Game Theory Basics, continued.

B

A

	party	stay home
party	$2, 2$	$-1, 0$
stay home	$0, -1$	$0, 0$

Definition 2.2. A strategy s^* for player i is a **best response** to the strategies s_{-i} of others if it maximizes i's utility/payoff. That is

$$u_i(s^*, s_{-i}) \geq u_i(s, s_{-i})$$

for all $s \in S_i$.

Definition 2.8. A strategy profile $(s_1, \ldots, s_n)$ is a Nash equilibrium if for every i, s_i is a best response to s_{-i}.

Inspector

Parker

	don't inspect	inspect
legal (p)	$(0, 0)$	$(0, -1)$
illegal ($1-p$)	$(10, -10)$	$(-90, -6)$

$$\boxed{-10(1-p)} = \boxed{-p - 6(1-p)}$$

these are equal when $p = \frac{4}{5}$

	don't inspect (q)	inspect ($1-q$)
legal	$(0, 0)$	$(0, -1)$
illegal	$(10, -10)$	$(-90, -6)$

0

$10q - 90(1-q)$

these are equal when $q = \frac{9}{10}$

	don't inspect ($9/10$)	inspect ($1/10$)
$4/5$ legal	$(0, 0)$	$(0, -1)$
$1/5$ illegal	$(10, -10)$	$(-90, -6)$

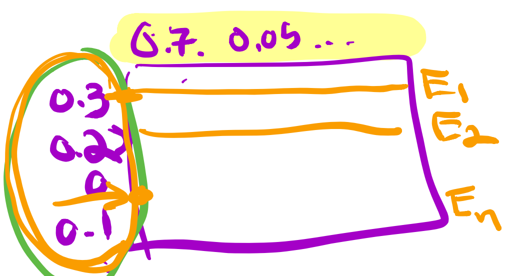

$x_R(\text{legal}) = \frac{4}{5}$
$x_R(\text{illegal}) = \frac{1}{5}$

$x_i(s) = \Pr(\text{player } i \text{ plays strategy } s)$

x_i

Definition 4.1. A **mixed strategy** is a probability distribution over pure strategies.

Definition 4.2. A strategy profile $\boxed{(x_1, \ldots, x_n)}$ where each x_i is possibly a mixed strategy $x_i : S_i \to$ probability distribution and $\sum_{s \in S_i} x_i(s) = 1$ is a (mixed) Nash equilibrium if for each i,

$$\sum_{s \in S_i, s_{-i} \in S_{-i}} x_i(s) \operatorname{Prob}(s_{-i}) u_i(s, s_{-i}) \geq$$

$$\sum_{s_{-i} \in S_{-i}} \operatorname{Prob}(s_{-i}) u_i(s_i, s_{-i}).$$

Exp payoff to player i when she plays x_i, x_{-i}

$\forall s_i \in S_i$

Exp payoff to player i when she plays $s_i \in S_i$, x_{-i}

	party	stay home
party	2, 2	−1, 0
stay home	0, −1	0, 0

p

$1-p$

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

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

(P_a, P_a)

(SH, SH)

$([\frac{1}{3},\frac{2}{3}], [\frac{1}{3},\frac{2}{3}])$

Fact 4.4. **Indifference principle:** If a mixed strategy of player i in a Nash equilibrium randomizes over a set of pure strategies $T_i \subset S_i$, then the expected payoff to the player from each pure strategy in T_i must be the same. And the payoff from any other strategy must be at most this high.

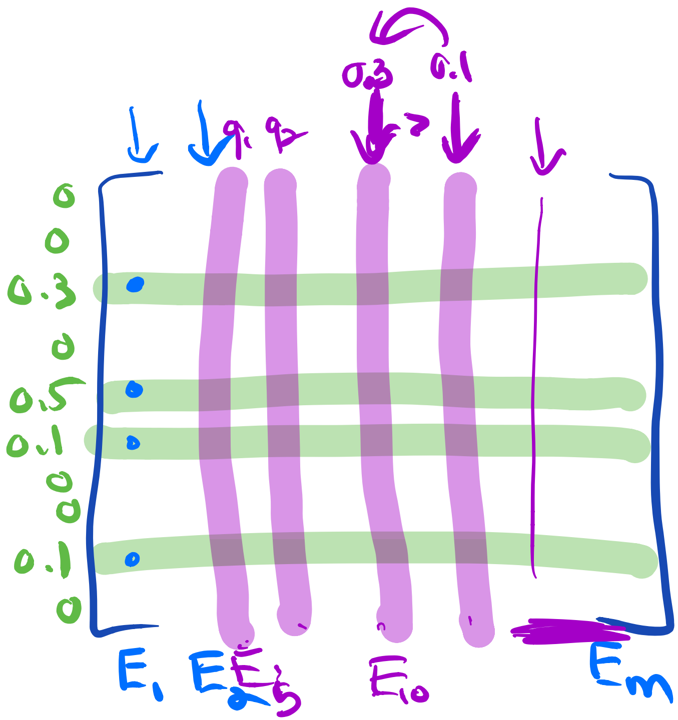

col player puts pos prob on cols $T_i \subseteq S_i$

$$\boxed{\begin{array}{l} E_i = c \quad \forall i \in T_i \\ E_i \le c \quad \forall i \notin T_i \end{array}}$$

Summary so far

- A Nash equilibrium is a set of stable (possibly mixed) strategies.
- Stable means that no player has an incentive to deviate given what the other players are doing.
- Pure equilibrium: there may be none, unique or multiple. Can be identified with "best response diagrams".
- A joint mixed strategy for n players:
  - A probability distribution for each player (possibly different)
- It is an equilibrium if
  - For each player, their distribution is a best response to the others.
  - Only consider unilateral deviations.
  - Everyone knows all the distributions (but not the outcomes of the coin flips).
- Nash's famous theorem: every game has a mixed strategy equilibrium.

Issues

- Does not suggest how players might choose between different equilibria
- Does not suggest how players might learn to play equilibrium.
- Does not allow for bargains, side payments, threats, collusions, "pre-play" communication.

- Computing Nash equilibria for large games is computationally difficult.

"if your laptop can't find it,
then how should the market"
(players/agents)

Other issues

- Relies on assumptions that might be violated in the real world
  - Rationality is common knowledge.
  - Agents are computationally unbounded.
  - Agents have full information about other players, payoffs, etc.

  - behavioral questions

Zero-sum games

Goalie

Kicker

	L	R
L	0.5	1
R	0.9	0.8

payoffs to row player (Kicker)

	L	R
L	(0.5, -0.5)	(1, -1)
R	(0.9, -0.9)	...

Penalty Kicks:

Row player is **Kicker**: chooses to kick either to left or right of other player.

Column player is **Goalie**: simultaneously, chooses to dive left or right.

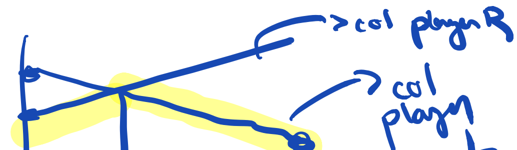

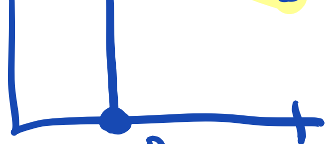

Penalty Kicks

		Goalie q	$1-q$
		L	R
Kicker p	L	0.5	1
$1-p$	R	0.9	0.8

Row player 1'st
she will go R
0.8

Goalie 1st
goes L
0.9

Payoff to Kicker higher
if go 2nd

Kicker goes first (announces p first)

$$V_I = \max_p \left[\min \left(\underbrace{\overset{0.9-0.4p}{0.5p+0.9(1-p)}}_{L},\ \underbrace{\overset{0.8+0.2p}{p+0.8(1-p)}}_{R} \right) \right]$$

$p^* \leftarrow$ chose p to maximize worst case payoff.

$$V_{II} = \min_q \left[\max \left(0.5q + (1-q),\ 0.9q + 0.8(1-q) \right) \right]$$

$q^* \leftarrow$

player I can certainly guarantee himself an exp pay $\geq V_I$

player II can guarantee an exp loss $\leq V_{II}$

$V_I \leq V_{II}$

$\boxed{V_I = V_{II}}$

Theorem 4.5 (John von Neumann, 1928). *Let* V_1 *be the expected gain that player I (maximizer) can guarantee herself in the worst case, and let* p^* *be the mixed strategy that achieves* V_1. *Let* V_2 *be the* ~~*lowest*~~ *expected loss that player II (minimizer) can limit his loss to in the worst case and let* q^* *be the mixed strategy that achieves* V_2.

Then for any 2-player, zero sum game $V_1 = V_2 = V$ *(called the minimax value of the game) and* (p^*, q^*) *is a Nash equilibrium.*

$\boxed{V = V_1 = V_2}$

value of the game.

If player I plays p^* guarantees exp pay $\geq V = V_1$

II q^* exp loss $\leq V = V_2$

Summary – zero-sum games

- Zero-sum games have a "value".
- Optimal strategies are well-defined.
- Maximizer can guarantee a gain of at least V by playing p*
- Minimizer can guarantee a loss of at most V by playing q*.
- This is a Nash equilibrium.
- In contrast to general-sum games, optimal strategies in zero-sum games can be computed efficiently (using linear programming).

Actual data   1500 penalty kicks.

0.423 0.577 real

0.42 0.58 optimal.

		kicker		L	R
0.4	0.38		L	0.58	0.95
0.6	0.62		R	0.93	0.7

Extensive Form Games

Definition 5.1. A k-player finite **extensive-form game** is defined by a finite, rooted tree T.

- Each node in T represents a possible state in the game, with leaves representing terminal states.
- Each internal (nonleaf) node v in T is associated with one of the players, indicating that it is his turn to play if/when v is reached.
- The edges from an internal node to its children are labeled with **actions**, the possible moves the corresponding player can choose from when the game reaches that state.
- Each leaf/terminal state results in a certain payoff for each player.

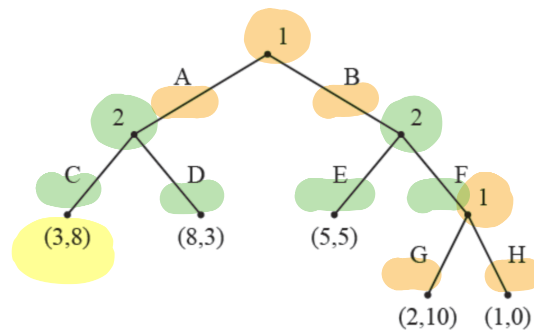

Tic-Tac-Toe

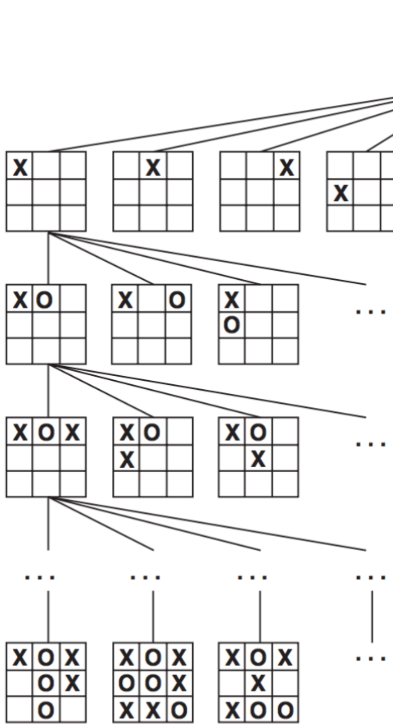

A **pure strategy** for a player in an extensive-form game specifies an action to be taken at each of that player's nodes.

A **mixed strategy** is a probability distribution over pure strategies.

Extensive-form games with perfect information

- When moving, each player is aware of all previous moves (perfect information).
- A (pure) strategy for player i is a mapping from player i's nodes to actions.
- Nash equilibrium, as before.
- In finite, perfect info game, can find one by backwards induction.

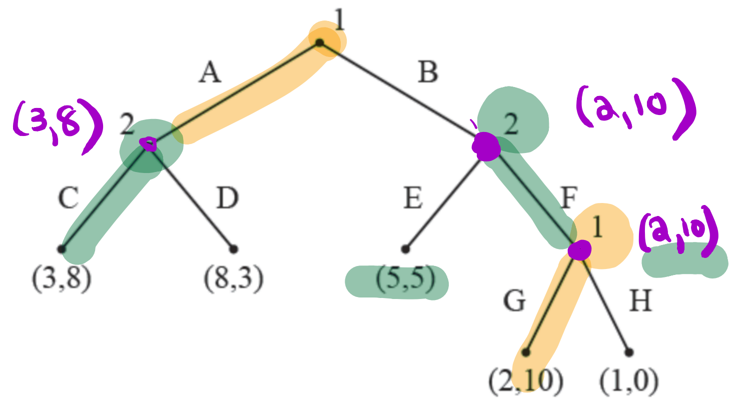

subgame perfect equilibrium

NE in each subtree.

Conversion to normal form

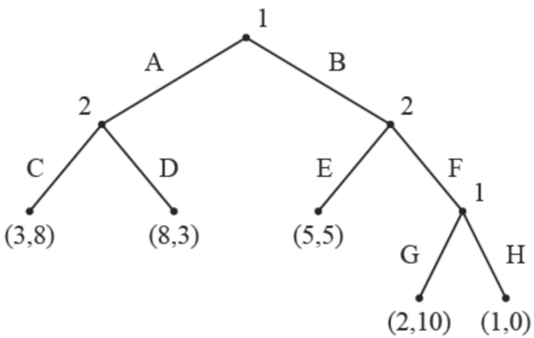

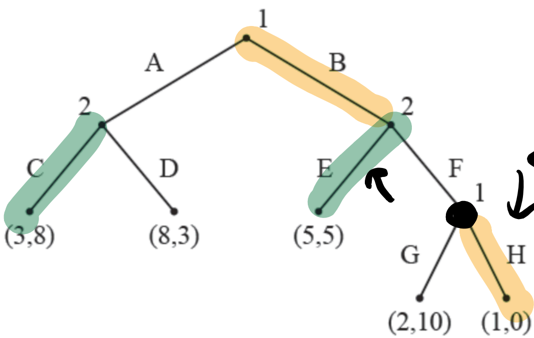

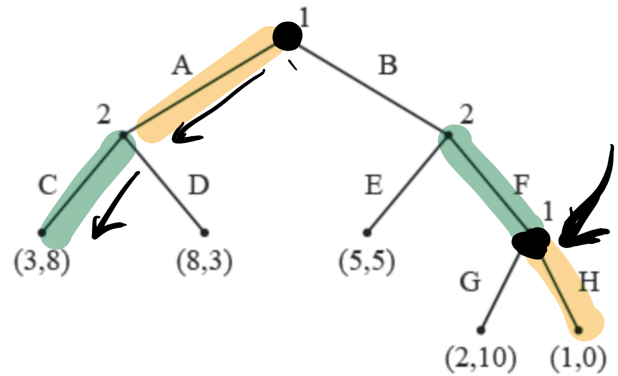

pure strategies for each agent:

	CE	CF	DE	DF
AG	$3,8$	$3,8$	$8,3$	$8,3$
AH	$3,8$	$3,8$	$8,3$	$8,3$
BG	$5,5$	$2,10$	$5,5$	$2,10$
BH	$5,5$	$1,0$	$5,5$	$1,0$

Mutual Assured Destruction

- Two countries, A and B, each possess nuclear weapons
- A is aggressive, B is benign
- A chooses between two options:
  - Escalate arms race
  - Do nothing/maintain the peace.
- If A escalates, then B has two options:
  - Retaliate
  - Back down

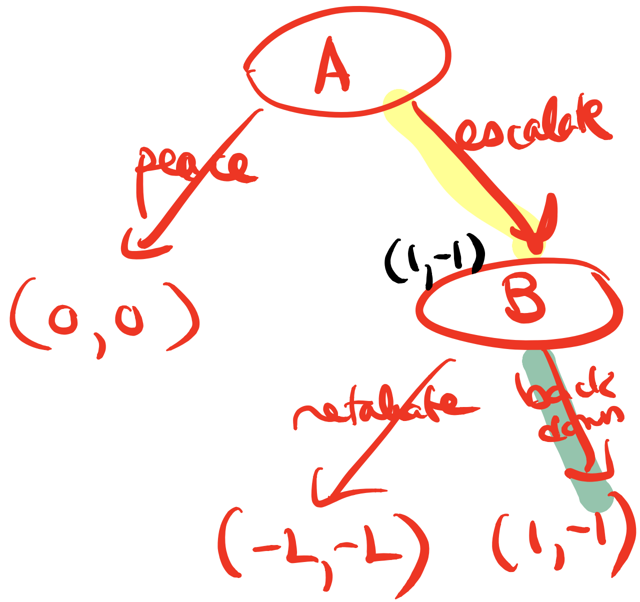

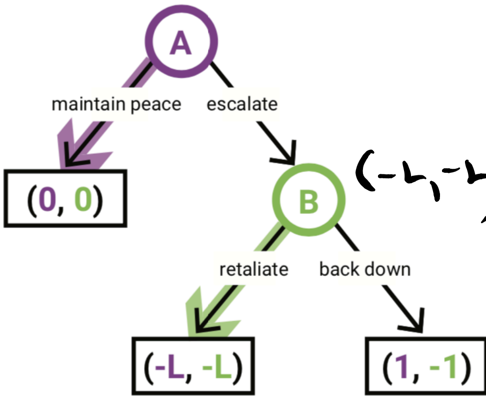

(-L,-L)

FIGURE 6.3. In the MAD game, (maintain peace, escalate) is a Nash equilibrium which is not subgame-perfect.

A **pure strategy** for a player in an extensive-form game specifies an action to be taken at each of that player's nodes.

A **mixed strategy** is a probability distribution over pure strategies.

The kind of equilibrium that is computed by backward induction is called a **subgame-perfect equilibrium** because the behavior in each **subgame**, is also an equilibrium.

Centipede: Pot of money that starts out with $4, and increases by $1 each round.

Two players take turns: The player whose turn it is can split the pot in his favor (and end the game) or allow the game to continue.

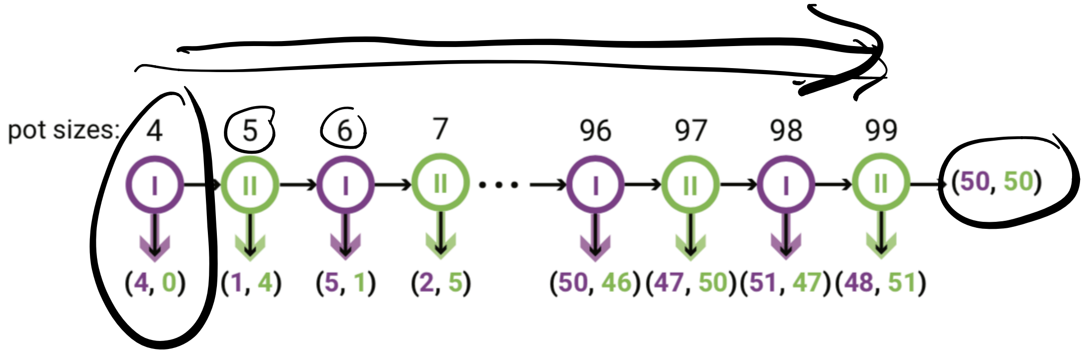

FIGURE 6.4. The top part of the figure shows the game and the resulting payoffs at each leaf. At each node, the "greedy" strategy consists of following the downward arrow, and the "continue" strategy is represented by the arrow to the right. Backward induction from the node with pot-size 99 shows that at each step the player is better off being greedy.

uncertainty about game