Introduction to Game Theory

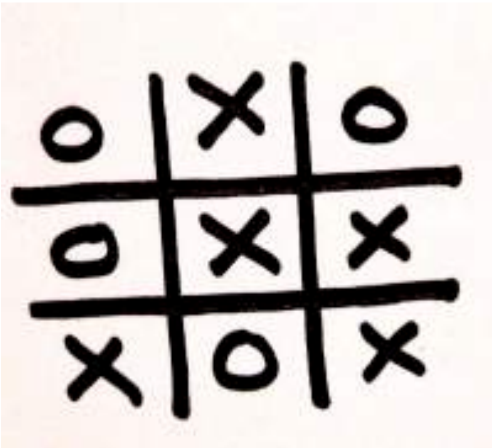

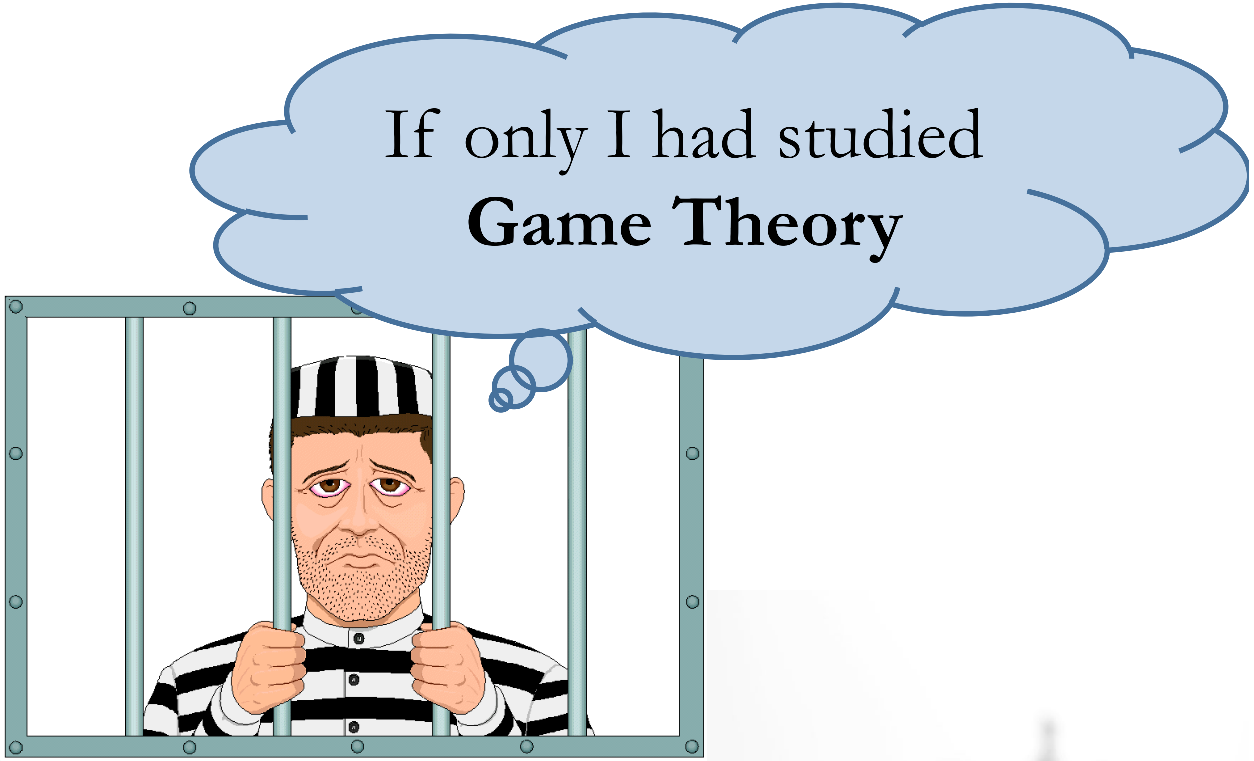

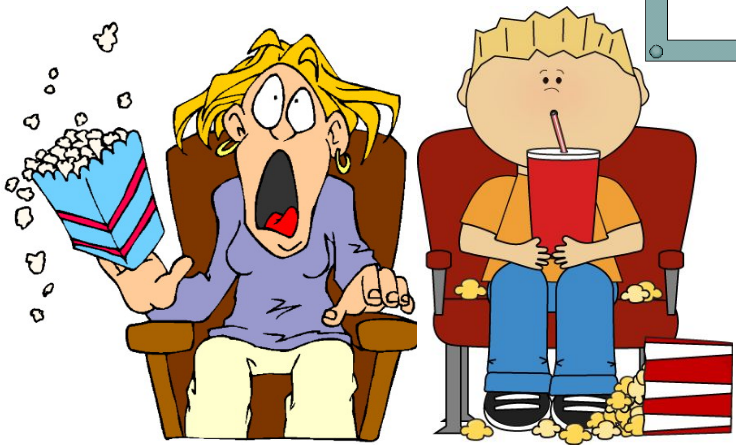

Introduction to Game Theory

Game theory is the mathematical study of interaction among independent, self-interested agents.

Outlook

- **Games in Normal Form**

Perfect-information Sequential Actions Games

Imperfect-information SA Games

Repeated Games

Bayesian Games

Coalitional Games

Normal Form

Definition 1.2.1 (Normal-form game). *A (finite,* n*-person)* normal-form game *is a tuple* (N, A, u)*, where:*

- *N is a finite set of n players, indexed by i;*
- $A = A_1 \times \cdots \times A_n$*, where A_i is a finite set of* actions *available to player i. Each vector* $a = (a_1, \ldots, a_n) \in A$ *is called an* action profile*;*
- $u = (u_1, \ldots, u_n)$ *where $u_i : A \mapsto \mathbb{R}$ is a real-valued* utility *(or* payoff*)* function *for player i.*

Actually, it's a table :)
such as this one

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Prisoner's dilemma

2 players → 2 dimensions

Payoff for particular action profile

Available actions

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Special cases

Definition 1.3.1 (Common-payoff game). *A* common-payoff game *is a game in which for all action profiles $a \in A_1 \times \cdots \times A_n$ and any pair of agents i, j, it is the case that $u_i(a) = u_j(a)$.*

Definition 1.3.2 (Constant-sum game). *A two-player normal-form game is* constant-sum *if there exists a constant c such that for each strategy profile $a \in A_1 \times A_2$ it is the case that $u_1(a) + u_2(a) = c$.*

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

	Left	Right
Left	$1, 1$	$0, 0$
Right	$0, 0$	$1, 1$

Strategies

Definition 1.4.1 (Mixed strategy). *Let (N, A, u) be a normal-form game, and for any set X let $\Pi(X)$ be the set of all probability distributions over X. Then the set of* mixed strategies *for player i is $S_i = \Pi(A_i)$.*

Definition 1.4.2 (Mixed-strategy profile). *The set of* mixed-strategy profiles *is simply the Cartesian product of the individual mixed-strategy sets, $S_1 \times \cdots \times S_n$.*

Definition 1.4.3 (Support). *The* support *of a mixed strategy s_i for a player i is the set of pure strategies $\{a_i | s_i(a_i) > 0\}$.*

Definition 1.4.4 (Expected utility of a mixed strategy). *Given a normal-form game (N, A, u), the expected utility u_i for player i of the mixed-strategy profile $s = (s_1, \ldots, s_n)$ is defined as*

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j).$$

Strategy profile dominance

How to measure "goodness" of a strategy profile?

Is strategy profile A better than B?

Definition 2.1.1 (Pareto domination). *Strategy profile* s Pareto dominates *strategy profile* s' *if for all* $i \in N$, $u_i(s) \geq u_i(s')$, *and there exists some* $j \in N$ *for which* $u_j(s) > u_j(s')$.

Definition 2.1.2 (Pareto optimality). *Strategy profile* s *is* Pareto optimal, *or* strictly Pareto efficient, *if there does not exist another strategy profile* $s' \in S$ *that Pareto dominates* s.

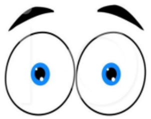

There is always some Pareto optimal (pure) strategy profile.

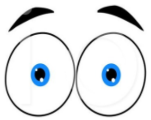

In common-payoff games, all Pareto optimal strategy profiles have the same payoffs.

Nash

Definition 2.2.1 (Best response). *a mixed*
strategy $s_i^* \in S_i$ *such that* $u_i(s_i^*, s_{-i})$

Best response which … esponses.

Definition 2.2.2 (Nash equilib … equilibrium
if, for all agents i, s_i *is a best res*

Definition 2.2.3 (Strict Nas … *agents* i *and for*
all strategies $s_i' \neq s_i$, $u_i(s_i, s$

Definition 2.2.4 (Weak … *all agents* i *and for*
all strategies $s_i' \neq s_i$, $u_i(s$ … *rium.*

Mixed strategy ca…

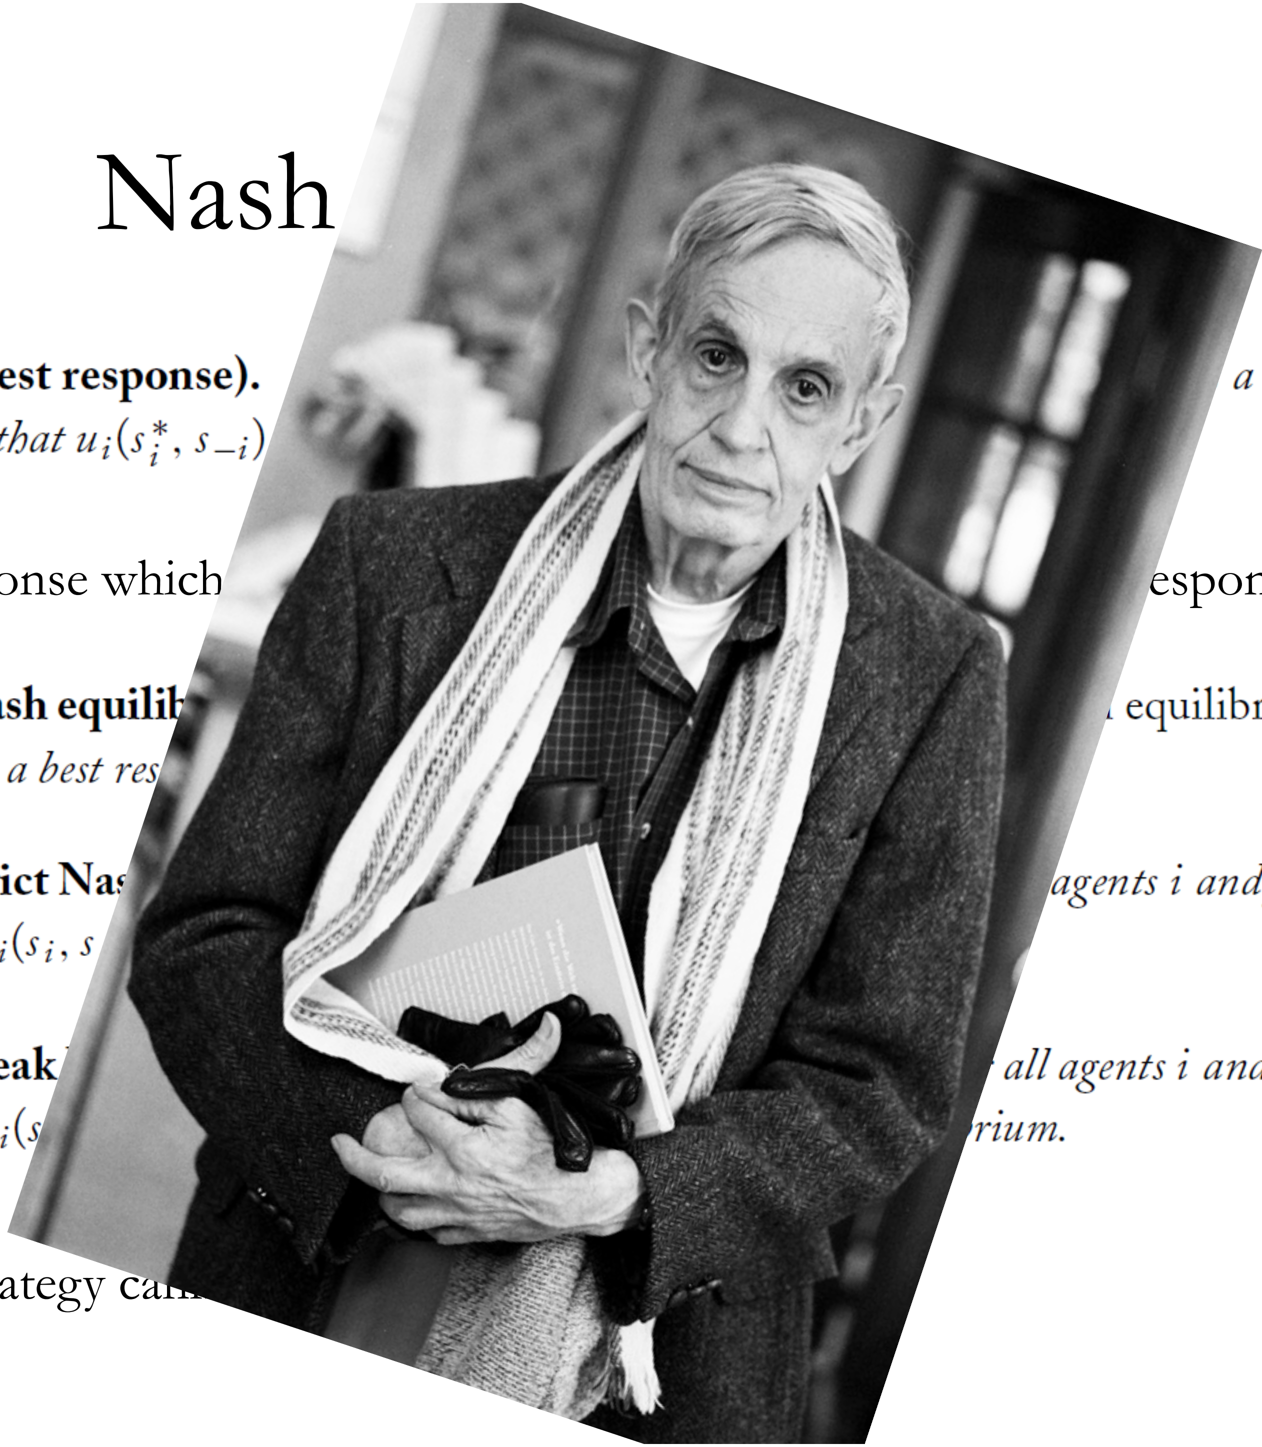

Nash Equilibrium example

	LW	WL
LW	2, 1	0, 0
WL	0, 0	1, 2

Are these the only equilibria?

Unfortunately no…

$$U_{\text{wife}}(\text{LW}) = U_{\text{wife}}(\text{WL})$$

$$2 * p + 0 * (1 - p) = 0 * p + 1 * (1 - p)$$

$$p = \frac{1}{3}$$

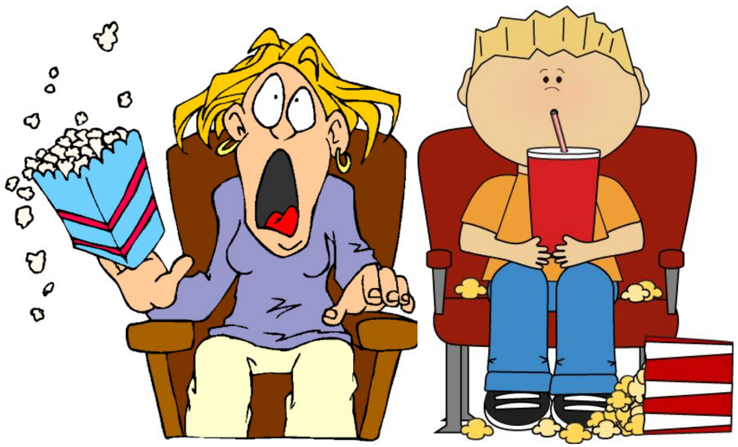

Existence of Nash Equilibrium

Does every game have at least one Nash Equilibrium?

Yes!, now give me a Nobel prize please.

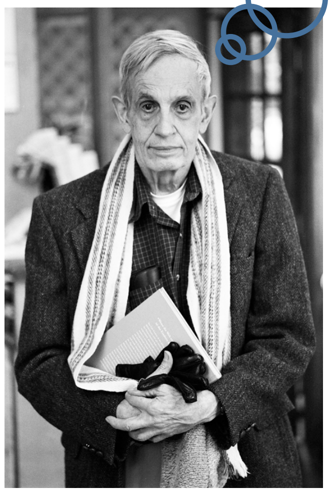

MaxMin & MinMax

What if the other players really don't like me?

Definition 3.1.1 (Maxmin). *The* maxmin strategy *for player* i *is* $\arg\max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$, *and the* maxmin value *for player* i *is* $\max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$.

Definition 3.1.2 (Minmax, two-player). *In a two-player game, the* minmax strategy *for player* i *against player* $-i$ *is* $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$, *and player* $-i$*'s* minmax value *is* $\min_{s_i} \max_{s_{-i}} u_{-i}(s_i, s_{-i})$.

Definition 3.1.3 (Minmax, n-player). *In an n-player game, the* minmax strategy *for player* i *against player* $j \neq i$ *is* i*'s component of the mixed-strategy profile* s_{-j} *in the expression* $\arg\min_{s_{-j}} \max_{s_j} u_j(s_j, s_{-j})$, *where* $-j$ *denotes the set of players other than* j. *As before, the* minmax value *for player* j *is* $\min_{s_{-j}} \max_{s_j} u_j(s_j, s_{-j})$.

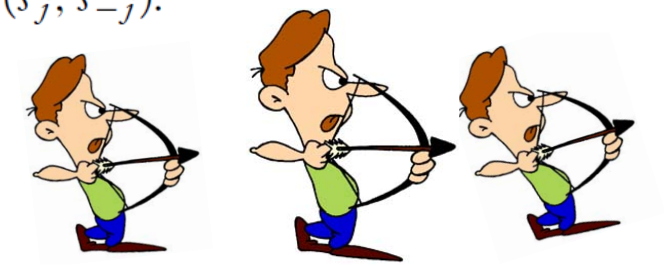

MinMax/MaxMin and Nash

Theorem 3.1.4 (Minimax theorem (von Neumann, 1928)). *In any finite, two-player, zero-sum game, in any Nash equilibrium*[1] *each player receives a payoff that is equal to both his maxmin value and his minmax value.*

Strategy dominance

Definition 3.3.1 (Domination). *Let s_i and s'_i be two strategies of player i, and S_{-i} the set of all strategy profiles of the remaining players. Then*

1. s_i strictly dominates s'_i *if for all $s_{-i} \in S_{-i}$, it is the case that $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$.*
2. s_i weakly dominates s'_i *if for all $s_{-i} \in S_{-i}$, it is the case that $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$, and for at least one $s_{-i} \in S_{-i}$, it is the case that $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$.*
3. s_i very weakly dominates s'_i *if for all $s_{-i} \in S_{-i}$, it is the case that $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.*

If one strategy dominates all others, we say that it is (strongly, weakly or very weakly) *dominant*.

Definition 3.3.2 (Dominant strategy). *A strategy is* strictly (resp., weakly; very weakly) dominant *for an agent if it strictly (weakly; very weakly) dominates any other strategy for that agent.*

Definition 3.3.3 (Dominated strategy). *A strategy s_i is* strictly (weakly; very weakly) dominated *for an agent i if some other strategy s'_i strictly (weakly; very weakly) dominates s_i.*

Strategy dominance

We can remove strictly dominated pure strategies as they will not be a part of any best response.

	L	C	R
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

R is strictly dominated

	L	C
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

M is dominated by mixed strategy of U and D

	L	C
U	3, 1	0, 1
D	0, 1	4, 1

Outlook

Games in Normal Form

➢ **Perfect-information Sequential Actions Games**

Imperfect-information SA Games

Repeated Games

Bayesian Games

Coalitional Games

Perfect-information Game

Definition 4.1.1 (Perfect-information game). *A (finite) perfect-information game (in extensive form) is a tuple* $G = (N, A, H, Z, \chi, \rho, \sigma, u)$*, where:*

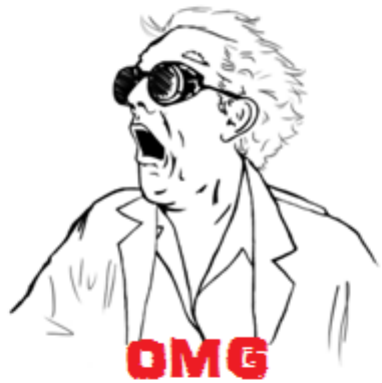

- N *is a set of* n *players;*
- A *is a (single) set of actions;*
- H *is a set of nonterminal choice nodes;*
- Z *is a set of terminal nodes, disjoint from* H*;*
- $\chi : H \mapsto 2^A$ *is the action function, which assigns to each choice node a set of possible actions;*
- $\rho : H \mapsto N$ *is the player function, which assigns to each nonterminal node a player* $i \in N$ *who chooses an action at that node;*
- $\sigma : H \times A \mapsto H \cup Z$ *is the successor function, which maps a choice node and an action to a new choice node or terminal node such that for all* $h_1, h_2 \in H$ *and* $a_1, a_2 \in A$*, if* $\sigma(h_1, a_1) = \sigma(h_2, a_2)$ *then* $h_1 = h_2$ *and* $a_1 = a_2$*; and*
- $u = (u_1, \ldots, u_n)$*, where* $u_i : Z \mapsto \mathbb{R}$ *is a real-valued utility function for player* i *on the terminal nodes* Z*.*

Perfect-information Game

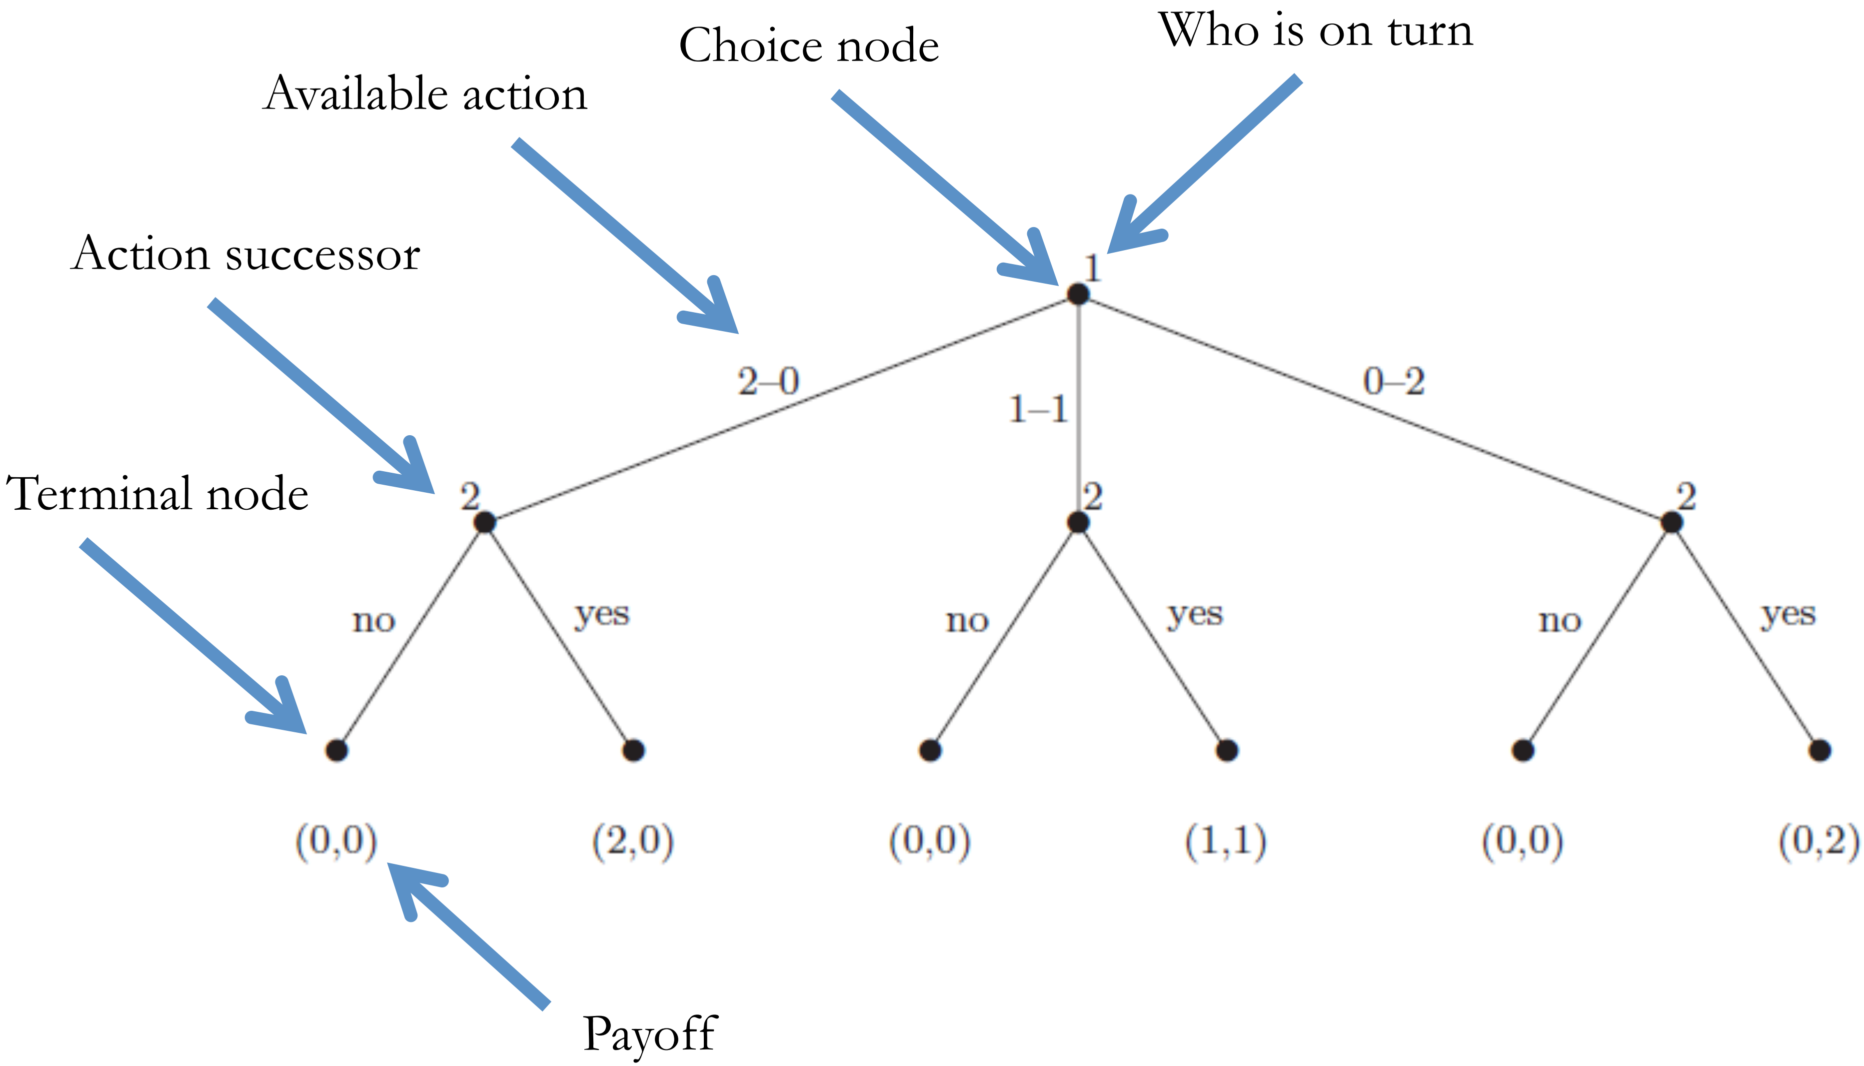

Perfect-information Game

Definition 4.1.1 (Perfect-information game). *A (finite) perfect-information game (in extensive form) is a tuple* $G = (N, A, H, Z, \chi, \rho, \sigma, u)$, *where:*

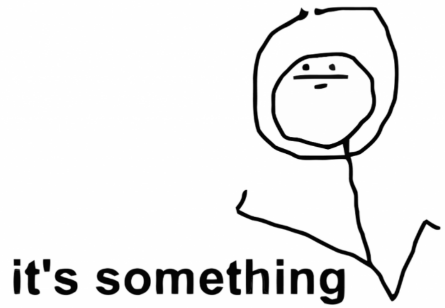

- N *is a set of* n *players;*
- A *is a (single) set of actions;*
- H *is a set of nonterminal choice nodes;*
- Z *is a set of terminal nodes, disjoint from* H;
- $\chi : H \mapsto 2^A$ *is the action function, which assigns to each choice node a set of possible actions;*
- $\rho : H \mapsto N$ *is the player function, which assigns to each nonterminal node a player* $i \in N$ *who chooses an action at that node;*
- $\sigma : H \times A \mapsto H \cup Z$ *is the successor function, which maps a choice node and an action to a new choice node or terminal node such that for all* $h_1, h_2 \in H$ *and* $a_1, a_2 \in A$, *if* $\sigma(h_1, a_1) = \sigma(h_2, a_2)$ *then* $h_1 = h_2$ *and* $a_1 = a_2$; *and*
- $u = (u_1, \ldots, u_n)$, *where* $u_i : Z \mapsto \mathbb{R}$ *is a real-valued utility function for player* i *on the terminal nodes* Z.

Strategies

Definition 4.2.1 (Pure strategies). *Let $G = (N, A, H, Z, \chi, \rho, \sigma, u)$ be a perfect-information extensive-form game. Then the pure strategies of player i consist of the Cartesian product $\prod_{h \in H, \rho(h)=i} \chi(h)$.*

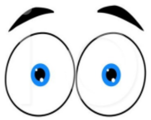

Strategy must define actions even in the unused branches.

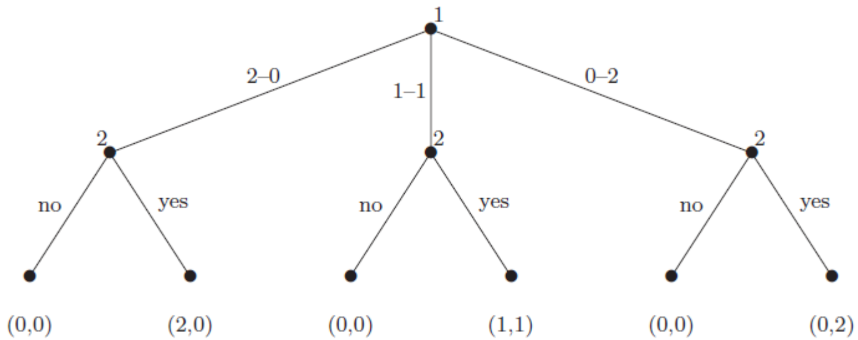

$S_1 = \{2\text{–}0, 1\text{–}1, 0\text{–}2\}$

$S_2 = \{(yes, yes, yes), (yes, yes, no), (yes, no, yes), (yes, no, no), (no, yes, yes), (no, yes, no), (no, no, yes), (no, no, no)\}$

Normal Games vs. PI Games

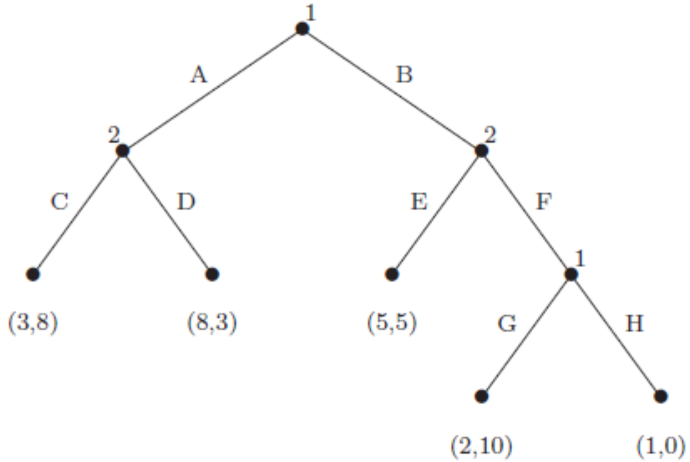

(A,G)				
(A,H)				
(B,G)				
(B,H)				


So this can be done for every Perfect-information Game?


Yes, but…


And what about the reverse?


No! Consider the Prisoner's dilemma.

PI Games Equilibrium

Theorem 4.2.2. *Every (finite) perfect-information game in extensive form has a pure-strategy Nash equilibrium.*

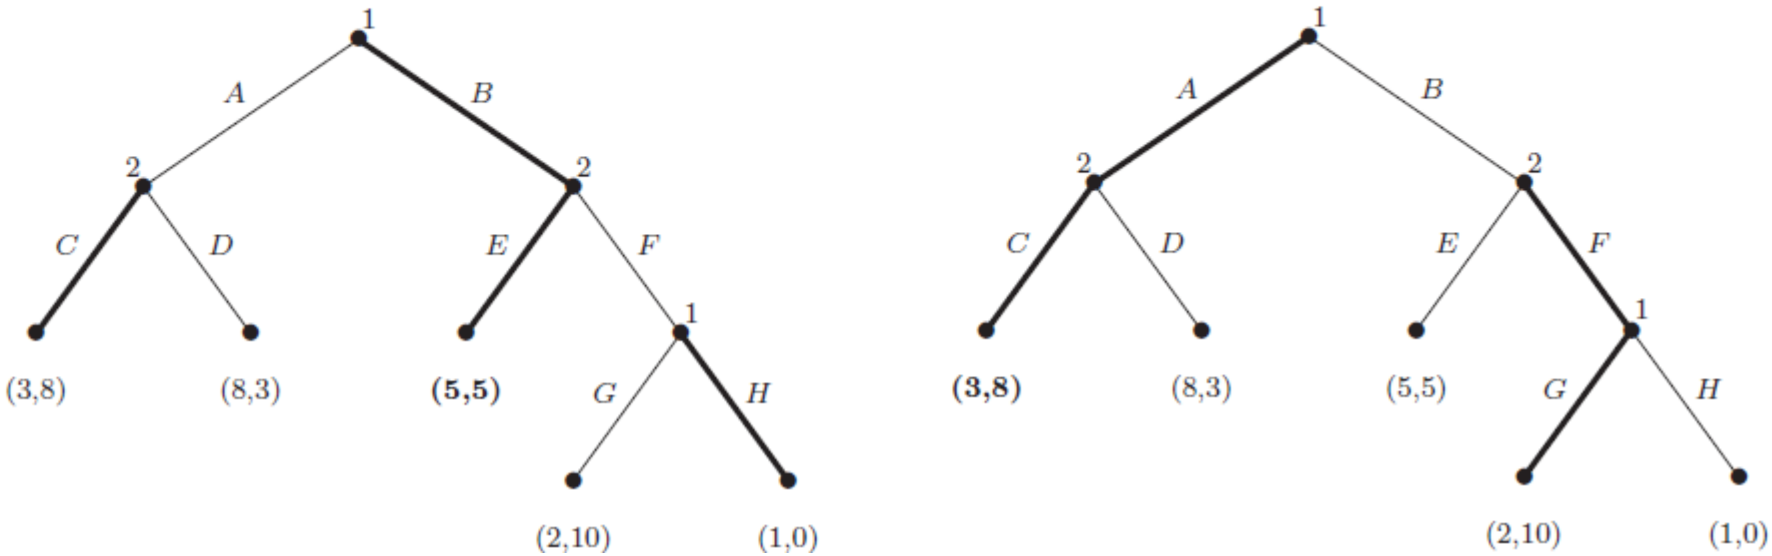

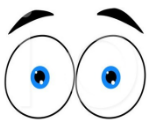
In the 1nd (left) equilibrium, player 1 plays H only as a **threat**.

Perfect-information Subgames

Definition 4.3.1 (Subgame). *Given a perfect-information extensive-form game G, the* subgame *of G rooted at node h is the restriction of G to the descendants of h. The set of subgames of G consists of all of subgames of G rooted at some node in G.*

Definition 4.3.2 (Subgame-perfect equilibrium). *The* subgame-perfect equilibria *(SPE) of a game G are all strategy profiles s such that for any subgame G' of G, the restriction of s to G' is a Nash equilibrium of G'.*

Subgame-perfect equilibrium is also a Nash equilibrium.

The “problematic” equilibrium is not subgame-perfect.

Backward Induction

We can compute the Subgame-perfect equilibrium efficiently.

function BackwardInduction (node h) **returns** $u(h)$
if $h \in Z$ **then**
    return $u(h)$ // h is a terminal node
$best\_util \leftarrow -\infty$
forall $a \in \chi(h)$ **do**
    $util\_at\_child \leftarrow$ BackwardInduction$(\sigma(h, a))$
    if $util\_at\_child_{\rho(h)} > best\_util_{\rho(h)}$ **then**
        $best\_util \leftarrow util\_at\_child$
return $best\_util$

It is actually minimax algorithm for two-player, zero-sum games.

Backward Induction

We have an efficient algorithm… any drawbacks?

Well, games can be big… e.g., chess game will have approx 10 ^150 nodes.

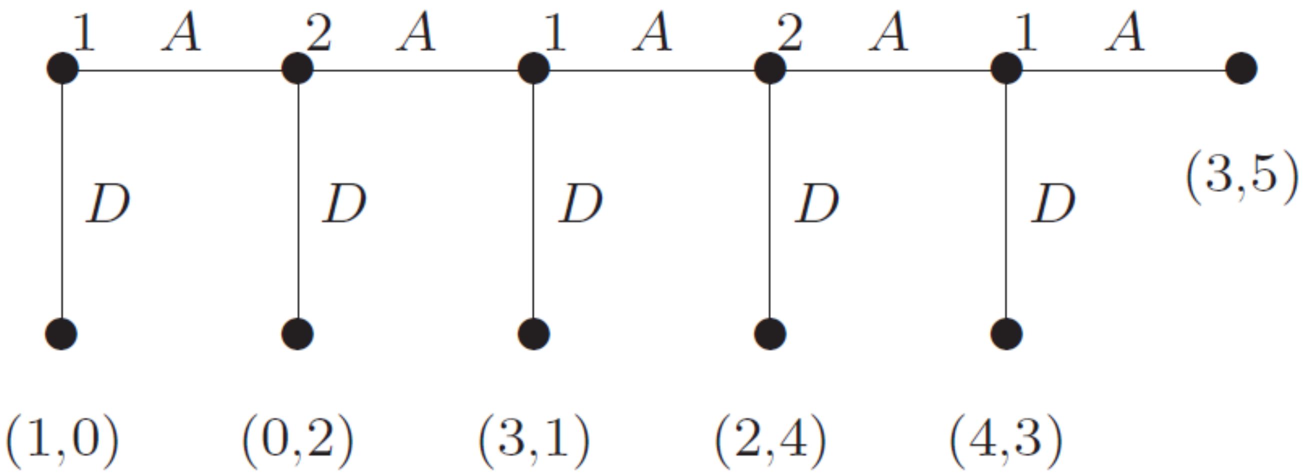

Outlook

Games in Normal Form

Perfect-information Sequential Actions Games

➢ **Imperfect-information SA Games**

Repeated Games

Bayesian Games

Coalitional Games

Imperfect-information Games

Definition 5.1.1 (Imperfect-information game). *An imperfect-information game (in extensive form) is a tuple* $(N, A, H, Z, \chi, \rho, \sigma, u, I)$, *where:*

- $(N, A, H, Z, \chi, \rho, \sigma, u)$ *is a perfect-information extensive-form game; and*
- $I = (I_1, \ldots, I_n)$, *where* $I_i = (I_{i,1}, \ldots, I_{i,k_i})$ *is an equivalence relation on (i.e., a partition of)* $\{h \in H : \rho(h) = i\}$ *with the property that* $\chi(h) = \chi(h')$ *and* $\rho(h) = \rho(h')$ *whenever there exists a* j *for which* $h \in I_{i,j}$ *and* $h' \in I_{i,j}$.

Definition 5.2.1 (Pure strategies). *Let* $G = (N, A, H, Z, \chi, \rho, \sigma, u, I)$ *be an imperfect-information extensive-form game. Then the pure strategies of player* i *consist of the Cartesian product* $\prod_{I_{i,j} \in I_i} \chi(I_{i,j})$.

Imperfect-information Games

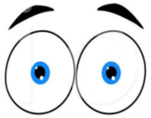 Prisoner's dilemma can be represented in the imperfect-information form.

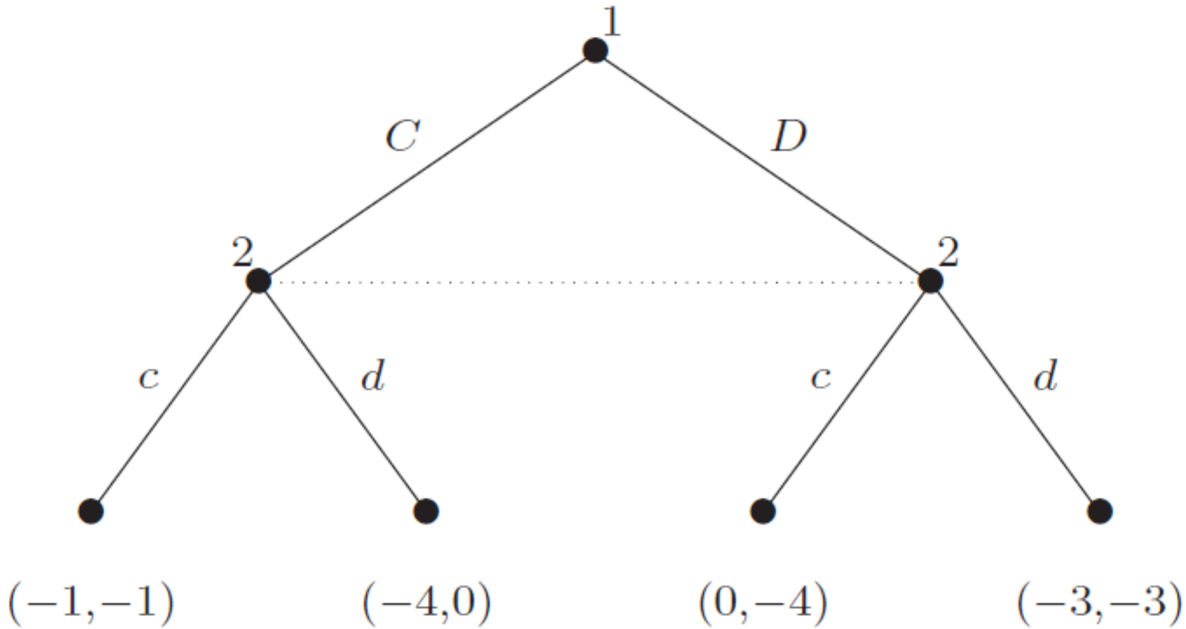

Behavioral vs. Mixed strategies

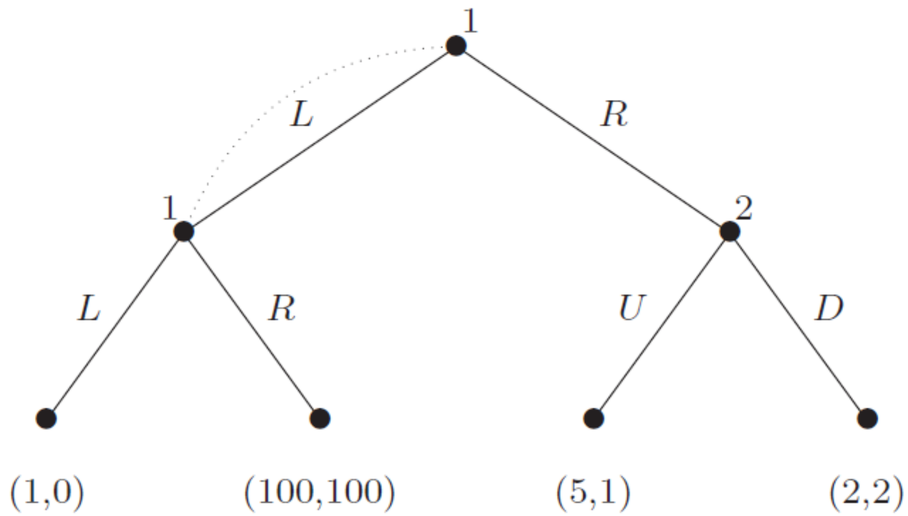

Mixed → Backward Induction → $((0, 1), (0, 1))$

Behavioral → $1 * p^2 + 100 * p(1 - p) + 2 * (1 - p)$ → $((98/198, 100/198), (0, 1))$


Behavioral and Mixed equilibrium sets might be not comparable

Outlook

Games in Normal Form

Perfect-information Sequential Actions Games

Imperfect-information SA Games

➢ **Repeated Games**

Bayesian Games

Coalitional Games

Finitely repeated Games

	C	D
C	−1, −1	−4, 0
D	0, −4	−3, −3

$\Rightarrow$

	C	D
C	−1, −1	−4, 0
D	0, −4	−3, −3

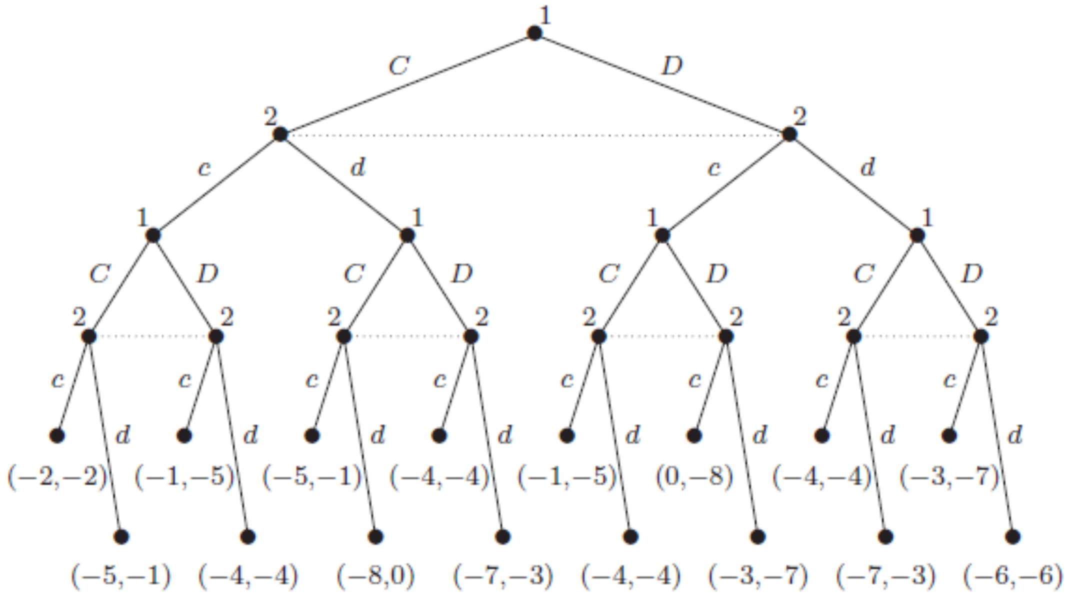

Infinitely repeated Games

What is the proper definition of the payoffs in IR Games?

Definition 6.2.1 (Average reward). *Given an infinite sequence of payoffs* $r_i^{(1)}, r_i^{(2)}, \ldots$ *for player* i, *the* average reward *of* i *is*

$$\lim_{k \to \infty} \frac{\sum_{j=1}^{k} r_i^{(j)}}{k}.$$

Definition 6.2.2 (Discounted reward). *Given an infinite sequence of payoffs* $r_i^{(1)}, r_i^{(2)}, \ldots$ *for player* i, *and a discount factor* β *with* $0 \leq \beta \leq 1$, *the* future discounted reward *of* i *is* $\sum_{j=1}^{\infty} \beta^j r_i^{(j)}$.

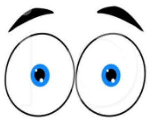

Either we do not care about the future or the game may end at some point.

Folk theorem

So what about the Nash Equilibria in IR Games?

Let us define: $v_i = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_{-i}, s_i)$

Definition 6.2.3 (Enforceable). *A payoff profile* $r = (r_1, r_2, \ldots, r_n)$ *is* enforceable *if* $\forall i \in N$, $r_i \geq v_i$.

Definition 6.2.4 (Feasible). *A payoff profile* $r = (r_1, r_2, \ldots, r_n)$ *is* feasible *if there exist rational, nonnegative values* α_a *such that for all* i, *we can express* r_i *as* $\sum_{a \in A} \alpha_a \, u_i(a)$, *with* $\sum_{a \in A} \alpha_a = 1$.

In other words, a payoff profile is feasible if it is a convex, rational combination of the outcomes in G.

Theorem 6.2.5 (Folk Theorem). *Consider any* n*-player normal-form game* G *and any payoff profile* $r = (r_1, r_2, \ldots, r_n)$.

1. *If* r *is the payoff profile for any Nash equilibrium* s *of the infinitely repeated* G *with average rewards, then for each player* i, r_i *is enforceable.*
2. *If* r *is both feasible and enforceable, then* r *is the payoff profile for some Nash equilibrium of the infinitely repeated* G *with average rewards.*

Outlook

Games in Normal Form

Perfect-information Sequential Actions Games

Imperfect-information SA Games

Repeated Games

➢ **Bayesian Games**

Coalitional Games

Even more uncertainty

We know what to do when we are uncertain about the history; however, what about if we are not sure about the available actions or even which game we are playing?

We can reduce these uncertainties to uncertain payoffs.

	L	R
U	1, 1	1, 3
D	0, 5	1, 13

	L	C	R
U	1, 1	0, 2	1, 3
D	0, 5	2, 8	1, 13

	L	C	R
U	1, 1	0, −100	1, 3
D	0, 5	2, −100	1, 13

Bayesian Games I

Definition 7.1.1 (Bayesian game: information sets). *A* Bayesian game *is a tuple* (N, G, P, I) *where:*

- *N is a set of agents;*
- *G is a set of games with N agents each such that if $g, g' \in G$ then for each agent $i \in N$ the strategy space in g is identical to the strategy space in g';*
- *$P \in \Pi(G)$ is a common prior over games, where $\Pi(G)$ is the set of all probability distributions over G; and*
- *$I = (I_1, ..., I_N)$ is a tuple of partitions of G, one for each agent.*

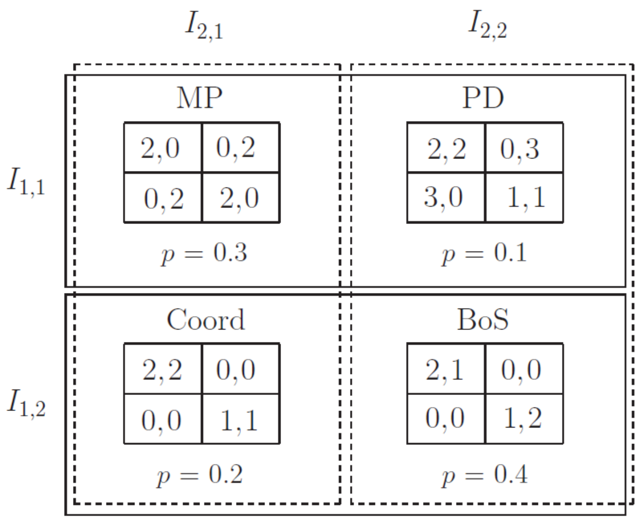

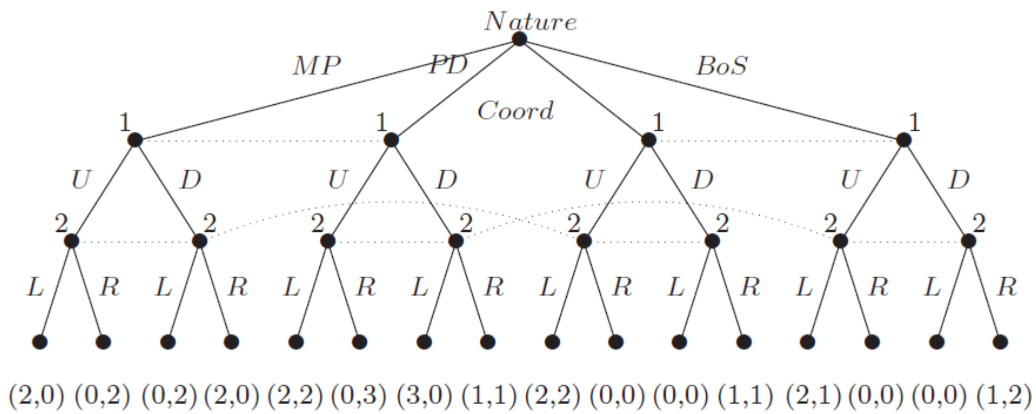

Bayesian Games II

Definition 7.1.2 (Bayesian game: types). *A* Bayesian game *is a tuple* (N, A, Θ, p, u) *where:*

- N *is a set of agents;*
- $A = A_1 \times \cdots \times A_n$, *where* A_i *is the set of actions available to player* i;
- $\Theta = \Theta_1 \times \ldots \times \Theta_n$, *where* Θ_i *is the type space of player* i;
- $p : \Theta \mapsto [0, 1]$ *is a common prior over types; and*
- $u = (u_1, \ldots, u_n)$, *where* $u_i : A \times \Theta \mapsto \mathbb{R}$ *is the utility function for player* i.

Agent knows his own type while does not know types of the others.

ISBG to TBG

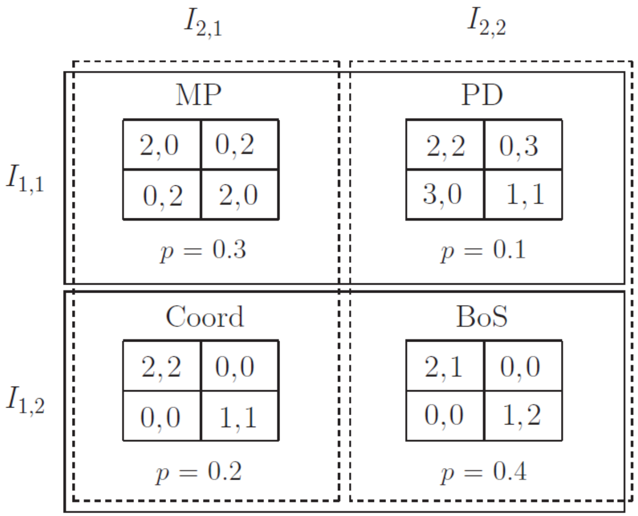

$p(\theta_{1,1}, \theta_{2,1}) = .3$

$p(\theta_{1,1}, \theta_{2,2}) = .1$

$p(\theta_{1,2}, \theta_{2,1}) = .2$

$p(\theta_{1,2}, \theta_{2,2}) = .4$

a_1	a_2	θ_1	θ_2	u_1	u_2
U	L	$\theta_{1,1}$	$\theta_{2,1}$	2	0
U	L	$\theta_{1,1}$	$\theta_{2,2}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,1}$	2	2
U	L	$\theta_{1,2}$	$\theta_{2,2}$	2	1
U	R	$\theta_{1,1}$	$\theta_{2,1}$	0	2
U	R	$\theta_{1,1}$	$\theta_{2,2}$	0	3
U	R	$\theta_{1,2}$	$\theta_{2,1}$	0	0
U	R	$\theta_{1,2}$	$\theta_{2,2}$	0	0

a_1	a_2	θ_1	θ_2	u_1	u_2
D	L	$\theta_{1,1}$	$\theta_{2,1}$	0	2
D	L	$\theta_{1,1}$	$\theta_{2,2}$	3	0
D	L	$\theta_{1,2}$	$\theta_{2,1}$	0	0
D	L	$\theta_{1,2}$	$\theta_{2,2}$	0	0
D	R	$\theta_{1,1}$	$\theta_{2,1}$	2	0
D	R	$\theta_{1,1}$	$\theta_{2,2}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,1}$	1	1
D	R	$\theta_{1,2}$	$\theta_{2,2}$	1	2

Expected Utility in BG

Definition 7.2.1 (*Ex post* expected utility). *Agent i's* ex post *expected utility in a Bayesian game* (N, A, Θ, p, u), *where the agents' strategies are given by s and the agent' types are given by* θ, *is defined as*

$$EU_i(s,\theta) = \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j|\theta_j) \right) u_i(a,\theta). \qquad (7.1)$$

Definition 7.2.2 (*Ex interim* expected utility). *Agent i's* ex interim *expected utility in a Bayesian game* (N, A, Θ, p, u), *where i's type is* θ_i *and where the agents' strategies are given by the mixed-strategy profile s, is defined as*

$$EU_i(s,\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\theta_i) EU_i(s,(\theta_i,\theta_{-i})). \qquad (7.3)$$

Definition 7.2.3 (*Ex ante* expected utility). *Agent i's* ex ante *expected utility in a Bayesian game* (N, A, Θ, p, u), *where the agents' strategies are given by the mixed-strategy profile s, is defined as*

$$EU_i(s) = \sum_{\theta \in \Theta} p(\theta) EU_i(s,\theta) \qquad (7.5)$$

Towards Equilibrium in BG

Definition 7.2.4 (Best response in a Bayesian game). *The set of agent i's* best responses *to mixed-strategy profile s_{-i} are given by*

$$BR_i(s_{-i}) = \arg\max_{s'_i \in S_i} EU_i(s'_i, s_{-i}). \tag{7.7}$$

Definition 7.2.5 (Bayes–Nash equilibrium). *A* Bayes–Nash equilibrium *is a mixed-strategy profile s that satisfies $\forall i \;\; s_i \in BR_i(s_{-i})$.*

Towards Equilibrium in BG

We can reduce BG to NG with *ex ante* expected utility.

$$
\begin{aligned}
u_2(UU, LL) &= \sum_{\theta \in \Theta} p(\theta) u_2(U, L, \theta) \\
&= p(\theta_{1,1}, \theta_{2,1}) u_2(U, L, \theta_{1,1}, \theta_{2,1}) + p(\theta_{1,1}, \theta_{2,2}) u_2(U, L, \theta_{1,1}, \theta_{2,2}) \\
&\quad + p(\theta_{1,2}, \theta_{2,1}) u_2(U, L, \theta_{1,2}, \theta_{2,1}) + p(\theta_{1,2}, \theta_{2,2}) u_2(U, L, \theta_{1,2}, \theta_{2,2}) \\
&= 0.3(0) + 0.1(2) + 0.2(2) + 0.4(1) = 1.
\end{aligned}
$$

	LL	LR	RL	RR
UU	2, 1	1, 0.7	1, 1.2	0, 0.9
UD	0.8, 0.2	1, 1.1	0.4, 1	0.6, 1.9
DU	1.5, 1.4	0.5, 1.1	1.7, 0.4	0.7, 0.1
DD	0.3, 0.6	0.5, 1.5	1.1, 0.2	1.3, 1.1

ex interim after observing certain type

	LL	LR	RL	RR
UU	2, 0.5	1.5, 0.75	0.5, 2	0, 2.25
UD	2, 0.5	1.5, 0.75	0.5, 2	0, 2.25
DU	0.75, 1.5	0.25, 1.75	2.25, 0	1.75, 0.25
DD	0.75, 1.5	0.25, 1.75	2.25, 0	1.75, 0.25

Towards Equilibrium in BG

Is a Nash Equilibrium in *ex-interim* induced game meaningfull?

No! The other agents can not observe the player's type.

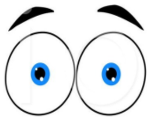

Now, Nash Equilibria in *ex-ante* induced game correspond to the Bayes-Nash Equilibria; thus, there is always at least one Bayes-Nash Equlibrium.

Beware! There are alternative algorithms for finding BN Eq. (e.g. expectimax).

Outlook

Games in Normal Form

Perfect-information Sequential Actions Games

Imperfect-information SA Games

Repeated Games

Bayesian Games

➢ **Coalitional Games**

Coalitional Games

Definition 8.1.1 (Coalitional game with transferable utility). *A* coalitional game with transferable utility *is a pair* (N, v), *where*

- N *is a finite*[2] *set of players, indexed by* i; *and*
- $v : 2^N \mapsto \mathbb{R}$ *associates with each coalition* $S \subseteq N$ *a real-valued payoff* $v(S)$ *that the coalition's members can distribute among themselves. The function* v *is also called the*

Which coalition will form?

How should that coalition divide its payoff among its members?

Example 8.1.2 (Voting game). A parliament is made up of four political parties, A, B, C, and D, which have 45, 25, 15, and 15 representatives, respectively. They are to vote on whether to pass a \$100 million spending bill and how much of this amount should be controlled by each of the parties. A majority vote, that is, a minimum of 51 votes, is required in order to pass any legislation, and if the bill does not pass then every party gets zero to spend.

Game Types

Definition 8.2.1 (Superadditive game). *A game* $G = (N, v)$ *is* superadditive *if for all* $S, T \subset N$, *if* $S \cap T = \emptyset$, *then* $v(S \cup T) \geq v(S) + v(T)$.

Definition 8.2.2 (Additive game). *A game* $G = (N, v)$ *is* additive *(or* inessential*) if for all* $S, T \subset N$, *if* $S \cap T = \emptyset$, *then* $v(S \cup T) = v(S) + v(T)$.

Definition 8.2.4 (Convex game). *A game* $G = (N, v)$ *is* convex *if for all* $S, T \subset N$, $v(S \cup T) \geq v(S) + v(T) - v(S \cap T)$.

Definition 8.2.5 (Simple game). *A game* $G = (N, v)$ *is* simple *if for all* $S \subset N$, $v(S) \in \{0, 1\}$.

Definition 8.2.3 (Constant-sum game). *A game* $G = (N, v)$ *is* constant sum *if for all* $S \subset N$, $v(S) + v(N \setminus S) = v(N)$.

Super-additive ⊃ Convex ⊃
Additive
Constant-sum ⊃ ⊃
Proper-simple
Simple ⊃

Some extra definitions

Definition 8.3.1 (Feasible payoff). *Given a coalitional game* (N, v), *the* feasible payoff set *is defined as* $\{x \in \mathbb{R}^N \mid \sum_{i \in N} x_i \leq v(N)\}$.

Definition 8.3.2 (Pre-imputation). *Given a coalitional game* (N, v), *the* pre-imputation set, *denoted* $\mathcal{P}$, *is defined as* $\{x \in \mathbb{R}^N \mid \sum_{i \in N} x_i = v(N)\}$.

Definition 8.3.3 (Imputation). *Given a coalitional game* (N, v), *the* imputation set, $\mathcal{C}$, *is defined as* $\{x \in \mathcal{P} \mid \forall i \in N, x_i \geq v(i)\}$.

How to divide payoffs?

How to define "fair" division?

Axiom 8.3.4 (Symmetry). *For any v, if i and j are interchangeable then $\psi_i(N, v) = \psi_j(N, v)$.*

Axiom 8.3.5 (Dummy player). *For any v, if i is a dummy player then $\psi_i(N, v) = v(\{i\})$.*

Axiom 8.3.6 (Additivity). *For any two v_1 and v_2, we have for any player i that $\psi_i(N, v_1 + v_2) = \psi_i(N, v_1) + \psi_i(N, v_2)$, where the game $(N, v_1 + v_2)$ is defined by $(v_1 + v_2)(S) = v_1(S) + v_2(S)$ for every coalition S.*

Theorem 8.3.7. *Given a coalitional game (N, v), there is a unique pre-imputation $\phi(N, v) = \phi(N, v)$ that satisfies the Symmetry, Dummy player, Additivity axioms.*

Shapley value

Definition 8.3.8 (Shapley value). *Given a coalitional game* (N, v), *the* Shapley value *of player i is given by*

$$\phi_i(N, v) = \frac{1}{N!} \sum_{S \subseteq N \setminus \{i\}} |S|!(|N| - |S| - 1)! \Big[v(S \cup \{i\}) - v(S) \Big].$$

Example 8.1.2 (Voting game). A parliament is made up of four political parties, A, B, C, and D, which have 45, 25, 15, and 15 representatives, respectively. They are to vote on whether to pass a \$100 million spending bill and how much of this amount should be controlled by each of the parties. A majority vote, that is, a minimum of 51 votes, is required in order to pass any legislation, and if the bill does not pass then every party gets zero to spend.

$$\begin{aligned}
\phi_A &= (3)\frac{(4-1)!(2-1)!}{4!}(100-0) + (3)\frac{(4-3)!(3-1)!}{4!}(100-0) \\
&\quad + (1)\frac{(4-4)!(4-1)!}{4!}(100-100) \\
&= (3)\frac{2}{24}(100) + (3)\frac{(1)(2)}{24}(00-0) + 0 \\
&= 25 + 25 = \$50 \text{ million.}
\end{aligned}$$

$$\begin{aligned}
\phi_B &= \frac{(4-2)!(2-1)!}{4!}(100-0) + \frac{(4-3)!(3-1)!}{4!}(100-0) \\
&= \frac{2}{24}(100) + \frac{2}{24}(100-0) \\
&= 8.33 + 8.33 = \$16.66 \text{ million.}
\end{aligned}$$

Core

Definition 8.3.9 (Core). *A payoff vector x is in the* core *of a coalitional game* (N, v) *if and only if*

$$\forall S \subseteq N, \sum_{i \in S} x_i \geq v(S).$$

Is core always nonempty or unique?

No to both…

Theorem 8.3.10. *Every constant-sum game that is not additive has an empty core.*

We say that a player i is a *veto player* if $v(N \setminus \{i\}) = 0$.

Theorem 8.3.11. *In a simple game the core is empty iff there is no veto player. If there are veto players, the core consists of all payoff vectors in which the nonveto players get zero.*

Theorem 8.3.12. *Every convex game has a nonempty core.*

Theorem 8.3.13. *In every convex game, the Shapley value is in the core.*

That was theory, what about reality?

Split or Steal?

St. Petersburg paradox

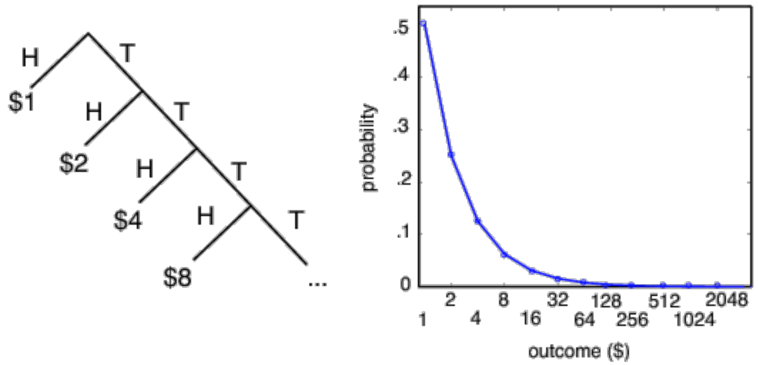

Thank you for your attention