CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

Polynomial time (P)

- An algorithm is said to be of **polynomial time** if its running time is upper bounded by a polynomial expression in the size of the input for the algorithm, i.e., $T(n) = O(n^k)$ for some constant k. Problems for which a deterministic polynomial time algorithm exists belong to the complexity class **P**, which is central in the field of computational complexity theory. Cobham's thesis states that polynomial time is a synonym for "tractable", "feasible", "efficient", or "fast".

- Logarithmic time: Binary Search
- Linearithmic time (O nlogn): MergeSort
- Linear time (O(n)): finding smallest or largest item in an unsorted array.
- Quadratic time: bubble sort, insertion sort
  - **Bubble sort** is a simple sorting algorithm that repeatedly steps through the list to be sorted, compares each pair of adjacent items and swaps them if they are in the wrong order. The pass through the list is repeated until no swaps are needed, which indicates that the list is sorted. The algorithm, which is a comparison sort, is named for the way smaller or larger elements "bubble" to the top of the list.
- Cubic time: Naive multiplication of two $n \times n$ matrices.

P vs. NP

- John Nash letter to NSA: https://www.nsa.gov/news-features/declassified-documents/nash-letters/assets/files/nash_letters1.pdf
  - “Nash’s letters detail an encryption technique based on the difficulty of computing certain mathematical functions – an idea that underlies modern cryptography, but was not developed publicly until the mid-1970s.”

- Informally, **NP** is the set of all decision problems for which the instances where the answer is "yes" have efficiently *verifiable* proofs. More precisely, these proofs have to be verifiable by deterministic computations that can be performed in polynomial time.

NP-complete

- Graph Isomorphism: Is graph G_1 isomorphic to graph G_2?
- Hamiltonian path: a path in an undirected or directed graph that visits each vertex exactly once
- Traveling salesman: "Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city exactly once and returns to the origin city?"
- Graph coloring: In its simplest form, it is a way of coloring the vertices of a graph such that no two adjacent vertices share the same color; this is called a **vertex coloring**

Hardness

- **NP-hardness** (*n*on-deterministic *p*olynomial-time hard), in computational complexity theory, is a class of problems that are, informally, "at least as hard as the hardest problems in NP." More precisely, a problem *H* is NP-hard when every problem *L* in NP can be reduced in polynomial time to *H*, that is: assuming a solution for *H* takes 1 unit time, we can use *H*'s solution to solve *L* in polynomial time. As a consequence, finding a polynomial algorithm to solve any NP-hard problem would give polynomial algorithms for all the problems in NP, which is unlikely as many of them are considered hard.
- NP-completeness: in NP and NP-hard.

Decision problem

- A *decision problem* is any arbitrary yes-or-no question on an infinite set of inputs. Because of this, it is traditional to define the decision problem equivalently as: the set of possible inputs together with the set of inputs for which the problem returns *yes*.

- A classic example of a decidable decision problem is the set of prime numbers. It is possible to effectively decide whether a given natural number is prime by testing every possible nontrivial factor. Although much more efficient methods of primality testing are known, the existence of any effective method is enough to establish decidability.

- Computing a Nash equilibrium in multiplayer (or two-player non zero-sum games): “a most fundamental computational problem whose complexity is wide open” and “together with factoring, the most important concrete open question on the boundary of P today”

PPAD

- In computer science, **PPAD** (*"Polynomial Parity Arguments on Directed graphs"*) is a complexity class introduced by Christos Papadimitriou in 1994. PPAD is a subclass of TFNP based on functions that can be shown to be total by a parity argument.
- “The class PPAD is defined using directed graphs based on PPA. Formally PPAD is defined by its complete problem (from an earlier post): Given an exponential-size directed graph with every node having in-degree and out-degree at most one described by a polynomial-time computable function f(v) that outputs the predecessor and successor of v, and a vertex s with a successor but no predecessors, find a t≠s that either has no successors or no predecessors.”
- Besides two player Nash, arbitrary k-player Nash and a discrete version of finding Brouwer fixed points are also complete for PPAD. “

- PPAD is a class of problems that are believed to be hard, but obtaining PPAD-completeness is a weaker evidence of intractability than that of obtaining NP-completeness. PPAD problems cannot be NP-complete, for the technical reason that NP is a class of decision problems, but the answer of PPAD problems is always yes, as a solution is known to exist, even though it might be hard to find that solution.
  - By Nash's Theorem, we know a Nash equilibrium always exists.

- Polynomial time:
  - Computing Nash equilibrium in two-player zero-sum strategic-form games
  - Computing Minmax and Maxmin values in two-player general-sum games
  - Computing Subgame Perfect Nash equilibrium in two-player zero-sum perfect-information games
  - Computing Nash equilibrium in two-player zero-sum extensive-form games
  - Various forms of domination

- NP-hard:
  - strategy elimination, reduction identity, uniqueness and reduction size problems for iterated weak dominance
  - Finding Nash equilibrium that maximizes “social welfare,” or satisfies other properties https://users.cs.duke.edu/~conitzer/nashGEB08.pdf
- PPAD-hard:
  - Computing Nash equilibrium in two-player general-sum games, or in games with >= 2 players, for both strategic-form and extensive-form games

Two-player zero-sum extensive-form games with imperfect information

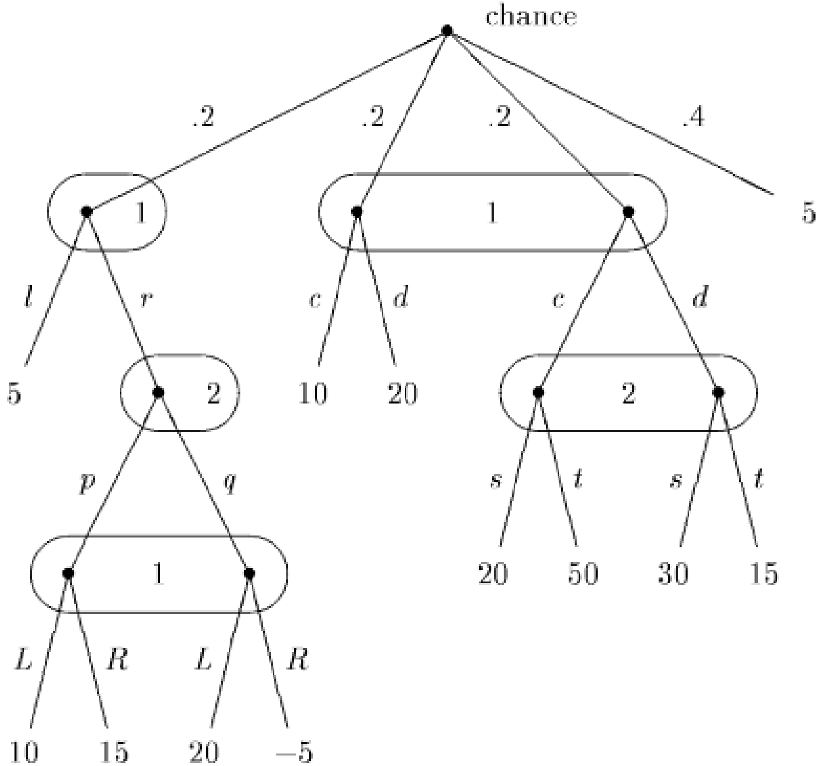

Figure 1. A zero-sum game in extensive form. From the root of the game tree, there is a chance move with the indicated probabilities. The ovals denote information sets, with the player to move written as a number inside. The leaves show the payoffs to player 1.

	(p,s)	(p,t)	(q,s)	(q,t)
(l,L,c)	9	15	9	15
(l,L,d)	13	10	13	10
(l,R,c)	9	15	9	15
(l,R,d)	13	10	13	10
(r,L,c)	8	14	10	16
(r,L,d)	10	7	12	9
(r,R,c)	9	15	5	11
(r,R,d)	11	8	7	4

Figure 2. Normal form of the game in Fig. 1. The rows and columns denote the pure strategies of player 1 and 2. The matrix entries are the expected payoffs to player 1. One of the first two pairs of rows is redundant because pure strategies with the choices l, L and l, R have the same effect. They are identified in the reduced normal form.

$\boldsymbol{E} =$

	$\emptyset$	l	r	rL	rR	c	d
	1						
	−1	1	1				
			−1	1	1		
	−1					1	1

$\boldsymbol{e} =$

1
0
0
0

$\boldsymbol{F} =$

	$\emptyset$	p	q	s	t
	1				
	−1	1	1		
	−1			1	1

$\boldsymbol{f} =$

1
0
0

$\boldsymbol{A} =$

	$\emptyset$	p	q	s	t
$\emptyset$	2				
l	1				
r					
rL		2	4		
rR		3	−1		
c	2			4	10
d	4			6	3

Figure 3. Constraint matrices and payoff matrix for the realization weights of the sequences in the game in Fig. 1. The matrices are sparse, zero entries are omitted.

“Sequence form” LP formulation

$$\begin{aligned} \underset{\boldsymbol{y},\boldsymbol{p}}{\text{minimize}} \quad & \boldsymbol{e}^T\boldsymbol{p} \\ \text{subject to} \quad & -\boldsymbol{A}\boldsymbol{y} + \boldsymbol{E}^T\boldsymbol{p} \geq \boldsymbol{0} \\ & -\boldsymbol{F}\boldsymbol{y} = -\boldsymbol{f}, \\ & \boldsymbol{y} \geq \boldsymbol{0}. \end{aligned} \tag{8}$$

$$\begin{aligned} \underset{\boldsymbol{x},\boldsymbol{q}}{\text{maximize}} \quad & -\boldsymbol{q}^T\boldsymbol{f} \\ \text{subject to} \quad & \boldsymbol{x}^T(-\boldsymbol{A}) - \boldsymbol{q}^T\boldsymbol{F} \leq \boldsymbol{0} \\ & \boldsymbol{x}^T\boldsymbol{E}^T = \boldsymbol{e}^T, \\ & \boldsymbol{x} \geq \boldsymbol{0}. \end{aligned} \tag{9}$$

Support enumeration algorithm

forall *support size profiles* $x = (x_1, x_2)$*, sorted in increasing order of, first,* $|x_1 - x_2|$ *and, second,* $(x_1 + x_2)$ **do**

- **forall** $\sigma_1 \subseteq A_1$ *s.t.* $|\sigma_1| = x_1$ **do**
  - $A'_2 \leftarrow \{a_2 \in A_2$ not conditionally dominated, given $\sigma_1 \}$
  - **if** $\nexists a_1 \in \sigma_1$ *conditionally dominated, given* A'_2 **then**
    - **forall** $\sigma_2 \subseteq A'_2$ *s.t.* $|\sigma_2| = x_2$ **do**
      - **if** $\nexists a_1 \in \sigma_1$ *conditionally dominated, given* σ_2 **and***TGS is satisfiable for* $\sigma = (\sigma_1, \sigma_2)$ **then**
        - **return** *the solution found; it is a NE*

Figure 4.6: The SEM algorithm

n-player general-sum games

- For n-player games with n >= 3, the problem of computing an NE can no longer be expressed as an LCP. While it can be expressed as a *nonlinear complementarity problem*, such problems are often hopelessly impractical to solve exactly.
- Can solve sequence of LCPs (generalization of Newton's method).
  - Not globally convergent
- Formulate as constrained optimization (minimization of a function), but also not globally convergent (e.g., hill climbing, simulated annealing can get stuck in local optimum)
- Simplicial subdivision algorithm (Scarf)
  - Divide space into small regions and search separately over the regions.
- Homotopy method (Govindan and Wilson)
  - n-player extension of Lemke-Howson Algorithm

Critiques of Nash equilibrium

- Is it too strict?
  - Does not exist in all games
  - Might rule out some more “reasonable” strategies
- Not strict enough?
  - Potentially many equilibria to select through
- Just right?

Definition 5.1.1 (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*.*

Backward induction

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 $best\_util$

Figure 5.6: Procedure for finding the value of a sample (subgame-perfect) Nash equilibrium of a perfect-information extensive-form game.

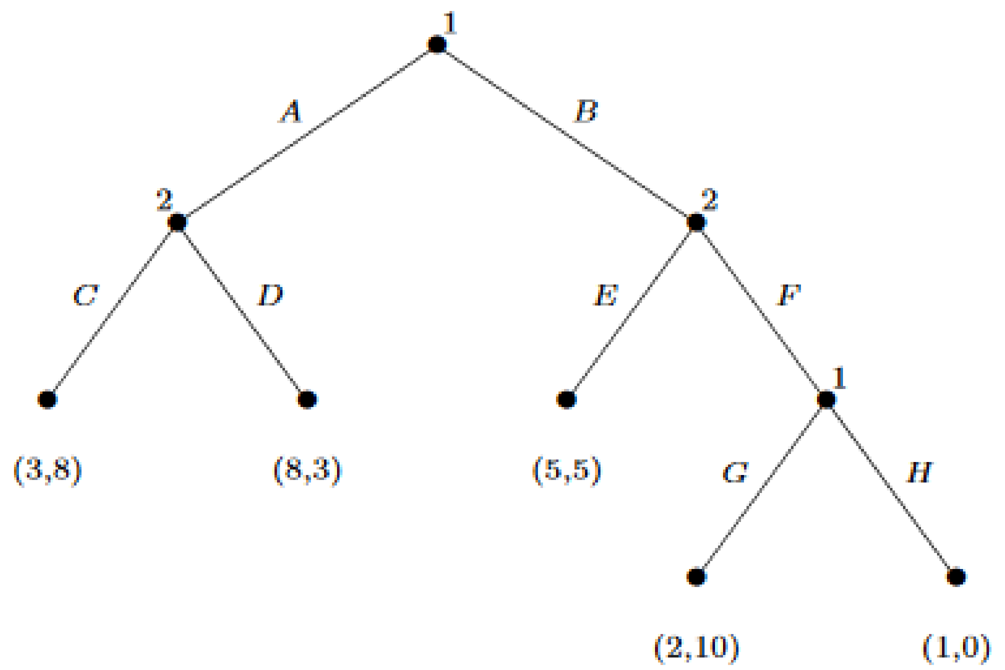

Figure 5.2: A perfect-information game in extensive form.

	(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

Figure 5.4: Equilibria of the game from Figure 5.2.

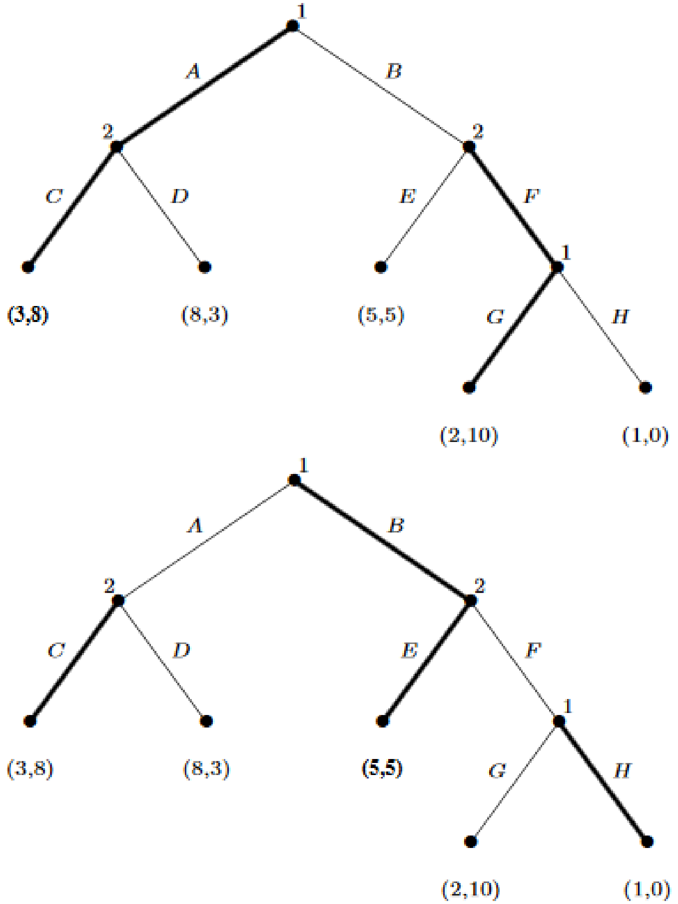

Figure 5.5: Two out of the three equilibria of the game from Figure 5.2: $\{(A,G),(C,F)\}$ and $\{(B,H),(C,E)\}$. Bold edges indicate players' choices at each node.

- First consider {(A,G),(C,F)}.
- {(B,H), (C,E)} less intuitive.
- Why is {(B,G), (C,E)} not an equilibrium??
- Player 1's decision to play H is a **threat**. But is it **credible**?

- Definition: 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 subgames of G rooted at some node in G.

- The **subgame-perfect equilibrium (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'.

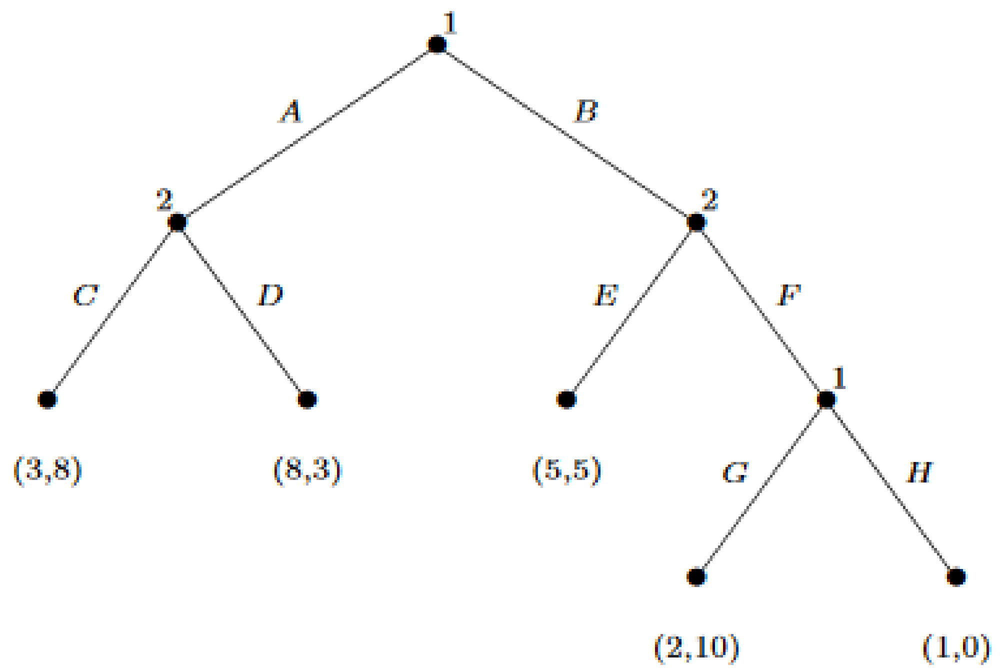

Figure 5.2: A perfect-information game in extensive form.

	(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

Figure 5.4: Equilibria of the game from Figure 5.2.

Subgame perfect equilibrium

- Every SPE is also a Nash equilibrium
- Furthermore, although SPE is a stronger concept than Nash equilibrium (i.e., every SPE is a Nash equilibrium, but not every NE is a SPE) it is still the case that every perfect-information extensive-form game has at least one subgame-perfect equilibrium.
- Rules out “noncredible threats.” The only SPE is {(A,G, (C,F)}. Consider the subgame rooted at player 1’s second choice node …

Another SPE Example

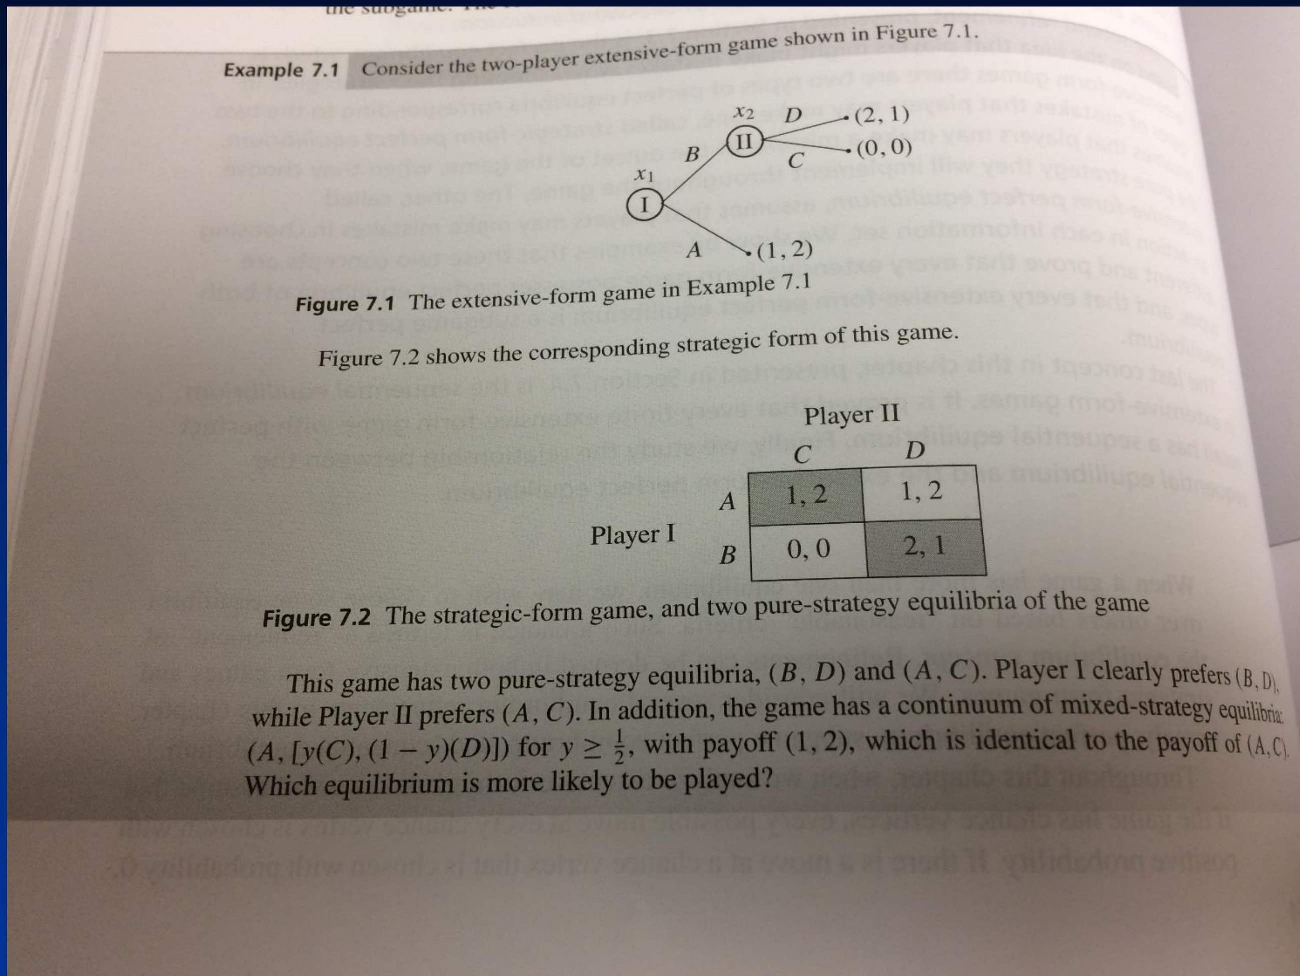

Example 7.1 Consider the two-player extensive-form game shown in Figure 7.1.

Figure 7.1 The extensive-form game in Example 7.1

Figure 7.2 shows the corresponding strategic form of this game.

Player II

Player I	C	D
A	1, 2	1, 2
B	0, 0	2, 1

Figure 7.2 The strategic-form game, and two pure-strategy equilibria of the game

This game has two pure-strategy equilibria, (B, D) and (A, C). Player I clearly prefers (B, D), while Player II prefers (A, C). In addition, the game has a continuum of mixed-strategy equilibria: $(A, [y(C), (1-y)(D)])$ for $y \geq \frac{1}{2}$, with payoff $(1, 2)$, which is identical to the payoff of (A, C). Which equilibrium is more likely to be played?

- What is the SPE?
- Is playing C a “credible threat”?

- Theorem: Every finite extensive-form game with perfect information has a subgame perfect equilibrium in pure strategies.

- Theorem: Every extensive-form game with "perfect recall" has a subgame perfect equilibrium in mixed strategies.

Repeated Prisoner's dilemma?

Centipede game

- Two players take turns choosing either to take a slightly larger share of an increasing pot, or to pass the pot to the other player. The payoffs are arranged so that if one passes the pot to one's opponent and the opponent takes the pot on the next round, one receives slightly less than if one had taken the pot on this round. Although the traditional centipede game had a limit of 100 rounds (hence the name), any game with this structure but a different number of rounds is called a centipede game.

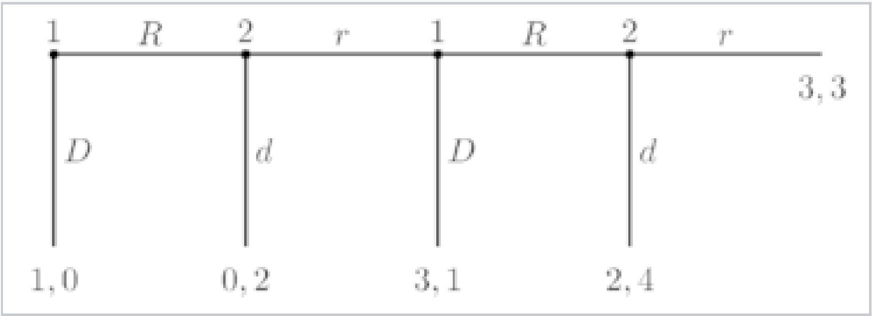

An Extensive Form representation of a four-stage centipede game, which ends after four rounds with the money being split. Passing the coins across the table is represented by a move of **R** (going across the row of the lattice, sometimes also represented by **A** for across) and pocketing the coins is a move **D** (down the lattice). The numbers **1** and **2** along the top of the diagram show the alternating decision-maker between two players denoted here as 1 and 2, and the numbers at the bottom of each branch show the payoff for players 1 and 2 respectively.

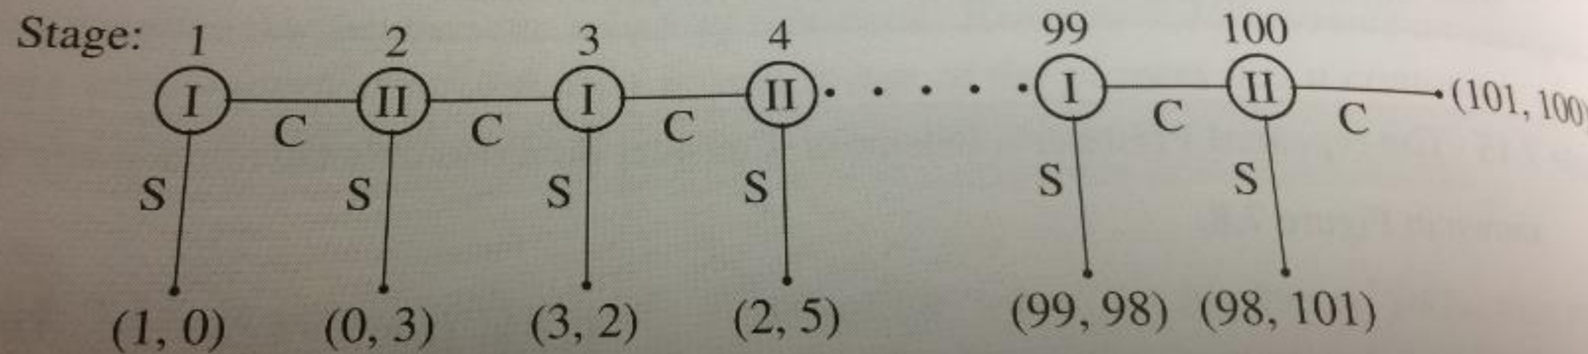

Figure 7.9 The Centipede game

- Consider two players: Alice and Bob. Alice moves first. At the start of the game, Alice has two piles of coins in front of her: one pile contains 4 coins and the other pile contains 1 coin. Each player has two moves available: either "take" the larger pile of coins and give the smaller pile to the other player or "push" both piles across the table to the other player. Each time the piles of coins pass across the table, the quantity of coins in each pile doubles. For example, assume that Alice chooses to "push" the piles on her first move, handing the piles of 1 and 4 coins over to Bob, doubling them to 2 and 8. Bob could now use his first move to either "take" the pile of 8 coins and give 2 coins to Alice, or he can "push" the two piles back across the table again to Alice, again increasing the size of the piles to 4 and 16 coins. The game continues for a fixed number of rounds or until a player decides to end the game by pocketing a pile of coins.

- Standard game theoretic tools predict that the first player will defect on the first round, taking the pile of coins for himself. There are several pure strategy Nash equilibria of the centipede game and infinitely many mixed strategy Nash equilibria. However, there is only one subgame perfect equilibrium (a popular refinement to the Nash equilibrium concept).
- In the unique subgame perfect equilibrium, each player chooses to defect at every opportunity. This, of course, means defection at the first stage. In the Nash equilibria, however, the actions that would be taken after the initial choice opportunities (even though they are never reached since the first player defects immediately) may be cooperative.

- Several studies have demonstrated that the Nash equilibrium (and likewise, subgame perfect equilibrium) play is rarely observed. Instead, subjects regularly show partial cooperation, playing "R" (or "r") for several moves before eventually choosing "D" (or "d"). It is also rare for subjects to cooperate through the whole game. For examples see McKelvey and Palfrey (1992) and Nagel and Tang (1998). As in many other game theoretic experiments, scholars have investigated the effect of increasing the stakes. As with other games, for instance the ultimatum game, as the stakes increase the play approaches (but does not reach) Nash equilibrium play.

Explanations?

- Altruism
- Error
- Degree of financial incentives?
- Palacios-Huerta and Volij (2009) find that expert chess players play differently from college students. With a rising Elo, the probability of continuing the game declines; all Grandmasters in the experiment stopped at their first chance. They conclude that chess players are familiar with using backward induction reasoning and hence need less learning to reach the equilibrium. However, in an attempt to replicate these findings, Levitt, List, and Sadoff (2010) find strongly contradictory results, with zero of sixteen Grandmasters stopping the game at the first node.

Trembling-hand perfect equilibrium

	L	**R**
U	1, 1	2, 0
D	0, 2	2, 2

- Two pure strategy equilibria (U,L) and (D,R).
- Assume row player is playing $(1-\varepsilon, \varepsilon)$ for $0 < \varepsilon < 1$ …

- In game theory, **trembling hand perfect equilibrium** is a refinement of Nash equilibrium due to Reinhard Selten. A trembling hand perfect equilibrium is an equilibrium that takes the possibility of off-the-equilibrium play into account by assuming that the players, through a “slip of the hand” or **tremble,** may choose unintended strategies, albeit with negligible probability.

- First we define a **perturbed game**. A perturbed game is a copy of a base game, with the restriction that only totally mixed strategies are allowed to be played. A totally mixed strategy is a mixed strategy where *every* pure strategy is played with non-zero probability. This is the "trembling hands" of the players; they sometimes play a different strategy than the one they intended to play. Then we define a strategy set S (in a base game) as being trembling hand perfect if there is a sequence of perturbed games that converge to the base game in which there is a series of Nash equilibria that converge to S.

Extensive-form games

- Two ways of defining trembling hand perfect equilibrium:
  - every strategy of the extensive-form game must be played with non-zero probability. This leads to the notion of a **normal-form trembling hand perfect equilibrium**.
  - every move at every information set is taken with non-zero probability. Limits of equilibria of such perturbed games as the tremble probabilities goes to zero are called **extensive-form trembling hand perfect equilibria**.
- These two notions are incomparable.

- Theorem: Every finite strategic-form game has at least one perfect equilibrium.
- Theorem: In every perfect equilibrium, every (weakly) dominated strategy is chosen with probability zero.
- Theorem: Every equilibrium in completely mixed strategies in a strategic-form game is a perfect equilibrium.

- Theorem: Every extensive-form game has a strategic-form perfect equilibrium.
- Theorem: Every extensive-form perfect equilibrium of extensive-form game Γ is a subgame perfect equilibrium.
- Every finite extensive-form game with perfect recall has an extensive-form perfect equilibrium.
- Every finite extensive-form game with perfect recall has a subgame perfect equilibrium in behavior strategies.

Critiques of Nash equilibrium

- Is it too strict?
  - Does not exist in all games
  - Might rule out some more "reasonable" strategies
- Not strict enough?
  - Potentially many equilibria to select through
- Just right?

Next lecture

- Wrap up refinements:
  - Evolutionarily stable strategies, sequential equilibrium, proper equilibrium.
- Clarify new concepts: perfect recall, behavior strategies
- Continue exploring repeated games.
  - What happens if you are playing prisoner's dilemma (or Golden Balls) repeatedly against the same opponent?

Assignment

- HW2 due 2/21.
- HW3 out 2/21 (due 3/2).
- Midterm on 3/7 (midterm review on 3/2).
  - Will cover material from lectures and homeworks (will not cover material from the textbooks that was not covered in lectures or homeworks).
  - 3 parts: multiple choice, true/false with explanation, analytical exercises
- Reading for next class: chapter 13 from main textbook