESSORES EMERITI

Family