CS325 Artificial Intelligence
Ch. 17.5–6, Game Theory

Cengiz Günay, Emory Univ.

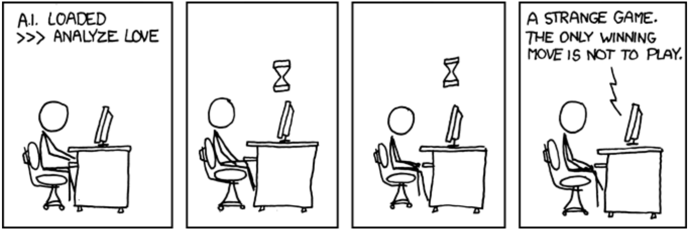

Spring 2013

State of the art subjects: build order planning, over state estimation, plan recognition. . .

Article on 2010 winner: Berkeley Overmind bot

MDPs and RL for games: Civilization

High-level Reinforcement Learning in Strategy Games

Christopher Amato *
Department of Computer Science
University of Massachusetts
Amherst, MA 01003 USA
camato@cs.umass.edu

Guy Shani *
Department of Computer Science
Ben-Gurion University
Beer-Sheva 84105 Israel
shanigu@bgu.ac.il

2010 Paper on playing Civilization IV; uses:

- Markov Decision Processes
- Reinforcement Learning, a model-based Q-learning approach

Compares strategies and parameters on winning outcomes.

Figure 3: Leaders in Civilization IV (clockwise from top left): Frederick II, Mahatma Gandhi, Genghis Kahn and George Washington

And Now, Game Theory

Game theory applies when:

- partially-observable, or
- with simultaneous moves (e.g., StarCraft).

And Now, Game Theory

Game theory applies when:

- partially-observable, or
- with simultaneous moves (e.g., StarCraft).

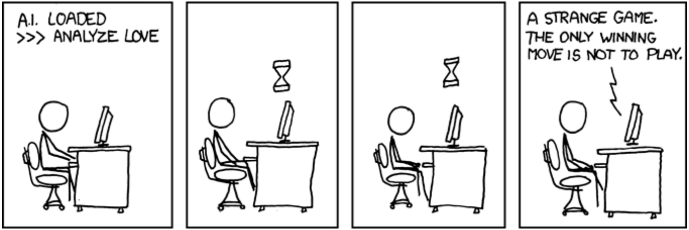

And Now, Game Theory

Game theory applies when:

- partially-observable, or
- with simultaneous moves (e.g., StarCraft).

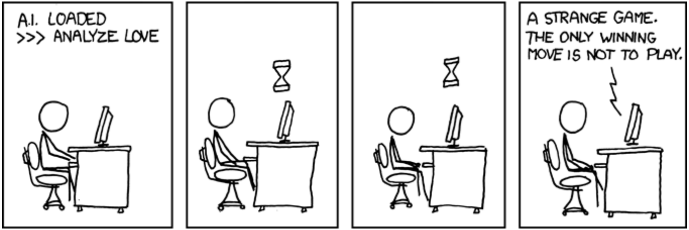

Game theory deals more with cases like:

- Diplomacy/war between enemies
- Bidding
- Creating win-win scenarios

Entry/Exit Surveys

Exit survey: Adversarial Games

- How do you reduce the tree search complexity of a turn-by-turn game like chess?
- Give an example for a game that we haven't studied in class which can be solved with the minimax algorithm. Suggest an evaluation function at the cutoff nodes.

Entry survey: Game Theory (0.25 points of final grade)

- Can we use minimax tree search work in simultaneous moves? Briefly explain why or why not?
- Think that you will have to make the move of the US side in a Cold War scenario. How would you consider the opponent's move, uncertainty, and secrecy?

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B?

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B? **Testifying** is dominant.

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B? **Testifying** is dominant.

Pareto optimal: If no better solution for both players exist.
Which condition?

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B? **Testifying** is dominant.

Pareto optimal: If no better solution for both players exist.
Which condition? Three of them.

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B? **Testifying** is dominant.

Pareto optimal: If no better solution for both players exist.
Which condition? Three of them.

Nash equilibrium: Local minima; single player switch does not improve.
Is there one?

Terminology: 2 Prisoners Dilemma

More like in Law and Order or The Closer:

- 2 suspects in separate interrogation rooms.

Each can either:

1. Testify against the other, or
2. Refuse to speak.

	A: testify	A: refuse
B: testify	$A = -5, B = -5$	$A = -10, B = 0$
B: refuse	$A = 0, B = -10$	$A = -1, B = -1$

Dominant strategy: Selfish decision that is always better.
For A and B? **Testifying** is dominant.

Pareto optimal: If no better solution for both players exist.
Which condition? Three of them.

Nash equilibrium: Local minima; single player switch does not improve.
Is there one? **Testifying**, again.

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.

For A and B?

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.

For A and B? **None!**

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.
For A and B? **None!**

Pareto optimal: If no better solution for both players exist.
Which condition?

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.
For A and B? **None!**

Pareto optimal: If no better solution for both players exist.
Which condition? **Only one**.

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.
For A and B? **None!**

Pareto optimal: If no better solution for both players exist.
Which condition? **Only one**.

Equilibrium: Local minima; single player switch does not improve.
Is there one?

Terminology (2): Game Console Game

Console producer (A) vs. game developer (B), need to decide between:

- Blu-ray vs. DVD

	A: bluray	A: dvd
B: bluray	$A = +9, B = +9$	$A = -4, B = -1$
B: dvd	$A = -3, B = -1$	$A = +5, B = +5$

Dominant strategy: Selfish decision that is always better.
For A and B? **None!**

Pareto optimal: If no better solution for both players exist.
Which condition? **Only one**.

Equilibrium: Local minima; single player switch does not improve.
Is there one? **Two cases**.

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better.

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better. **None!**

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better. **None!**

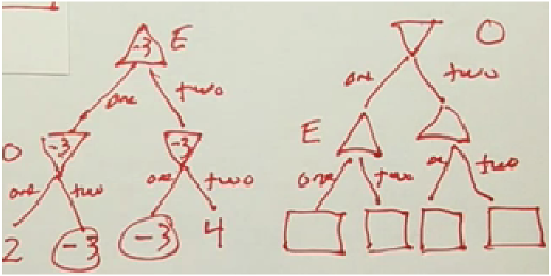

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better. **None!**

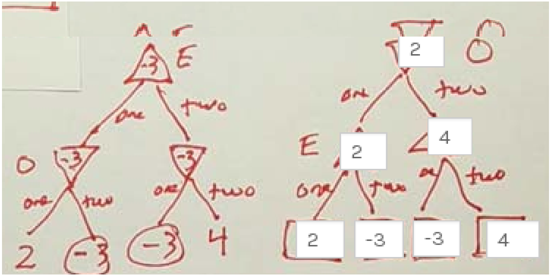

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better. **None!**

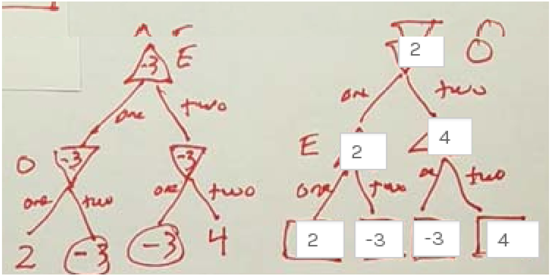

- Utility of E: $-3 \leq U_E \leq 2$

Strategies: 2 Finger Morra Game

Difficult, zero-sum betting game:

1. Show a number of fingers
2. Player betting on odd (O) or even (E) wins based on total fingers

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Dominant strategy: Selfish decision that is always better. **None!**

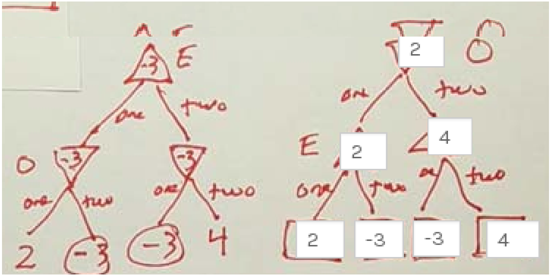

- Utility of E:
  $-3 \leq U_E \leq 2$
- Not very sure? Handicapped?
- Use **mixed strategy**

Mixed Strategy: 2 Finger Morra

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

Parameterize with **probabilities**:

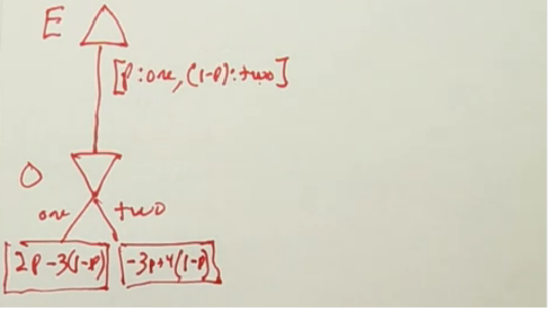

	O: one	*O*: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

- Calculate ***E*'s utility** for both players' mixed strategies.

Parameterize with **probabilities**:

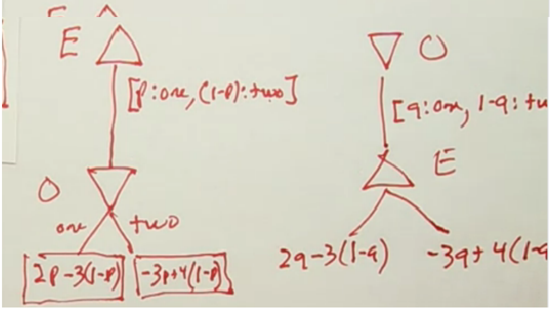

Mixed Strategy: 2 Finger Morra

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

- Calculate **E's utility** for both players' mixed strategies.
- **Mixed strategy:** Leave opponent no choice!

Parameterize with **probabilities**:

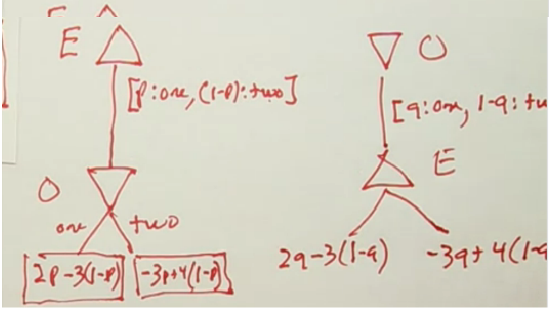

Mixed Strategy: 2 Finger Morra

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

- Calculate ***E*'s utility** for both players' mixed strategies.
- **Mixed strategy:** Leave opponent no choice!

Parameterize with **probabilities**:

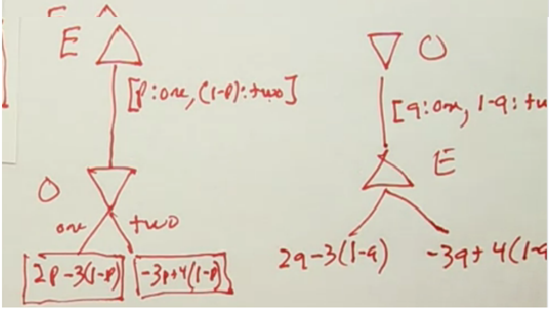

Optimal p for E:

$$
\begin{aligned}
2p - 3(1-p) &= -3p + 4(1-p) \\
p &= 7/12 \\
U_E &= 2p - 3(1-p) \\
&= -1/12
\end{aligned}
$$

Mixed Strategy: 2 Finger Morra

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

- Calculate **E's utility** for both players' mixed strategies.
- **Mixed strategy:** Leave opponent no choice!

Parameterize with **probabilities**:

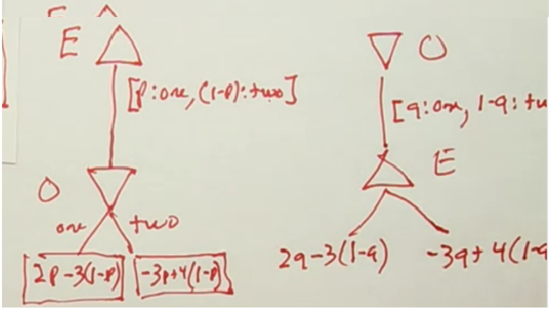

Optimal p for E:

$$\begin{aligned} 2p - 3(1-p) &= -3p + 4(1-p) \\ p &= 7/12 \\ U_E &= 2p - 3(1-p) \\ &= -1/12 \end{aligned}$$

Optimal q for O:

$$\begin{aligned} 2q - 3(1-q) &= -3q + 4(1-q) \\ q &= 7/12 \\ U_E &= 3q + 4(1-q) \\ &= -1/12 \end{aligned}$$

	O: one	O: two
E: one	$E = +2$	$E = -3$
E: two	$E = -3$	$E = +4$

- Calculate **E's utility** for both players' mixed strategies.
- **Mixed strategy:** Leave opponent no choice!

Parameterize with **probabilities**:

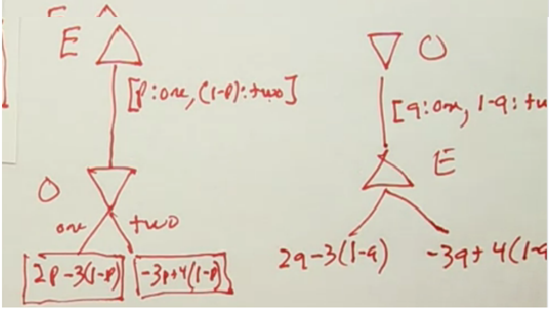

Optimal p for E:

$$\begin{aligned} 2p - 3(1-p) &= -3p + 4(1-p) \\ p &= 7/12 \\ U_E &= 2p - 3(1-p) \\ &= -1/12 \end{aligned}$$

$\leq U_E \leq$

Optimal q for O:

$$\begin{aligned} 2q - 3(1-q) &= -3q + 4(1-q) \\ q &= 7/12 \\ U_E &= 3q + 4(1-q) \\ &= -1/12 \end{aligned}$$

Mixed Strategy Issues

Secrecy and Rationality:

Secrecy: If a dominant strategy exists, your opponent can guess it!

Rationality: Sometimes it's better to **look crazy** to make your opponent believe you will do something irrational.

Mixed Strategy Issues

Secrecy and Rationality:

Secrecy: If a dominant strategy exists, your opponent can guess it!

Rationality: Sometimes it's better to **look crazy** to make your opponent believe you will do something irrational.

A riddle for you.

Mixed Strategy Example

Zero-sum game with min and max:

	$\nabla : 1$	$\nabla : 2$
$\triangle$: 1	$\triangle = 5$	3
$\triangle$: 2	4	2

Let's solve it?

Mixed Strategy Example

Zero-sum game with min and max:

	$\nabla : 1$	$\nabla : 2$
$\triangle : 1$	$\triangle = 5$	3
$\triangle : 2$	4	2

Let's solve it?

No, need! **Dominant strategies exist!**

Mixed Strategy Example

Zero-sum game with min and max:

	$\nabla : 1$	$\nabla : 2$
$\triangle : 1$	$\triangle = 5$	3
$\triangle : 2$	4	2

Let's solve it?

No, need! **Dominant strategies exist!**

1. We know what to do
2. We can guess the other rational player's move

Mixed Strategy Example

Zero-sum game with min and max:

	$\nabla : 1$	$\nabla : 2$
$\triangle : 1$	$\triangle = 5$	3
$\triangle : 2$	4	2

Let's solve it?

No, need! **Dominant strategies exist!**

1. We know what to do
2. We can guess the other rational player's move

Therefore,

$$U_E = 5 \, .$$

Another Mixed Strategy Example

	$\nabla : 1$	$\nabla : 2$
$\triangle : 1$	$\triangle = 3$	6
$\triangle : 2$	4	5

Dominant strategies?

Another Mixed Strategy Example

	$\bigtriangledown: 1$	$\bigtriangledown: 2$
$\triangle$: 1	$\triangle = 3$	6
$\triangle$: 2	4	5

Dominant strategies? None for $\triangle$.

- Need to calculate only probability p, because dominant $q = 0$.

Another Mixed Strategy Example

	$\nabla: 1$	$\nabla: 2$
$\triangle$: 1	$\triangle = 3$	6
$\triangle$: 2	4	5

Dominant strategies? None for $\triangle$.

- Need to calculate only probability p, because dominant $q = 0$.

$$\begin{aligned} 3p + 5(1-p) &= 6p + 4(1-p) \\ p &= 1/4 \\ U_\triangle &= 4.5 \end{aligned}$$

Another Mixed Strategy Example

	$\nabla:1$	$\nabla:2$
$\triangle$: 1	$\triangle$= 3	6
$\triangle$: 2	4	5

Dominant strategies? None for $\triangle$.

- Need to calculate only probability p, because dominant $q = 0$.

$$\begin{aligned} 3p + 5(1-p) &= 6p + 4(1-p) \\ p &= 1/4 \\ U_\triangle &= 4.5 \end{aligned}$$

Based on $\triangle$'s decision,

$$\begin{aligned} U_\triangle &= 3q + 6(1-q) \\ &= 6 \end{aligned}$$

or

$$\begin{aligned} U_\triangle &= 5q + 4(1-q) \\ &= 4 \end{aligned}$$

Another Mixed Strategy Example

	$\bigtriangledown : 1$	$\bigtriangledown : 2$
$\triangle$: 1	$\triangle = 3$	6
$\triangle$: 2	4	5

Dominant strategies? None for $\triangle$.

- Need to calculate only probability p, because dominant $q = 0$.

$$3p + 5(1-p) = 6p + 4(1-p)$$
$$p = 1/4$$
$$U_\triangle = 4.5$$

Based on $\triangle$'s decision,

$$U_\triangle = 3q + 6(1-q) = 6$$

or

$$U_\triangle = 5q + 4(1-q) = 4$$

Therefore,

$$4 \le U_\triangle \le 6 \,.$$

Poker

Simplified!

- Deck has only ace and kings: AAKK
- Deal: 1 card each

Rounds:

1. raise/check
2. call/fold

Sequential game/ extensive form

Poker

Simplified!

- Deck has only ace and kings: AAKK
- Deal: 1 card each

Rounds:

1. raise/check
2. call/fold

Sequential game/ extensive form

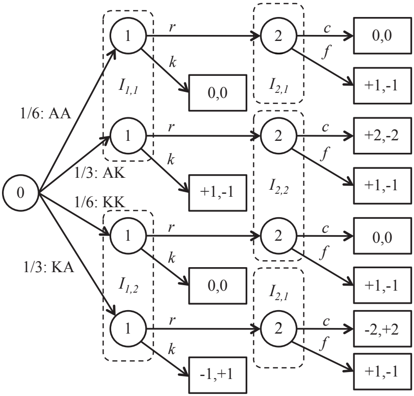

Poker

Simplified!

- Deck has only ace and kings: AAKK
- Deal: 1 card each

Rounds:

1. raise/check
2. call/fold

Sequential game/ extensive form

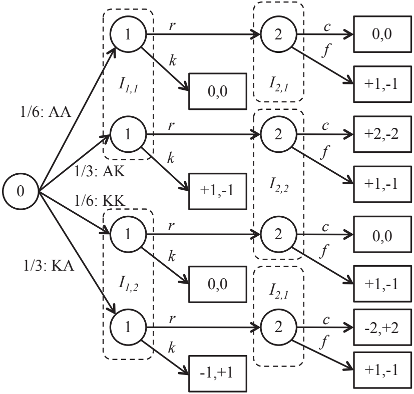

	2:*cc*	2:*cf*	2:*ff*	2:*fc*
1:*rr*	0	-1/6	1	7/6
1:*kr*	-1/3	-1/6	5/6	2/3
1:*rk*	1/3	**0**	1/6	1/2
1:*kk*	0	**0**	0	0

Poker

Simplified!

- Deck has only ace and kings: AAKK
- Deal: 1 card each

Rounds:

1. raise/check
2. call/fold

Sequential game/ extensive form

- Real game has how many states?

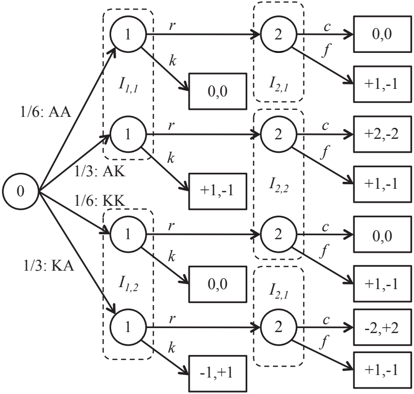

	2:*cc*	2:*cf*	2:*ff*	2:*fc*
1:*rr*	0	-1/6	1	7/6
1:*kr*	-1/3	-1/6	5/6	2/3
1:*rk*	1/3	**0**	1/6	1/2
1:*kk*	0	**0**	0	0

Poker

Simplified!

- Deck has only ace and kings: AAKK
- Deal: 1 card each

Rounds:

1. raise/check
2. call/fold

Sequential game/ extensive form

- Real game has how many states? $\sim 10^{18}$

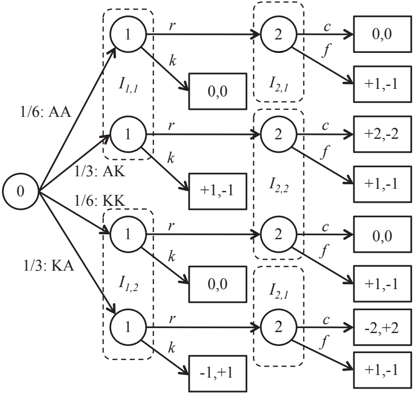

	2:cc	2:cf	2:ff	2:fc
1:rr	0	-1/6	1	7/6
1:kr	-1/3	-1/6	5/6	2/3
1:rk	1/3	**0**	1/6	1/2
1:kk	0	**0**	0	0

So How To Solve Non-Simplified Games?

Strategies:

- abstracting; lumping together:
  - Don't care about aces' suits, all aces equal
  - Lump similar cards together: cards 1–7 together
  - Bets: small and large
  - Deals: Monte Carlo sampling

So How To Solve Non-Simplified Games?

Strategies:

- abstracting; lumping together:
  - Don't care about aces' suits, all aces equal
  - Lump similar cards together: cards 1–7 together
  - Bets: small and large
  - Deals: Monte Carlo sampling

In summary, **game theory** is:

- Good for:

So How To Solve Non-Simplified Games?

Strategies:

- abstracting; lumping together:
  - Don't care about aces' suits, all aces equal
  - Lump similar cards together: cards 1–7 together
  - Bets: small and large
  - Deals: Monte Carlo sampling

In summary, **game theory** is:

- Good for: simultaneous moves, stochasticity, uncertainty, partial observability, multi-agent, sequential, dynamic

So How To Solve Non-Simplified Games?

Strategies:

- abstracting; lumping together:
  - Don't care about aces' suits, all aces equal
  - Lump similar cards together: cards 1–7 together
  - Bets: small and large
  - Deals: Monte Carlo sampling

In summary, **game theory** is:

- Good for: simultaneous moves, stochasticity, uncertainty, partial observability, multi-agent, sequential, dynamic
- Not good for:

So How To Solve Non-Simplified Games?

Strategies:

- abstracting; lumping together:
  - Don't care about aces' suits, all aces equal
  - Lump similar cards together: cards 1–7 together
  - Bets: small and large
  - Deals: Monte Carlo sampling

In summary, **game theory** is:

- Good for: simultaneous moves, stochasticity, uncertainty, partial observability, multi-agent, sequential, dynamic
- Not good for: unknown actions, continuous actions, irrational opponents, unknown utility

Game: Feds vs. Politicians

Game between:

1. Federal reserve and
2. Politicians

on controlling the budget.

Find equilibrium below:

Game between:

1. Federal reserve and
2. Politicians

on controlling the budget.

Find equilibrium below:

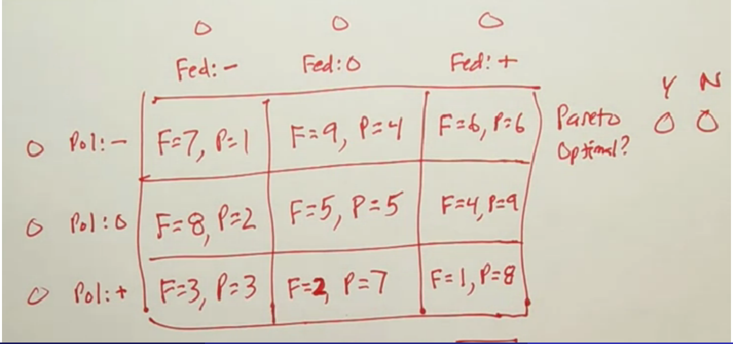

Game between:

1. Federal reserve and
2. Politicians

on controlling the budget.

Find equilibrium below:

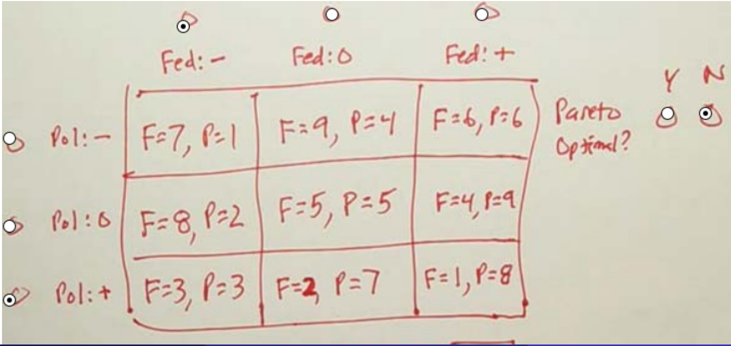

Alturistic Side of Game Theory: Mechanism Design

Mechanism Design:

- Design to get the most for the game for: players, game itself, or public
- Example: advertisements

Alturistic Side of Game Theory: Mechanism Design

Mechanism Design:

- Design to get the most for the game for: players, game itself, or public
- Example: advertisements

Strategy: design game to have dominant strategy

- Example: second-price auctions (like eBay)