

CSCE 631 — Algorithmic Game Theory

Lecture 10: Extensive-Form Games I

Alan Kuhnle

Summer 2026

Texas A&M University

Today's Plan

From Normal Form to Extensive Form

Game Trees

Perfect vs. Imperfect Information

Information Sets

Chance Nodes

Formal Definition of an Extensive-Form Game

Examples

Strategies in Extensive-Form Games

From Extensive Form to Normal Form

Summary

From Normal Form to Extensive Form

Limitations of Normal-Form Representation

- Normal-form (strategic-form) games specify:
 - Players $N = \{1, \dots, n\}$
 - Strategy sets $S_1, \dots, S_n$
 - Payoff functions $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$
- **Missing:** sequential structure, timing, information revelation
- Many strategic situations have a natural *temporal ordering*:
 - Negotiations, auctions, parlor games, military engagements
 - “I observe your move, then I respond”
- Extensive-form games capture:
 - The order of moves
 - What each player knows when they move
 - The role of chance (nature)

Why Extensive Form?


- **Richer modeling:** captures sequential interaction, imperfect information
- **More refined equilibria:** subgame perfection, sequential equilibrium
- **Computational structure:** tree-based algorithms (backward induction, CFR)
- **Practical applications:** poker, negotiations, mechanism design, security games

Key Insight

Every extensive-form game has a normal-form representation, but the extensive form is exponentially more compact and carries information about sequential rationality.

Game Trees

Game Trees: Informal Picture



- Circle nodes: decision points; square nodes: terminal (payoff) nodes
- Edges: actions available at each decision node

Components of a Game Tree

1. **Nodes** V : a finite set forming a rooted tree
 - **Root** v_0 : the starting point of the game
 - **Decision nodes** D : where a player chooses an action
 - **Terminal nodes** $Z \subset V$: leaves where the game ends
2. **Edges**: each edge from a decision node corresponds to an *action*
3. **Player assignment** $P : D \rightarrow N \cup \{c\}$
 - $P(v) = i$ means player i moves at node v
 - $P(v) = c$ means chance (nature) moves
4. **Action sets** $A(v)$ for each decision node v
5. **Payoff function** $u_i : Z \rightarrow \mathbb{R}$ for each player i

Terminal Nodes and Payoffs

- Each terminal node $z \in Z$ specifies a **payoff vector**

$$u(z) = (u_1(z), u_2(z), \dots, u_n(z)) \in \mathbb{R}^n$$

- The payoff encodes the outcome of a complete play of the game
- A **play** (or history) is a path from the root v_0 to some terminal node z
- The sequence of actions along this path is the *action history*

Convention

In two-player zero-sum games: $u_1(z) + u_2(z) = 0$ for all $z \in Z$.

Perfect vs. Imperfect Information

Definition

An extensive-form game has **perfect information** if every information set is a singleton: each player knows exactly which decision node they are at when it is their turn to move.

- Examples: chess, checkers, Go, tic-tac-toe
- The entire history of play is common knowledge at every decision point
- Every perfect-information game has a pure-strategy Nash equilibrium (Zermelo's theorem)

Imperfect Information Games

Definition

An extensive-form game has **imperfect information** if at least one information set contains more than one node: some player cannot distinguish between certain game histories when choosing their action.

- Examples: poker (hidden cards), Battleship, sealed-bid auctions
- A player may not know:
 - Earlier moves by opponents
 - Outcomes of chance events (private cards)
 - Even some of their own earlier moves (in exotic settings)
- Imperfect information is what makes games *strategically interesting* and *computationally hard*

Information Sets

Information Sets: Formal Definition

Definition


An **information set** I for player i is a partition element of player i 's decision nodes such that:

1. All nodes in I belong to the same player: $P(v) = i$ for all $v \in I$
2. The same actions are available at every node in I : $A(v) = A(v')$ for all $v, v' \in I$

We write $\mathcal{I}_i$ for the collection of all information sets of player i .

- Player i knows they are at *some* node in I but cannot distinguish which one
- The available actions at I are $A(I) := A(v)$ for any $v \in I$
- In perfect-information games: $|I| = 1$ for every information set

Information Sets: Visual Representation



- Dashed box: Player 2 cannot distinguish the two nodes
- Player 2 does not know whether Player 1 chose L or R
- Player 2 must choose the same (possibly mixed) action at both nodes

Chance Nodes

Chance Nodes (Nature)

Definition

A **chance node** is a decision node assigned to a special player called *nature* (or *chance*). At each chance node v , nature selects an action $a \in A(v)$ according to a fixed, commonly known probability distribution $\sigma_c(v)$.

- Model exogenous randomness: card deals, dice rolls, weather
- Nature has no payoff function and no strategic interests
- The probability distribution at each chance node is part of the game description

Example (Poker): Nature deals cards at the start; each deal is a chance event. The resulting card assignment is private information for each player.

Formal Definition of an Extensive-Form Game

Formal Definition

Definition (Extensive-Form Game)

An **extensive-form game** Γ is a tuple

$$\Gamma = (N, V, v_0, Z, A, P, \sigma_c, \mathcal{I}, u)$$

where:

- $N = \{1, \dots, n\}$: set of players
- (V, v_0) : game tree rooted at v_0 ; $Z \subseteq V$ are terminal nodes
- A : action function; $A(v)$ is the set of actions at decision node v
- $P : V \setminus Z \rightarrow N \cup \{c\}$: player function
- σ_c : probability distributions at chance nodes
- $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_n)$: information partition
- $u = (u_1, \dots, u_n)$: payoff functions $u_i : Z \rightarrow \mathbb{R}$


Definition

Player i has **perfect recall** if, for every information set $I \in \mathcal{I}_i$, every node $v \in I$ shares the same sequence of information sets visited by player i and actions taken by player i on the path from v_0 to v .

- Informally: a player never “forgets” their own past actions or past information
- Almost all games of practical interest satisfy perfect recall
- **Absent-minded driver** is a classic example of *imperfect recall*
- Perfect recall is required for Kuhn’s theorem (coming soon)

Examples

Example: Sequential Prisoner's Dilemma




Perfect information version:

- Player 1 moves first (C or D)
- Player 2 *observes* the choice, then moves
- Payoffs match standard PD matrix

Analysis:

- Player 2 plays *D* at both nodes (dominant)
- Player 1 anticipates this, plays *D*
- Unique SPE: (*D*, *DD*) with payoff (1, 1)
- Sequential play does not help

Example: Entry Deterrence Game



- **Entrant** chooses: enter market or stay out
- **Incumbent** responds (if entry): fight or accommodate
- Two NE:
 1. (In, Accommodate): payoff (1, 1)
 2. (Out, Fight): payoff (0, 2)
- NE #2 involves a **non-credible threat**: incumbent threatens to fight, but if entry actually occurs, fighting is irrational
- Only NE #1 is **subgame perfect**

Strategies in Extensive-Form Games

Pure Strategies

Definition

A **pure strategy** for player i is a function

$$s_i : \mathcal{I}_i \rightarrow \bigcup_{I \in \mathcal{I}_i} A(I)$$

such that $s_i(I) \in A(I)$ for every information set $I \in \mathcal{I}_i$.

- A pure strategy specifies an action at *every* information set, even those not reached during play
- The number of pure strategies is

$$\prod_{I \in \mathcal{I}_i} |A(I)|$$

- This grows **exponentially** in the depth of the tree

Mixed vs. Behavioral Strategies

Mixed Strategy

A **mixed strategy** σ_i is a probability distribution over player i 's pure strategies. If player i has $|S_i|$ pure strategies, $\sigma_i \in \Delta(S_i)$.

Behavioral Strategy

A **behavioral strategy** β_i assigns, at each information set $I \in \mathcal{I}_i$, a probability distribution over the actions $A(I)$:

$$\beta_i(I) \in \Delta(A(I))$$

The randomization is *independent* across information sets.

- Behavioral strategies are far more compact: $\sum_{I \in \mathcal{I}_i} |A(I)|$ parameters vs. $|S_i|$ for mixed

Mixed vs. Behavioral: Example

Setting: Player has two information sets I_1, I_2 , each with actions $\{L, R\}$.

Mixed strategy:

σ : probability p_1 on (L, L) , p_2 on (L, R) , p_3 on (R, L) , p_4 on (R, R)

Can encode *correlations* between choices at I_1 and I_2 .

Behavioral strategy:

$\beta(I_1)$: play L with prob q_1

$\beta(I_2)$: play L with prob q_2

Independence: $\Pr[\text{play } (L, L)] = q_1 q_2$

Cannot correlate choices across information sets.

Question: When does this independence restriction matter?

Kuhn's Theorem

Theorem (Kuhn, 1953)

In an extensive-form game with **perfect recall**, every mixed strategy is *realization-equivalent* to some behavioral strategy, and vice versa. That is, for every mixed strategy σ_i , there exists a behavioral strategy β_i that induces the same probability distribution over terminal nodes (for any fixed strategies of the other players), and conversely.

- **Realization-equivalent:** same distribution over outcomes for every opponent strategy profile
- Under perfect recall, the correlation structure that mixed strategies allow is never useful
- Consequence: it suffices to work with behavioral strategies
- This is crucial for computational tractability (Lecture 12–14)

Kuhn's Theorem: Proof Sketch

Behavioral $\Rightarrow$ Mixed: trivially, any behavioral strategy is a special case of a mixed strategy (with independent randomization).

Mixed $\Rightarrow$ Behavioral: Given mixed strategy σ_i , define behavioral strategy β_i by:


$$\beta_i(a \mid I) = \frac{\sum_{s_i: s_i(I)=a, s_i \text{ reaches } I} \sigma_i(s_i)}{\sum_{s_i: s_i \text{ reaches } I} \sigma_i(s_i)}$$

i.e., the conditional probability of playing a at I given that I is reached.

- Perfect recall ensures this is well-defined: reaching I depends only on player i 's earlier information sets and actions (not on the specific node within I)
- The key: under perfect recall, the probability of reaching any node factors into independent per-information-set contributions

Kuhn's Theorem: Why Perfect Recall Matters

Counterexample (absent-minded driver):



- Player forgets whether they are at v_1 or v_2 (imperfect recall)
- Behavioral strategy: exit with prob p . Expected payoff: $4(1 - p)^2$. Best: $p = 0$,

From Extensive Form to Normal Form

Converting Extensive Form to Normal Form

- Every extensive-form game induces a normal-form game:
 - Players: same N
 - Strategy sets: pure strategies S_i (mappings from info sets to actions)
 - Payoffs: expected payoff under the strategy profile (averaging over chance)
- **Problem:** exponential blowup
 - If player i has k information sets each with m actions: $|S_i| = m^k$
 - The normal-form representation has $\prod_i |S_i|$ entries

Takeaway

The normal form is **exponentially larger** than the extensive form. This motivates algorithms that work directly on the game tree (backward induction, sequence form, CFR).

Example: Normal-Form Representation

Entry deterrence game (from earlier):

- Entrant strategies: {In, Out}
- Incumbent strategies: {Fight, Accommodate}

	Fight	Accommodate
In	-1, -1	1, 1
Out	0, 2	0, 2

- Two NE: (In, Acc.) and (Out, Fight)
- Normal form hides the sequential structure
- Both look equally “valid” in normal form, but (Out, Fight) relies on a non-credible threat

Summary

Summary

- **Extensive-form games** model sequential interaction with a game tree
- Key components: players, actions, terminal nodes, payoffs, player assignment, information sets, chance nodes
- **Perfect information:** every info set is a singleton (chess, Go)
- **Imperfect information:** some info sets are non-singleton (poker)
- **Pure strategies** map each information set to an action — exponentially many
- **Behavioral strategies** assign local distributions at each info set — polynomially many parameters
- **Kuhn's theorem:** under perfect recall, behavioral strategies suffice (realization equivalence)
- Normal-form representation is exponentially larger

Next time: Subgame perfect equilibrium, backward induction, and sequential