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.

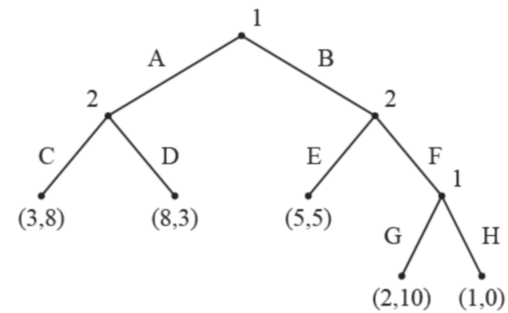

Extensive-form games with perfect information

- 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.

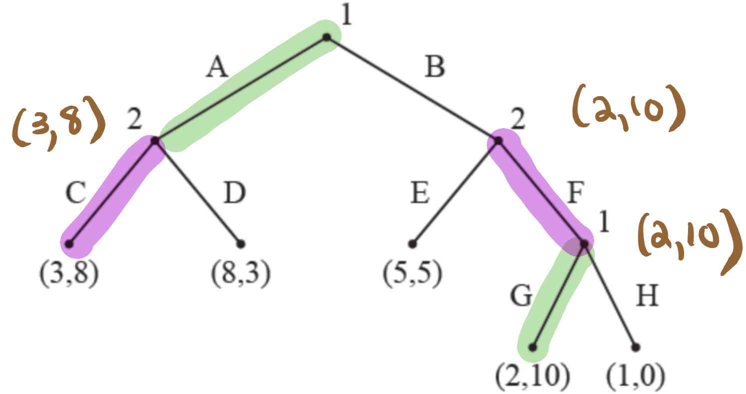

Centipede: Pot of money that starts out with $4, and increases by $1 each round.

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.

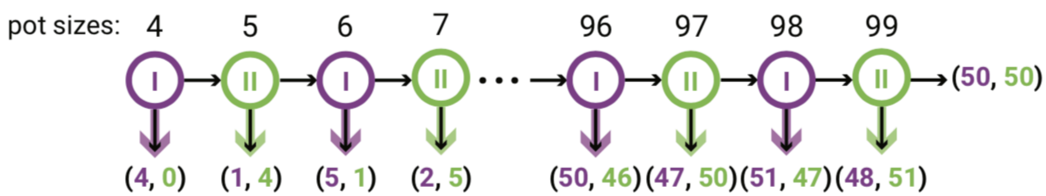

FIGURE 6.4. The top part of the figure shows the game and the resulting payoffs at each leaf. At each node, the "greedy" strategy consists of following the downward arrow, and the "continue" strategy is represented by the arrow to the right. Backward induction from the node with pot-size 99 shows that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

Finite games of perfect information

- At all times, a player knows the history of previous moves and hence current state
- For each possible sequence of actions, each player knows what payoffs each player will get.
- Any such game has a subgame-perfect Nash equilibrium which can be computed by backwards induction.

Checking that a strategy profile is a subgame-perfect equilibrium

- A **single deviation** from strategy s_i is a strategy s_i' that differs from s_i in the action prescribed by a single node in the game tree.
- A single deviation is **useful** if in the play from the subgame defined by that node, agent i's utility in s_i' is strictly better than in s_i, fixing all the others.

Lemma A strategy profile is a subgame-perfect equilibrium in a finite extensive-form game if and only if there is no useful single deviation.

Proof not SGP ⇒ useful single deviation.

improving deviation w/ smallest # changes.

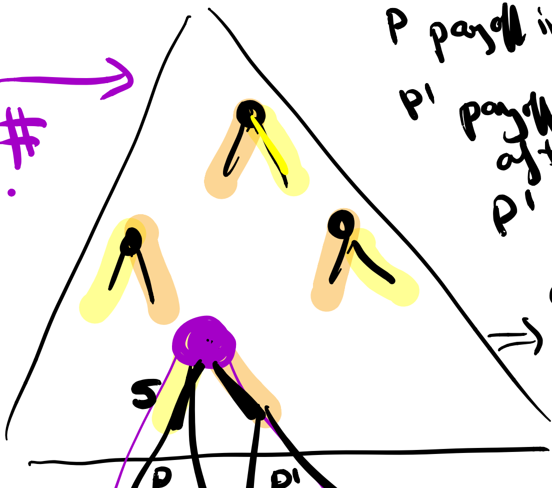

P payoff in this subtree originally

P' payoff in subtree after deviation

P' > P

deviation at ⇒ single node useful.

Other kinds of extensive form games

- Imperfect information (player may not know what node in the tree she is at)
- Incomplete information (number of players, moves available, payoffs)
- Moves by nature.

poker

node
- history of play
- whose turn it is.
- what cards each player has

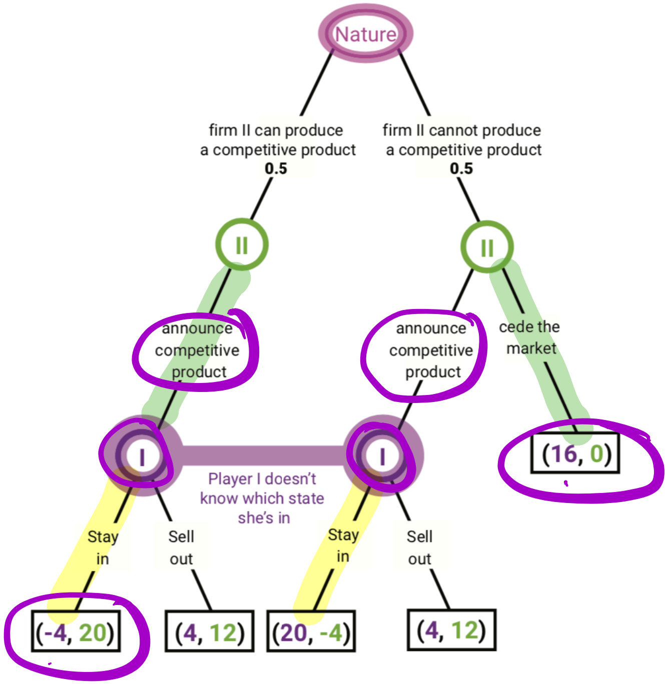

Player I: Startup

Player II: Large Company

I announces new technology threatening II's business.

II has a large research and development group so may be able to pull together competitive product. Prob it can 1/2

Regardless may announce competitive product, to intimidate startup and motivate it to accept buyout offer.

$0.5(-4,20) + 0.5(16,0)$

		player II	
		announce/cede	announce/announce
player I	stay in (I)	$(6, 10)$	$(8, 8)$
	sell out (O)	$(10, 6)$	$(4, 12)$

Repeated Games

		player II	
		cooperate (C)	defect (D)
player I	cooperate (C)	(6, 6)	(0, 8)
	defect (D)	(8, 0)	(2, 2)

- One-round game (e.g. PD) is played repeatedly for some number of rounds?
- What are the equilibria if it's played for n rounds? 100

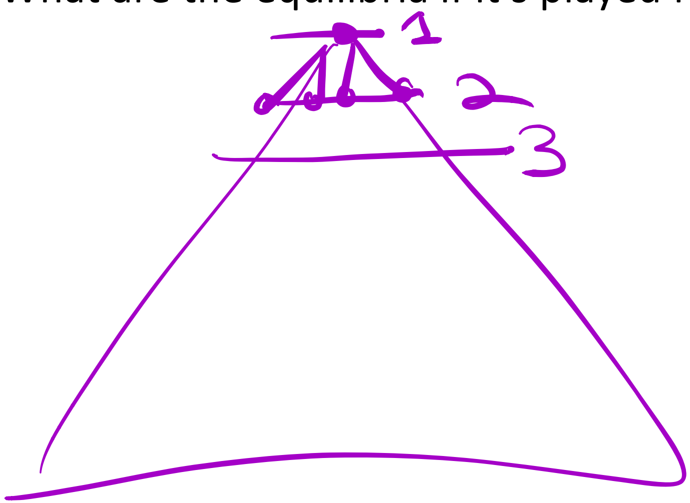

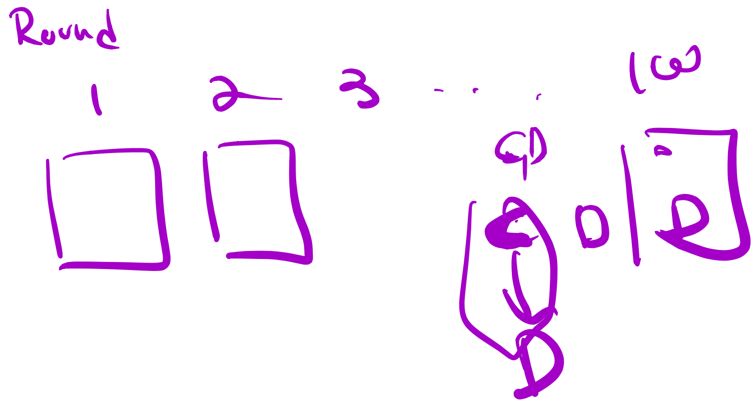

Repeated Games

		player II	
		cooperate (C)	defect (D)
player I	cooperate (C)	(6, 6)	(0, 8)
	defect (D)	(8, 0)	(2, 2)

- One-round game (e.g. PD) is played repeatedly for some number of rounds?
- What if we consider the **discounted payoff**?

$$\boxed{\sum_{t=1}^{\infty} \beta^t \text{ (payoff in round t)}}$$

$0 < \beta < 1$

Interpretations:

① in each round, stop with probability $1-\beta$

② discount future rewards.

		player II	
		cooperate (C)	defect (D)
player I	cooperate (C)	(6, 6)	(0, 8)
	defect (D)	(8, 0)	(2, 2)

Grim Trigger:

- Cooperate until a round in which the other player defects.
- Then defect from that point on.

					t	$t+1$	…	
Player I Grim Trigger	C	C	C	…	C	D	D	D
Player II deviates	C	C	C	…	D	D	D	D

Payoff vector: (6, 6) … (0, 8) (2, 2)

versus (6, 6) (6, 6) … with no deviation

From t on for II (deviating)

$$\boxed{\leq 8\beta^t + \sum_{j>t} \beta^j \cdot 2 < \sum_{j\geq t} \beta^j \cdot 6}$$

worse to deviate.

$$\underbrace{8 + \sum_{j=1}^{\infty} \beta^j \cdot 2}_{6 + \sum_{j=0}^{\infty} 2\beta^j} \overset{?}{<} \sum_{j=0}^{\infty} 6\beta^j$$

$$4 \underbrace{\sum_{j=0}^{\infty} \beta^j}_{\frac{1}{1-\beta}} > 6$$

$$4 > 6(1-\beta)$$

$$6\beta > 2$$

$\beta > \frac{1}{3}$

player I \ player II	cooperate (C)	defect (D)
cooperate (C)	$(6,6)$	$(0,8)$
defect (D)	$(8,0)$	$(2,2)$

Tit-for-tat:

- Cooperate in round 1.
- For every round k > 1, play what the opponent played in round k-1.

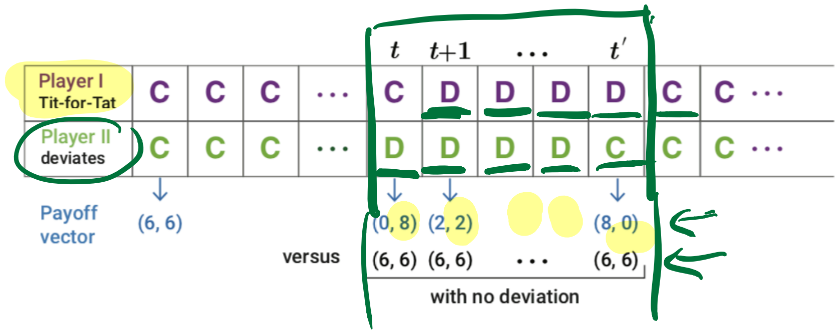

FIGURE 6.14. Illustration of deviation in Tit-for-Tat strategies

$\beta > \frac{1}{3}$

Q: for what β is

$$\sum_{j=t}^{t'} \beta^j 6 > 8\beta^t + \sum_{j=t+1}^{t'-1} \beta^j 2$$

Axelrod's Tournaments

- Robert Axelrod ran a tournament for computer programs playing repeated PD.
- 15 entrants, 200 rounds.
- The simplest of these, Tit-for-Tat won.

62 entrants. TfT

Applications

"Incentives do not build robustness in BitTorrent"

- Recall P2P file sharing
  - Fundamental problem: tendency of users to free ride – consume resources without contributing anything.
  - BitTorrent protocol for file sharing inspired by Tit for Tat.
  - Files broken up into pieces => think of transfer as repeated prisoner's dilemma.
  - In each round, protocol specifies that the peers a user should upload to are those from whom he has downloaded the most data from recently.

- Repeated PD also used to model what's going on in reputation systems. (See next homework).

BitTyrant.