Lecture 11
Other Solution Concepts and Game Models

CS598

Ruta Mehta

Some slides are borrowed from V. Conitzer's presentations.

So far

- Normal-form games
  - Multiple rational players, single shot, simultaneous move
- Nash equilibrium
  - Existence
  - Computation in two-player games.

Today:

- Issues with NE
  - Multiplicity
  - Selection: How players decide/reach any particular NE

- Possible Solutions
  - Dominance: Dominant Strategy equilibria
  - Arbitrator/Mediator: Correlated equilibria, Coarse-correlated equilibria
  - Communication/Contract: Stackelberg equilibria, Nash bargaining
- Other Games
  - Extensive-form Games, Bayesian Games

Dominance

- Strict dominance: For player move s_i strictly dominates s_i' if no matter what others play s_i gives better payoff than s_i'
  - for all s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$
- s_i weakly dominates s_i' if
  - for all 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"

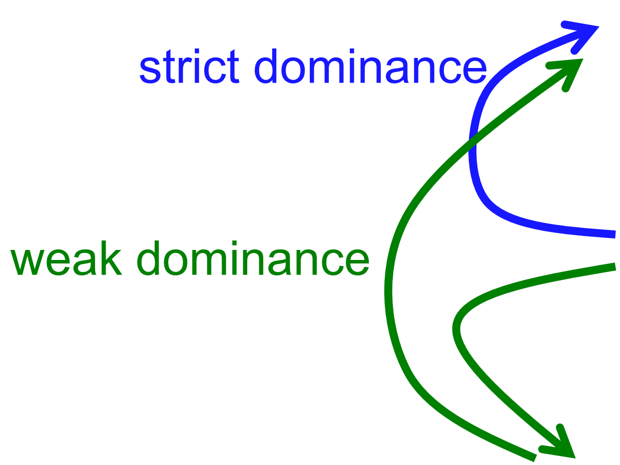

	L	M	R
U	0, 0	1, -1	1, -1
G	-1, 1	0, 0	-1, 1
B	-1, 1	1, -1	0, 0

Dominant Strategy Equilibrium

Playing move s is best for me, no matter what others play.

- For each player i, there is a (move) strategy s_i that (weakly) dominates all other strategies.
  - for all $\mathrm{i}, \mathrm{s}_{\mathrm{i}}', s_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$;

Example?

Dominance by Mixed strategies

- Example of dominance by a mixed strategy:

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

Iterated dominance: path (in)dependence

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

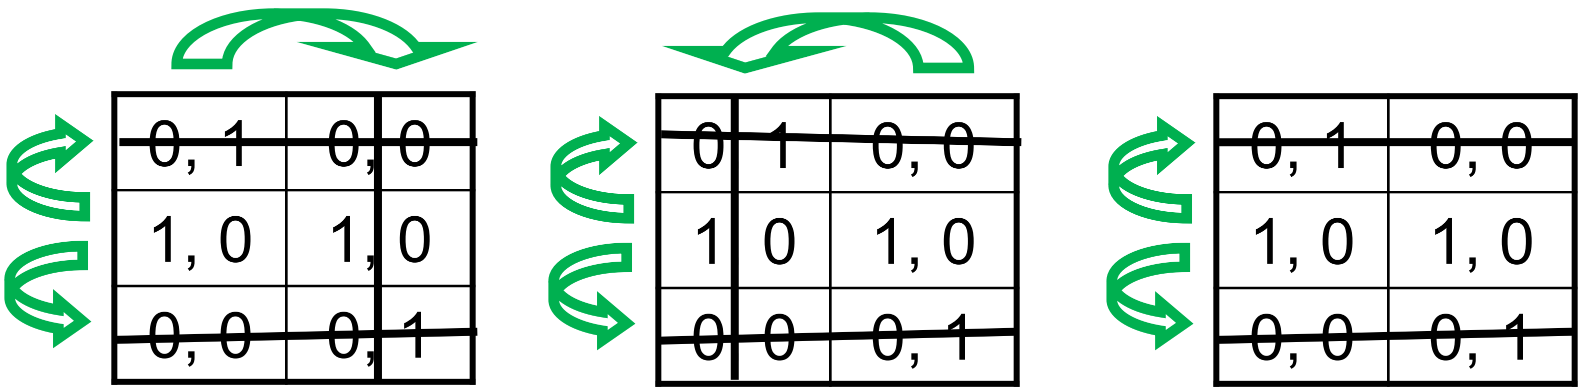

Iterated strict dominance is path-independent: elimination process will always terminate at the same point
(whether or not dominance by mixed strategies allowed)

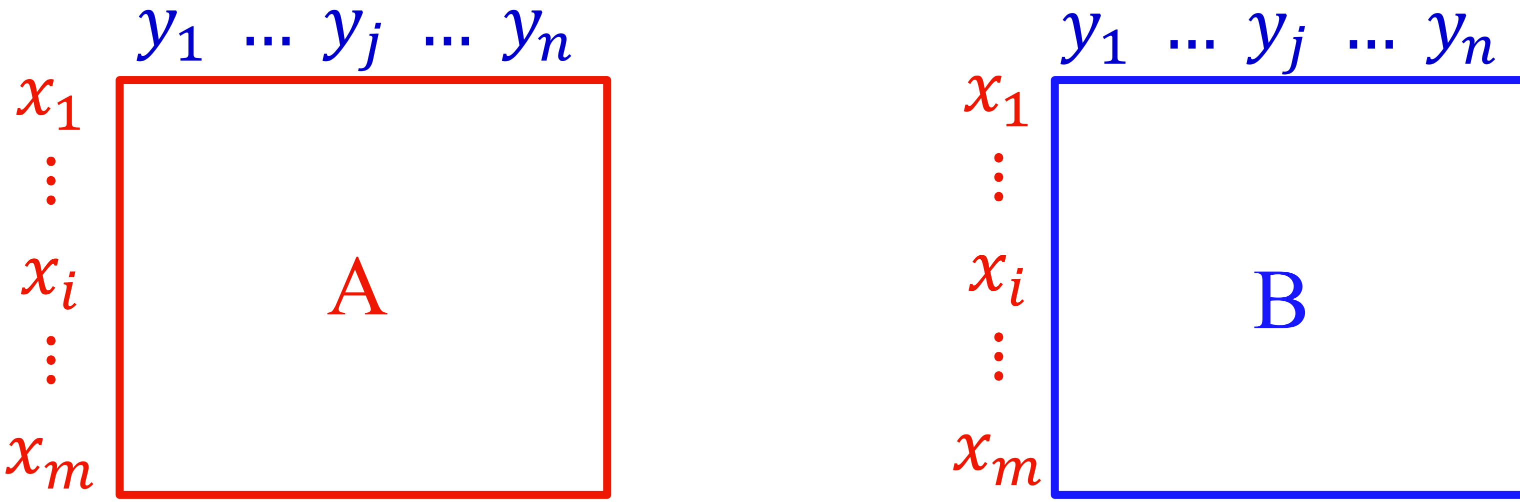

NE: $x^T A y \geq x'^T A y,\ \forall x'$ $\quad$ $x^T B y \geq x^T B y',\ \forall y'$

No one plays dominated strategies.

Why?

What if they can discuss beforehand?

Players: {Alice, Bob}

Two options: {Football, Tennis}

		2/3 F	1/3 T
1/3	F	1 2 (0.5)	0 0
2/3	T	0 0	2 1 (0.5)

At Mixed NE both get 2/3 < 1

Instead they agree on ½(F, T), ½(T, F)

Payoffs are (1.5, 1.5) Fair!

Needs a common coin toss!

Correlated Equilibrium – (CE)
(Aumann'74)

- **Mediator** declares a joint distribution P over S=$\times_i S_i$
- Tosses a coin, chooses $s = (s_1, \ldots, s_n) \sim P$.
- Suggests s_i to player i **in private**

- P is at equilibrium if each player wants to follow the suggestion when others do.
  - $U_i\left(s_i, P_{(s_i,.)}\right) \geq U_i\left(s_i', P_{(s_i,.)}\right), \forall s_i' \in S_1$

$\sum_{s_{-i} \in S_{-i}} P(s_i, s_{-i}) U_i(s_i, s_{-i})$ Linear in P variables!

Players: {Alice, Bob}

Two options: {Football, Shopping}

	F	S
F	1 2 (0.5)	0 0
S	0 0	2 1 (0.5)

Instead they agree on ½(F, S), ½(S, F)

Payoffs are (1.5, 1.5)

Fair!

CE!

Prisoner's Dilemma

	C	NC
C	-5, -5 (1)	0, -6 (0)
NC	-6, 0 (0)	-1, -1 (0)

NC is dominated

Rock-Paper-Scissors (Aumann)

	R	P	S
R	0, 0 (0)	0, 1 (1/6)	1, 0 (1/6)
P	1, 0 (1/6)	0, 0 (0)	0, 1 (1/6)
S	0, 1 (1/6)	1, 0 (1/6)	0, 0 (0)

When Alice is suggested R

Bob must be following $P_{(R,.)} = (0,1/6,1/6)$

Following the suggestion gives her 1/6

While P gives 0, and S gives 1/6.

Computation: Linear Feasibility Problem

Game (A, B). Find, J.D. $P = \begin{bmatrix} p_{11} & \cdots & p_{1n} \\ \vdots & \vdots & \vdots \\ p_{m1} & \cdots & p_{mn} \end{bmatrix}$

s.t.

$$\sum_j A_{ij} p_{ij} \geq \sum_j A_{i'j} p_{ij} \quad \forall i, i' \in S_1$$

$$\sum_i B_{ij} p_{ij} \geq \sum_i B_{ij'} p_{ij} \quad \forall j, j' \in S_2$$

$$\sum_{ij} p_{ij} = 1$$

N-player game: Find distribution P over $S = \times_{i=1}^{N} S_i$

s.t. $U_i\left(s_i, P_{(s_i,.)}\right) \geq U_i\left(s_i', P_{(s_i,.)}\right), \forall s_i, s_i' \in S_i$

$$\sum_{s \in S} P(s) = 1$$

↑

$\sum_{s_{-i} \in S_{-i}} U_i(s_i, s_{-i}) P(s_i, s_{-i})$ Linear in P variables!

Computation: Linear Feasibility Problem

N-player game: Find distribution P over $S = \times_{i=1}^{N} S_i$

s.t. $U_i\left(s_i, P_{(i,.)}\right) \geq U_i\left(s_i', P_{(s_i,.)}\right)$, $\forall s_i, s_i' \in S_i$

$\uparrow$ $\sum_{s \in S} P(s) = 1$

$\sum_{s_{-i} \in S_{-i}} U_i(s_i, s_{-i}) P(s_i, s_{-i})$ Linear in P variables!

Can optimize any convex function as well!

Coarse- Correlated Equilibrium

- After mediator declares P, each player opts in or out.
- Mediator tosses a coin, and chooses s ~ P.
- If player i opted in, then the mediator suggests her s_i in private, and she has to obey.
- If she opted out, then (knowing nothing about s) plays a fixed strategy $t \in S_i$
- At equilibrium, each player wants to opt in, if others are.

$$U_i(P) \geq U_i(t, P_{-i}), \ \forall t \in S_i$$

Where P_{-i} is joint distribution of all players except i.

Importance of (Coarse) CE

- Natural dynamics quickly arrive at approximation of such equilibria.
  - No-regret, Multiplicative Weight Update (MWU)

- Poly-time computable in the size of the game.
  - Can optimize a convex function too.

Show the following

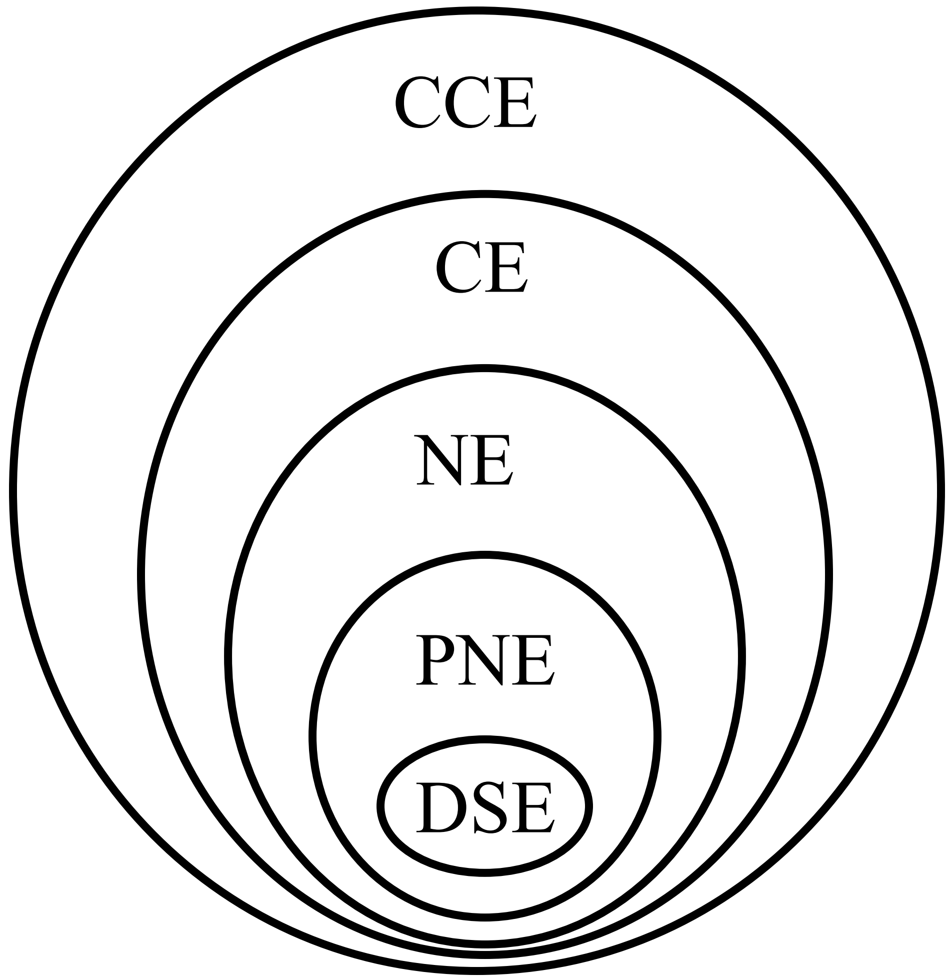

Extensive-form Game

- Players move one after another
  - Chess, Poker, etc.
  - Tree representation.

Strategy of a player:
What to play at each of its node.

	I	O
F	-1, 1	2, 0
A	1, 1	2, 0

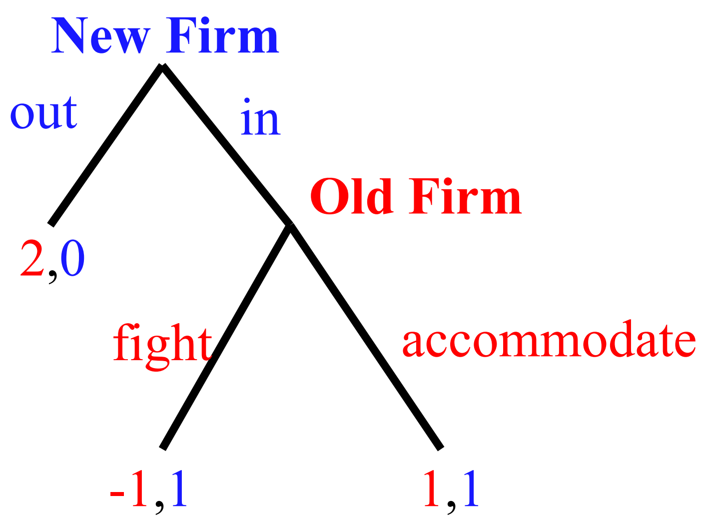

Entry game

A poker-like game

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

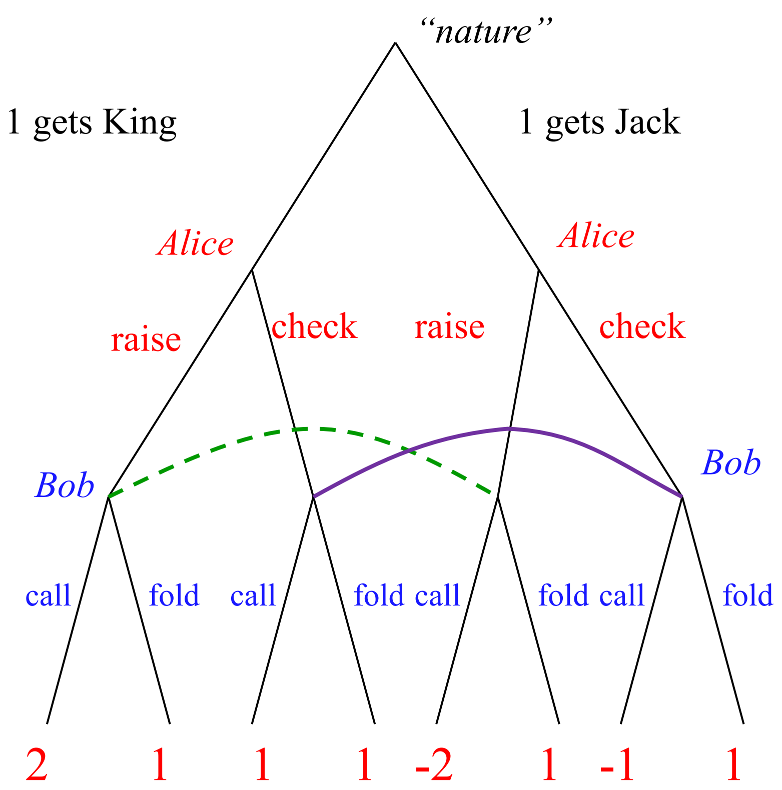

Poker-like game in normal form

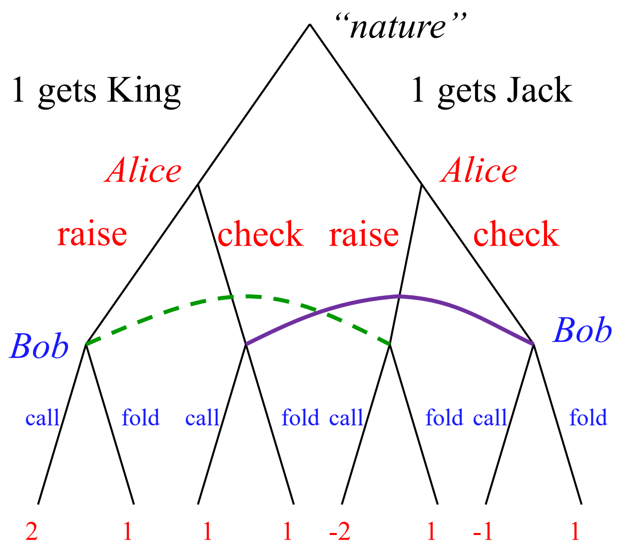

	cc	cf	fc	ff
rr	0, 0	0, 0	1, -1	1, -1
rc	.5, -.5	1.5, -1.5	0, 0	1, -1
cr	-.5, .5	-.5, .5	1, -1	1, -1
cc	0, 0	1, -1	0, 0	1, -1

Can be exponentially big!

Sub-Game Perfect Equilibrium

- Every sub-tree is at equilibrium
- Computation when perfect information (no nature/chance move): Backward induction

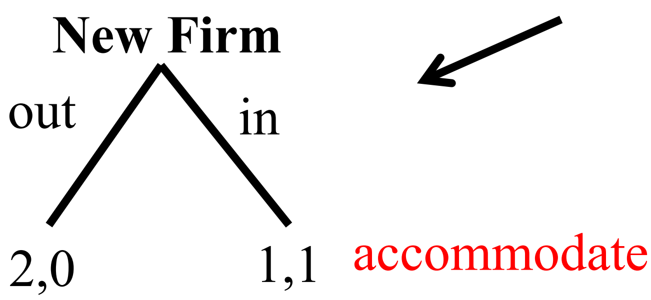

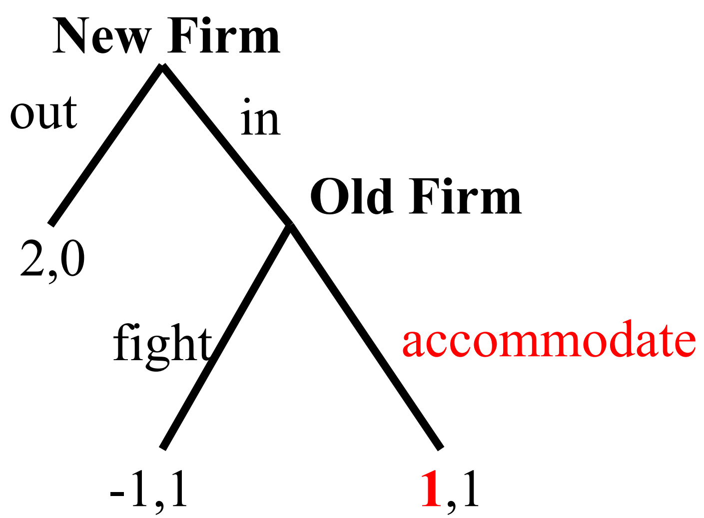

Entry game

Sub-Game Perfect Equilibrium

- Every sub-tree is at equilibrium
- Computation when perfect information (no nature/chance move): Backward induction

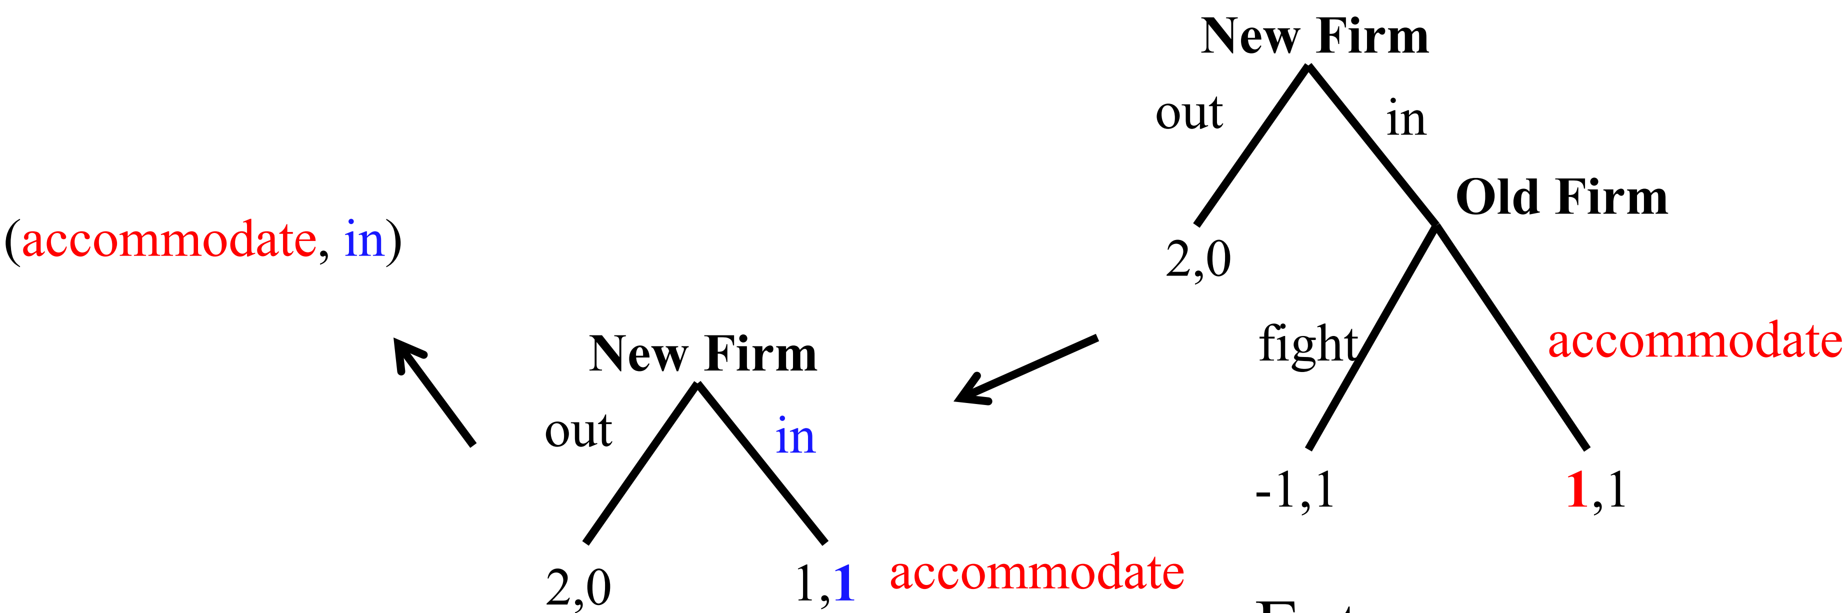

Entry game

Corr. Eq. in Extensive form Game

- How to define?
  - CE in its normal-form representation.
- Is it computable?
  - Recall: exponential blow up in size.
- Can there be other notions?

See "Extensive-Form Correlated Equilibrium: Definition and Computational Complexity" by von Stengel and Forges, 2008.

Commitment
(Stackelberg strategies)

Commitment

Unique Nash equilibrium (iterated strict dominance solution) → top-left cell

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

von Stackelberg

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

Commitment: an extensive-form game

For the case of committing to a pure strategy:

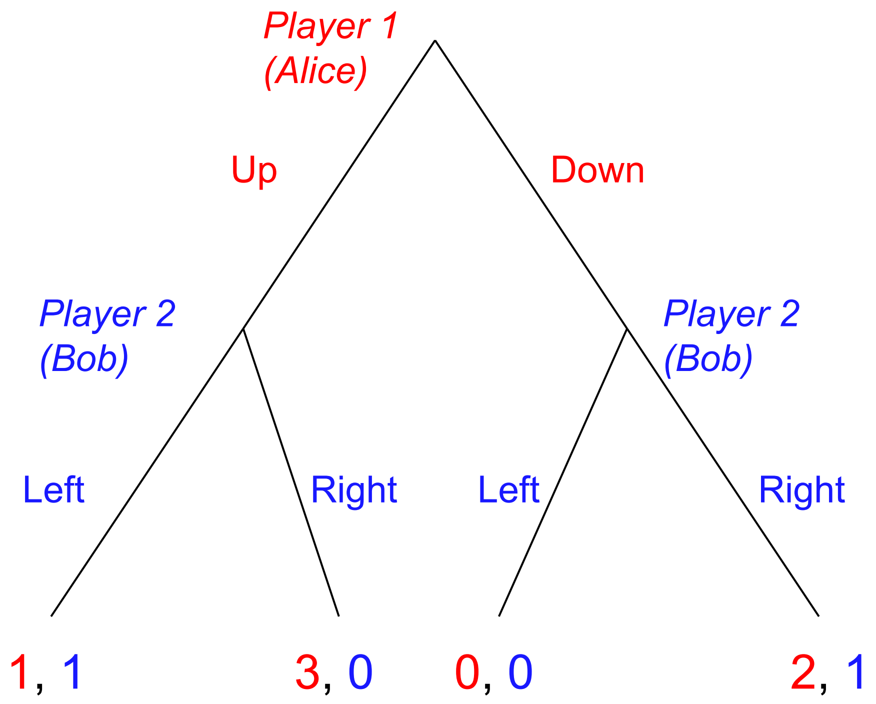

Commitment to mixed strategies

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

Also called a Stackelberg (mixed) strategy

Commitment: an extensive-form game

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

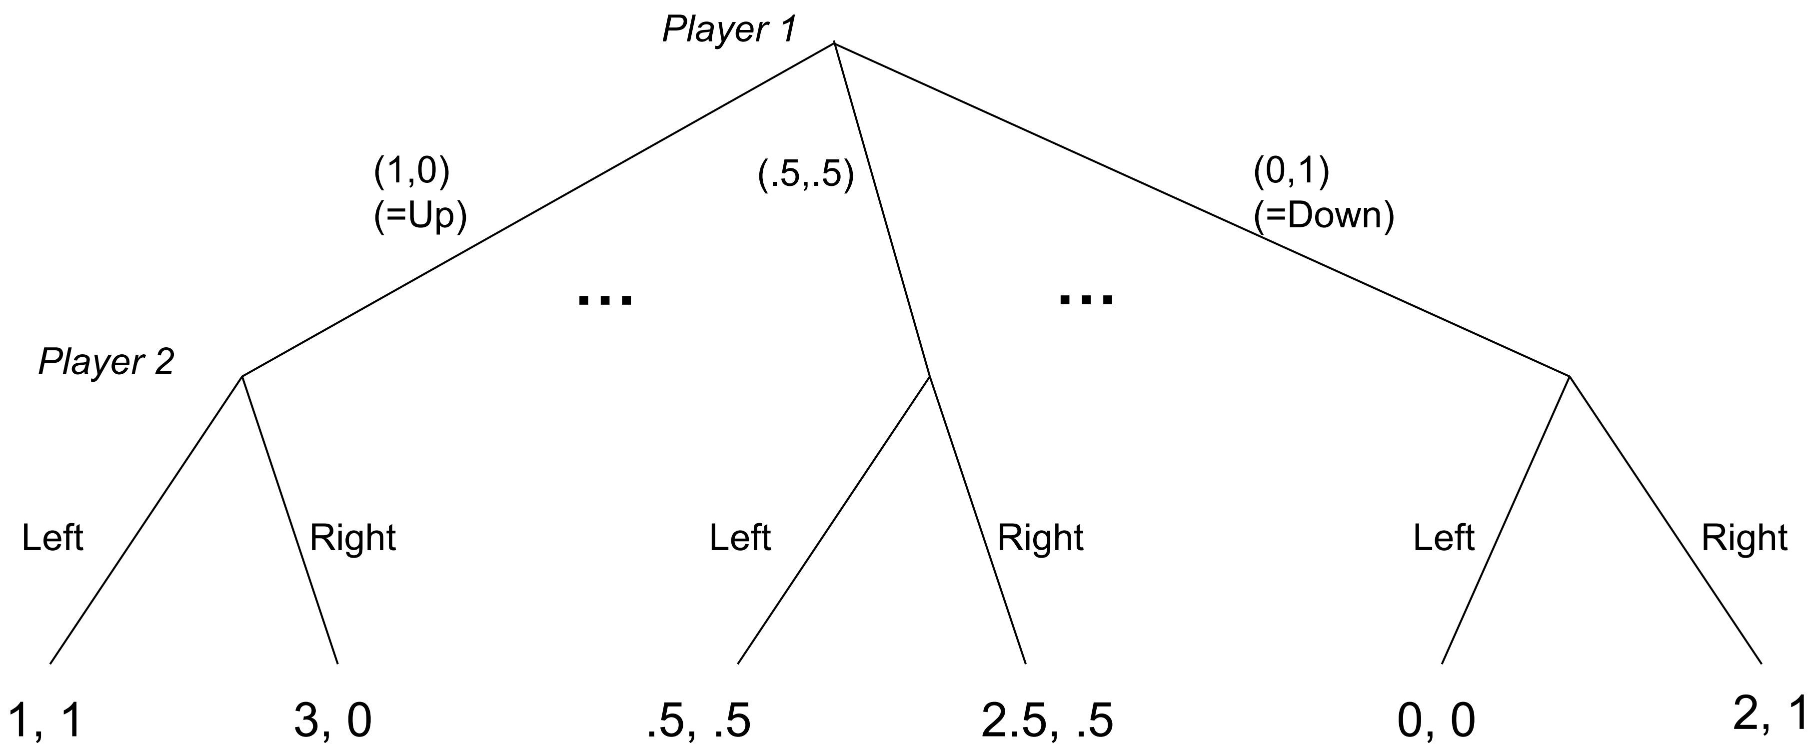

- Economist: Just an extensive-form game, nothing new here
- Computer scientist: Infinite-size game! Representation matters

Computing the optimal mixed strategy to commit to

[Conitzer & Sandholm EC'06]

- Player 1 (Alice) is a leader.
- Separate LP for every column $j^* \in S_2$:

maximize $\sum_i x_i A_{i\,j^*}$ Alice's utility when Bob plays j^*

subject to $\forall j,\ \ (x^T B)_{j^*} \geq (x^T B)_j$ Playing j^* is best for Bob

$x \geq 0,\ \sum_i x_i = 1$ x is a probability distribution

Among soln. of all the LPs,
pick the one that gives max utility.

On the game we saw before

x_1	1, 1	3, 0
x_2	0, 0	2, 1

$$\textit{maximize } \mathbf{1} x_1 + \mathbf{0}\, x_2$$

$$\textit{subject to}$$

$$\mathbf{1}\, x_1 + \mathbf{0}\, x_2 \geq \mathbf{0}\, x_1 + \mathbf{1}\, x_2$$

$$x_1 + x_2 = 1$$

$$x_1 \geq 0,\ x_2 \geq 0$$

$$\textit{maximize } \mathbf{3}\, x_1 + \mathbf{2}\, x_2$$

$$\textit{subject to}$$

$$\mathbf{0}\, x_1 + \mathbf{1}\, x_2 \geq \mathbf{1}\, x_1 + \mathbf{0}\, x_2$$

$$x_1 + x_2 = 1$$

$$x_1 \geq 0,\ x_2 \geq 0$$

Generalizing beyond zero-sum games

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

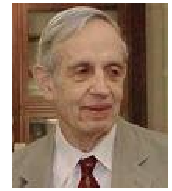

0, 0	-1, 1
-1, 1	0, 0

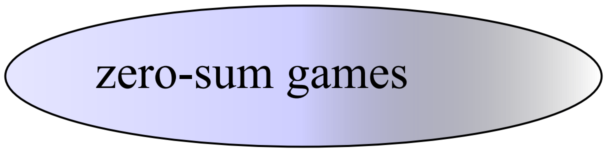

minimax strategies

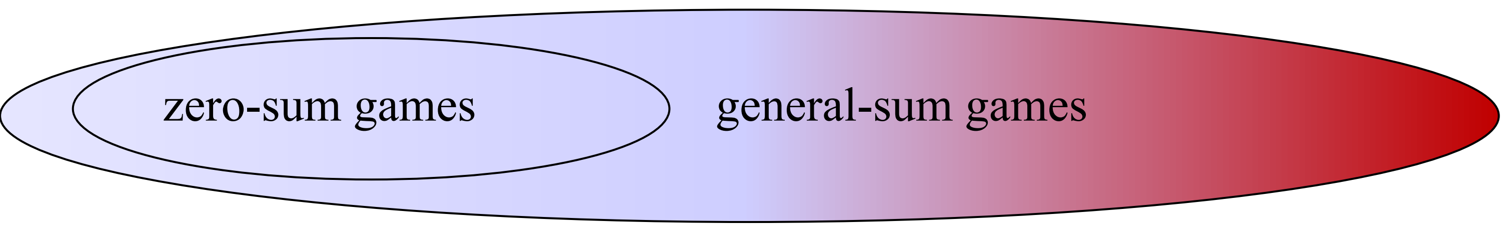

Nash equilibrium

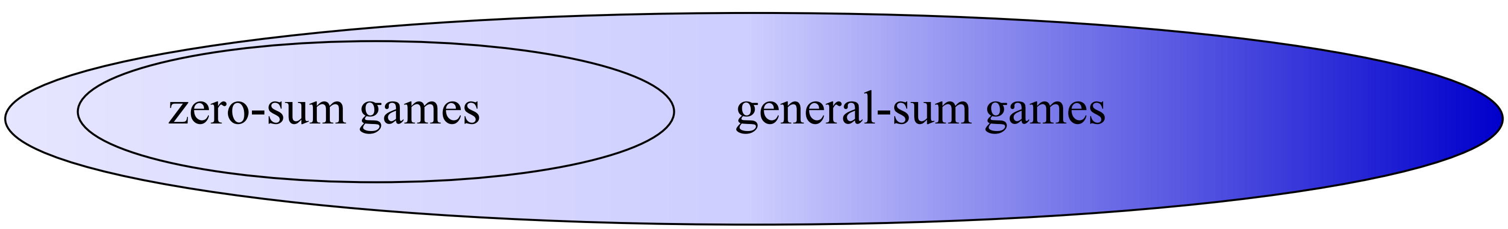

Stackelberg mixed strategies

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])

≥

Bayesian Games

So far in Games,

- Complete information (each player has perfect information regarding the element of the game).

Bayesian Game

- A game with incomplete information
- Each player has initial private information, **type**.
- Bayesian equilibrium: solution of the Bayesian game

Bayesian game

- Utility of a player depends on her **type** and the actions taken in the game
  - θ_i is player i's type, $\theta_i \sim \Theta_i$. Utilily when θ_i type and s play is $u_i(\theta_i, s)$
  - Each player knows/learns its own type, but only distribution of others (before choosing action)
    - Pure strategy $s_i : \Theta_i \rightarrow S_i$ (where S_i is i's set of actions)

(In general players can also receive signals about other players' utilities; we will not go into this)

row player (Alice)
type 1 (prob. 0.5)

	L	R
U	4	6
D	2	4

column player (Bob)
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

Car Selling Game

- A seller wants to sell a car
- A buyer has private value ‘v’ for the car w.p. P(v)
- Sellers knows P, but not v
- Seller sets a price ‘p’, and buyer decides to buy or not buy.
- If sell happens then the seller gets p, and buyer gets (v-p).

S_1=All possible prices, Θ_1={1}
S_2={buy, not buy}, Θ_2 =All possible ‘v’
$U_1(1,(p,\text{buy})) = p$, $U_1(1,(p,\text{not buy})) = 0$
$U_2(v,(p,\text{buy}))=v-p$, $U_2(v,(p,\text{not buy})) = 0$

Converting Bayesian games to normal form

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

	type 1: L type 2: L	type 1: L type 2: R	type 1: R type 2: L	type 1: R type 2: R
type 1: U type 2: U	3, 3	4, 3	4, 4	5, 4
type 1: U type 2: D	4, 3.5	4, 3	4, 4.5	4, 4
type 1: D type 2: U	2, 3.5	3, 3	3, 4.5	4, 4
type 1: D type 2: D	3, 4	3, 3	3, 5	3, 4

exponential blowup in size

Bayes-Nash equilibrium

- A profile of strategies is a Bayes-Nash equilibrium if it is a Nash equilibrium for the normal form of the game
  - Minor caveat: each type should have >0 probability

- Alternative definition:
  - Mixed strategy of player i, $\sigma_i : \Theta_i \to \Delta(S_i)$
  - for every i, for every type θ_i, for every alternative action s_i, we must have:

$$\Sigma_{\theta_{-i}} P(\theta_{-i})\, u_i(\theta_i, \sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})) \geq \Sigma_{\theta_{-i}} P(\theta_{-i})\, u_i(\theta_i, s_i, \sigma_{-i}(\theta_{-i}))$$

where $P(\theta_{-i}) = \Pi_{p \neq i} P(\theta_p)$

Again what about corr. eq. in Bayesian games?

Notion of signaling.

Look up the literature.