CMSC 474, Introduction to Game Theory

12. Perfect-Information Extensive Form Games

Mohammad T. Hajiaghayi

University of Maryland

The Sharing Game

- Suppose agents 1 and 2 are two children
- Someone offers them two cookies, but only if they can agree how to share them
- Agent 1 chooses one of the following options:
  - Agent 1 gets 2 cookies, agent 2 gets 0 cookies
  - They each get 1 cookie
  - Agent 1 gets 0 cookies, agent 2 gets 2 cookies
- Agent 2 chooses to accept or reject the split:
  - Accept => they each get their cookies(s)
  - Otherwise, neither gets any

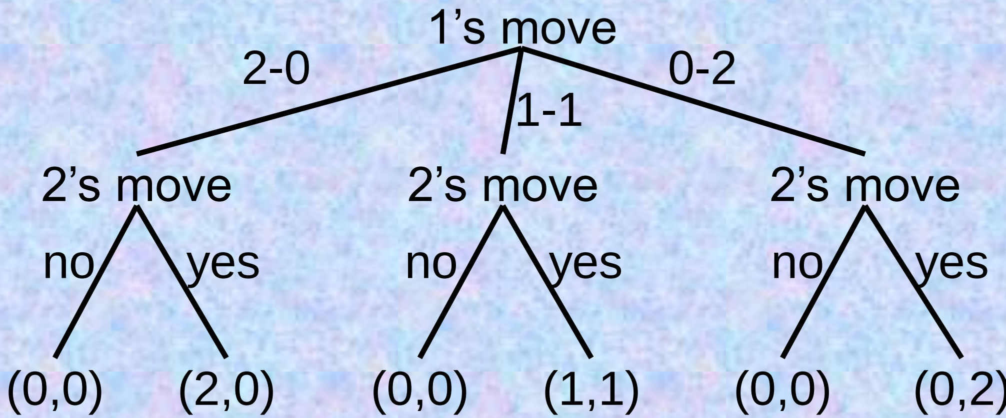

Extensive Form

- The sharing game is a game in **extensive form**
  - A game representation that makes the temporal structure explicit
  - Doesn't assume agents act simultaneously
- Extensive form can be converted to normal form
  - So previous results carry over
  - But there are additional results that depend on the temporal structure
- In a perfect-information game, the extensive form is a **game tree**:
  - **Choice** (or **nonterminal**) **node**: place where an agent chooses an action
  - **Edge**: an available **action** or **move**
  - **Terminal node**: a final outcome
  - At each terminal node *h*, each agent *i* has a utility $u_i(h)$

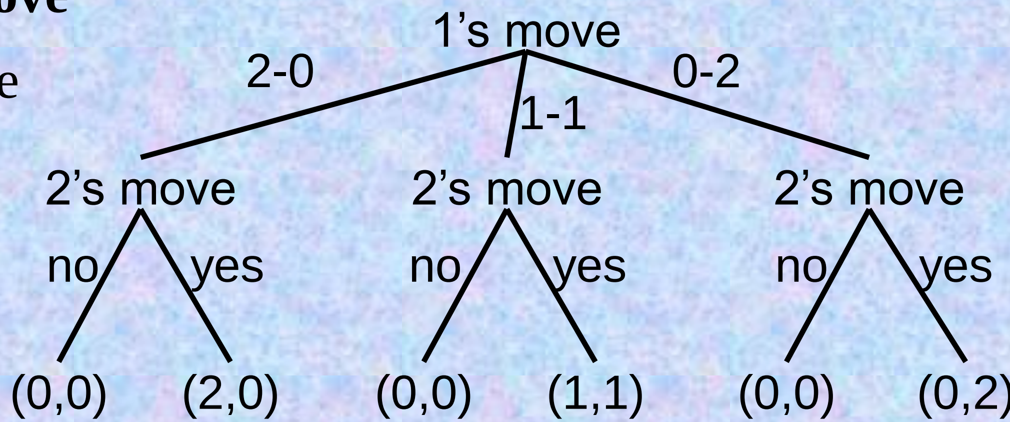

Notation from the Book (Section 4.1)

- H = {nonterminal nodes}
- Z = {terminal nodes}
- If h is a nonterminal node, then
  - $\rho(h)$ = the player to move at h
  - $\chi(h)$ = {all available actions at h}
  - $\sigma(h,a)$ = node produced by action a at node h
  - h's **children** or **successors** = $\{\sigma(h,a) : a \in \chi(h)\}$
- If h is a node (either terminal or nonterminal), then
  - h's **history** = the sequence of actions leading from the root to h
  - h's **descendants** = all nodes in the subtree rooted at h
- The book doesn't give the nodes names
  - The labels tell which agent makes the next move

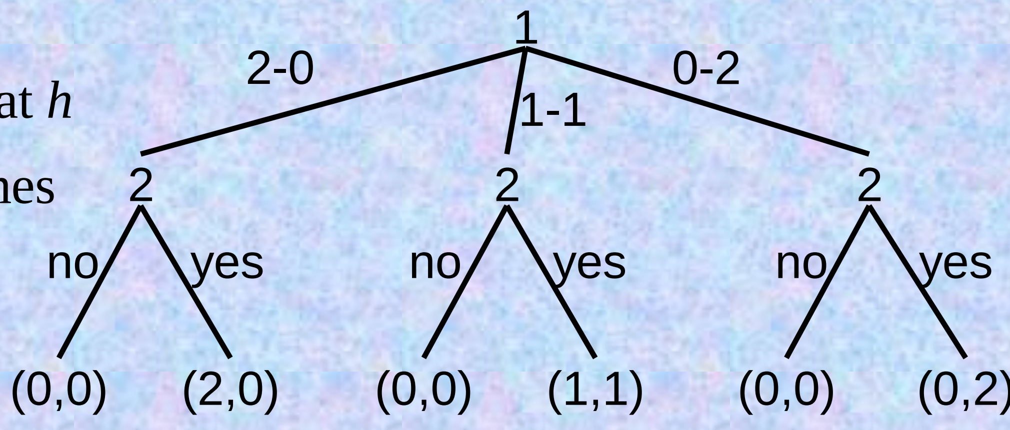

Pure Strategies

- Pure strategy for agent i in a perfect-information game:
  - Function telling what action to take at every node where it's i's choice
    - i.e., every node h at which $\rho(h) = i$
- The book specifies pure strategies as lists of actions
  - Which action at which node?
  - Either assume a canonical ordering on the nodes, or use different action names at different nodes

Sharing game:

- Agent 1 has 3 pure strategies: S_1 = {2-0, 1-1, 0-2}
- Agent 2 has 8 pure strategies:
- 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)}

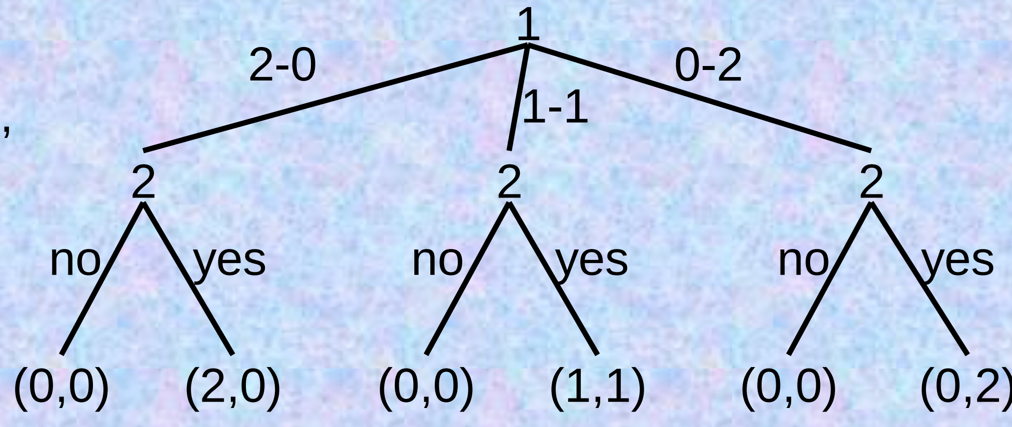

Extensive form vs. normal form

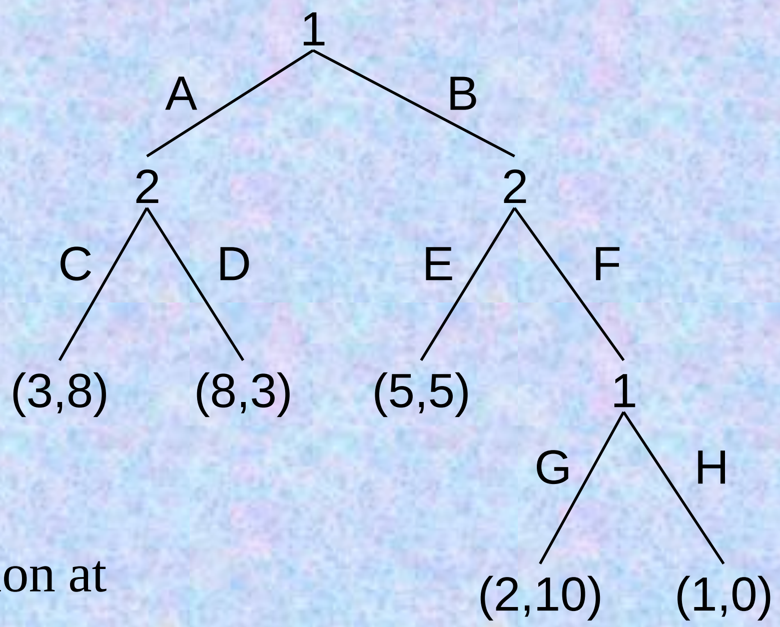

- Every game tree corresponds to an equivalent normal-form game
- The first step is to get all of the agents' pure strategies
- Each pure strategy for *i* must specify an action at every node where it's *i*'s move
- Example: the game tree shown here
  - Agent 1 has four pure strategies:
    - s_1 = {(A, G), (A, H), (B, G), (B, H)}
      - Mathematically, (A, G) and (A, H) are different strategies, even though action A makes the G-versus-H choice irrelevant
  - Agent 2 also has four pure strategies:
    - s_2 = {(C, E), (C, F), (D, E), (D, F)}

Extensive form vs. normal form

- Once we have all of the pure strategies, we can rewrite the game in normal form

- Converting to normal form introduces redundancy
  - 16 outcomes in the payoff matrix, versus 5 outcomes in the game tree
  - Payoff (3,8) occurs
    - once in the game tree
    - four times in the payoff matrix
- This can cause an exponential blowup

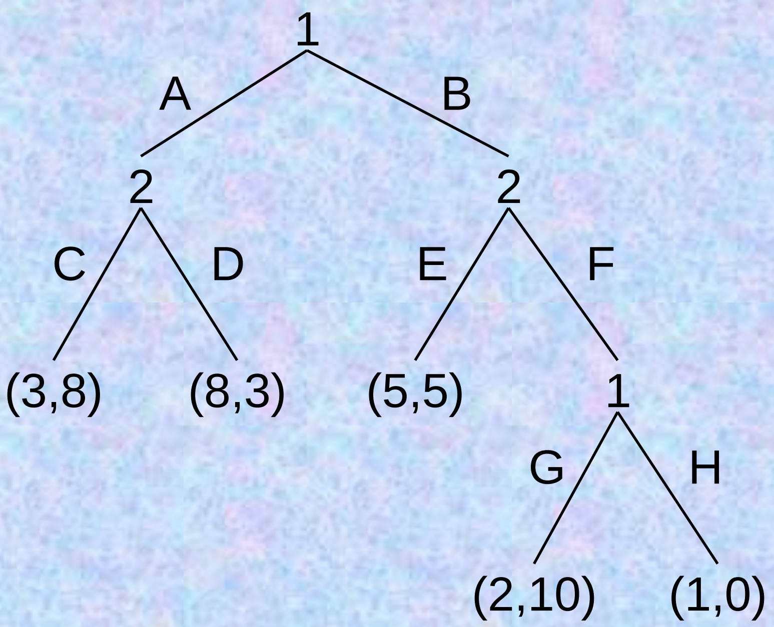

	(C,E)	(C,F)	(D,E)	(D,F)
(A,G)	3,8	3,8	8,3	8,3
(A,H)	3,8	3,8	8,3	8,3
(B,G)	5,5	2,10	5,5	2,10
(B,H)	5,5	1,0	5,5	1,0

Nash Equilibrium

- **Theorem.** Every perfect-information game in extensive form has a pure-strategy Nash equilibrium
  - This theorem has been attributed to Zermelo (1913), but there's some controversy about that
- Intuition:
  - Agents take turns, and everyone sees what's happened so far before making a move
  - So never need to introduce randomness into action selection to find an equilibrium
- In our example, there are three pure-strategy Nash equilibria

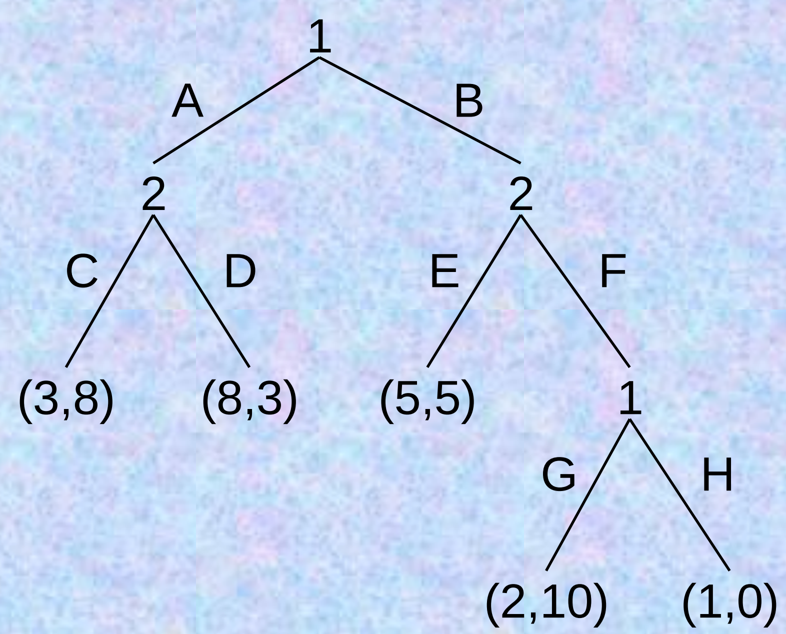

	(C,E)	(C,F)	(D,E)	(D,F)
(A,G)	3,8	3,8	8,3	8,3
(A,H)	3,8	3,8	8,3	8,3
(B,G)	5,5	2,10	5,5	2,10
(B,H)	5,5	1,0	5,5	1,0

Nash Equilibrium

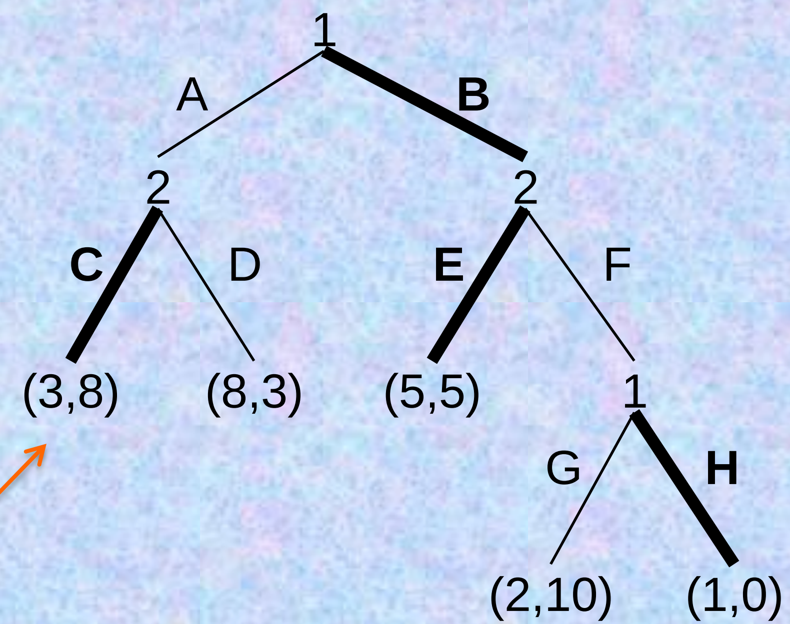

- The concept of a Nash equilibrium can be too weak for use in extensive-form games
- Recall that our example has three pure-strategy Nash equilibria:
  - {(A,G), (C,F)}
  - {(A,H), (C,F)}
  - {(B,H), (C,E)}
- Here is {(B,H), (C,E)} with the game in extensive form

	(C,E)	(C,F)	(D,E)	(D,F)
(A,G)	3,8	3,8	8,3	8,3
(A,H)	3,8	3,8	8,3	8,3
(B,G)	5,5	2,10	5,5	2,10
(B,H)	5,5	1,0	5,5	1,0

Nash Equilibrium

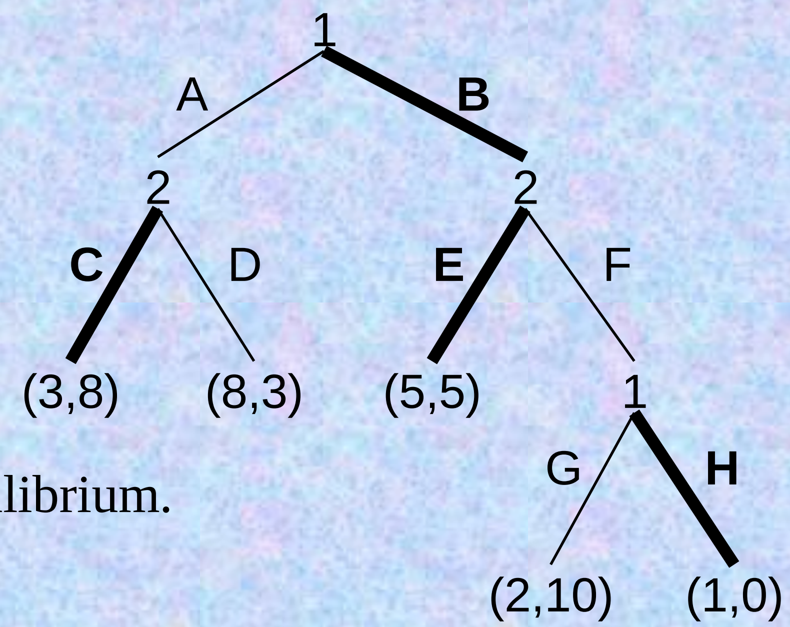

- If agent 1 used (*B*,*G*) instead of (*B*,*H*)
  - Then agent 2’s best response would be (*C*,*F*), not (*C*,*E*)
  - Thus {(B,G), (C,E)} is not a Nash Equilibrium.
- When agent 1 plays *B*
  - The only reason for agent 2 to choose *E* is if 2 knows that agent 1 has already committed to *H* rather than *G*
- This behavior by agent 1 is a *threat*:
  - By committing to choose *H*, which is harmful to agent 2, agent 1 can make agent 2 avoid that part of the tree
  - Thus agent 1 gets a payoff of 5 instead of 2
- But is agent 1’s threat credible?
  - If agent 2 plays *F*, would agent 1 *really* play *H* rather than *G*?
  - It would reduce agent 1’s own utility

Summary

- Extensive-form games
  - relation to normal-form games
  - Nash equilibria

- In extensive-form games, the game tree is often too big to search completely
  - E.g., game tree for chess: about 10^{150} nodes