CSCE 631 — Algorithmic Game Theory

Lecture 11: Extensive-Form Games II

Alan Kuhnle

Summer 2026

Texas A&M University

Today's Plan

Subgame Perfect Equilibrium

Motivation: Non-Credible Threats

Recall the **Entry Deterrence** game:

- Two Nash equilibria: (In, Accommodate) and (Out, Fight)
- (Out, Fight) relies on a *non-credible threat*: incumbent threatens to fight, but would never follow through if entry actually occurred
- The normal-form NE concept cannot distinguish credible from non-credible threats

Goal

Define an equilibrium refinement that eliminates Nash equilibria supported by non-credible threats in sequential games.

Subgames: Formal Definition

Definition

A **subgame** of an extensive-form game Γ rooted at node v is the restriction of Γ to the subtree at v, provided:

1. v is a *singleton* information set (i.e., $\{v\}$ is an information set for player $P(v)$)
2. The subtree rooted at v *does not cut any information set*: for every information set I, either I is entirely contained in the subtree or entirely outside it

- The entire game Γ (rooted at v_0) is always a subgame
- In a **perfect-information** game, every decision node starts a subgame
- In games with imperfect information, subgames can be rare

Subgame Perfect Equilibrium (SPE)

Definition (Selten, 1965)
A strategy profile σ^* is a **subgame perfect equilibrium (SPE)** if for every subgame G' of Γ, the restriction of σ^* to G' is a Nash equilibrium of G'.

- Every SPE is a NE, but not every NE is an SPE
- SPE is a **refinement** of NE
- SPE eliminates strategies that are not sequentially rational (i.e., not optimal given that we have reached a particular subgame)
- Every finite perfect-information game has a pure-strategy SPE

SPE vs. NE: The Refinement

	Nash Equilibrium	**Subgame Perfect Eq.**
Requirement	No player can profitably deviate *globally*	No player can profitably deviate in *any subgame*
Credibility	May rely on non-credible threats	All threats must be credible
Existence	Always exists (in mixed strategies)	Always exists in finite games (Kuhn)
Computation	Can be PPAD-hard	Backward induction for perfect info

Backward Induction

Backward Induction Algorithm

For **perfect-information** games, SPE can be computed by backward induction:

1. Start at the **terminal nodes**; their payoffs are known.
2. At each **penultimate decision node** v: the moving player $P(v)$ selects the action maximizing $u_{P(v)}$ among the children.
3. Replace v with the selected child's payoff vector and repeat.
4. Continue until the root is reached.

Theorem (Zermelo / Kuhn)

Backward induction computes a subgame perfect equilibrium in *pure strategies* for every finite perfect-information game. The resulting SPE is unique if all payoffs at terminal nodes are distinct for the moving player.

Backward Induction: Example

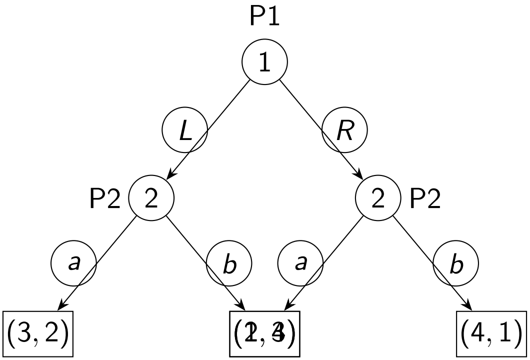

- **Step 1:** P2 at left node: $\max(2, 4) \Rightarrow b$, payoff $(1, 4)$
- **Step 2:** P2 at right node: $\max(3, 1) \Rightarrow a$, payoff $(2, 3)$
- **Step 3:** P1 at root: $\max(1, 2) \Rightarrow R$, payoff $(2, 3)$
- **SPE:** P1 plays R; P2 plays b at left, a at right

Backward Induction: Complexity

- **Time:** $O(|V|)$ — linear in the size of the game tree
- **Space:** $O(\text{depth})$ using depth-first traversal
- This makes perfect-information games "easy" computationally
- **Caveat:** the game tree itself may be exponentially large (e.g., chess: $\sim 10^{47}$ nodes)

Perfect-Information Games are "Solved"

Every finite, two-player, zero-sum perfect-information game has a deterministic minimax value computable by backward induction (alpha-beta pruning in practice).

SPE Computation: More Examples

Example: Centipede Game

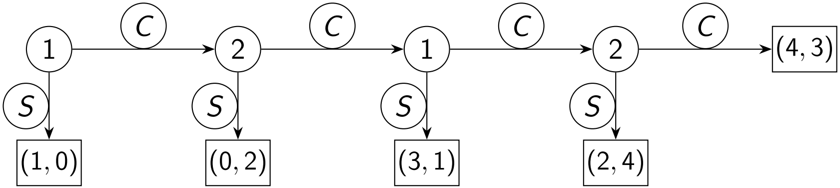

- Backward induction: P2 at last node stops $(4 > 3)$. P1 at node 3 stops $(3 > 2)$. P2 at node 2 stops $(2 > 1)$. P1 at node 1 stops $(1 > 0)$.
- **SPE:** Both players stop immediately. Payoff $(1, 0)$.
- Paradox: mutual cooperation yields $(4, 3)$, but backward induction unravels cooperation entirely.

Example: Stackelberg Competition

- **Setup:** Two firms; Firm 1 (leader) chooses quantity q_1 first, Firm 2 (follower) observes and chooses q_2
- Inverse demand: $P(Q) = a - Q$ where $Q = q_1 + q_2$
- Profit: $\pi_i = q_i(a - q_1 - q_2) - c\, q_i$

Backward induction:

1. Firm 2's best response: $q_2^*(q_1) = \frac{a-c-q_1}{2}$
2. Firm 1 maximizes $\pi_1 = q_1\big(a - c - q_1 - q_2^*(q_1)\big)$
3. SPE: $q_1^* = \frac{a-c}{2}$, $q_2^* = \frac{a-c}{4}$
4. Leader advantage: $\pi_1^* > \pi_2^*$ (first-mover advantage)

Compare with simultaneous Cournot: $q_1^C = q_2^C = \frac{a-c}{3}$.

Example: Ultimatum Game

- **Setup:** Player 1 proposes split $(x, 1-x)$ of \$1; Player 2 accepts or rejects
- If accept: payoffs $(x, 1-x)$. If reject: payoffs $(0, 0)$.

SPE analysis:

- Player 2 accepts any $x < 1$ (getting $1 - x > 0$ is better than 0)
- Player 1 offers $x = 1 - \epsilon$ (keep almost everything)
- Unique SPE outcome: Player 1 gets ≈ 1, Player 2 gets ≈ 0

Behavioral observation:

- Experiments show people reject "unfair" offers (e.g., 80/20 splits)
- SPE prediction fails descriptively — motivates behavioral game theory
- But SPE remains the correct *rational* benchmark

Imperfect Information: Sequential Equilibrium

Limitations of SPE with Imperfect Information

- SPE requires subgames, but imperfect-information games may have *very few* subgames
- If the only subgame is the entire game, SPE = NE (no refinement power)
- We need a concept that enforces sequential rationality even at information sets not reached in equilibrium

Key Idea

Augment the strategy profile with a **belief system**: at each information set, specify what the player believes about which node they are at. Then require both beliefs and strategies to be "consistent."

Beliefs at Information Sets

Definition

A **belief system** μ assigns to each information set $I \in \mathcal{I}_i$ a probability distribution $\mu(\cdot \mid I)$ over the nodes in I:

$$\mu(v \mid I) \geq 0 \quad \text{for all } v \in I, \qquad \sum_{v \in I} \mu(v \mid I) = 1$$

$\mu(v \mid I)$ represents player i's belief that the game is at node v, given that information set I has been reached.

- Beliefs encode a player's uncertainty about the game state
- At singleton information sets: $\mu(v \mid \{v\}) = 1$ (trivially)
- The interesting case: non-singleton information sets

Definition

An **assessment** is a pair (σ, μ) where:

- $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a behavioral strategy profile
- μ is a belief system

Two requirements for a "good" assessment:

1. **Sequential rationality:** at every information set I, player i's strategy σ_i maximizes expected payoff given beliefs $\mu(\cdot \mid I)$ and opponents' strategies σ_{-i}
2. **Consistency:** beliefs μ are derived from strategies σ via Bayes' rule wherever possible (i.e., at information sets reached with positive probability)

Sequential Rationality

Definition

An assessment (σ, μ) is **sequentially rational** if for every player i and every information set $I \in \mathcal{I}_i$:

$$\sum_{v \in I} \mu(v \mid I) \cdot U_i(\sigma \mid v) \geq \sum_{v \in I} \mu(v \mid I) \cdot U_i((\sigma'_i, \sigma_{-i}) \mid v)$$

for all alternative strategies σ'_i for player i, where $U_i(\sigma \mid v)$ is the expected payoff starting from node v.

- Player i maximizes expected payoff at *every* information set, not just at the start of the game
- The belief μ resolves uncertainty about which node in I is the actual game state

Consistency of Beliefs

Bayes Consistency (Weak)

At information sets reached with positive probability under σ, beliefs must follow Bayes' rule:

$$\mu(v \mid I) = \frac{\Pr^{\sigma}(v)}{\sum_{v' \in I} \Pr^{\sigma}(v')} \qquad \text{if } \overset{\sigma}{\Pr}(I) > 0$$

Problem: What about information sets reached with zero probability? Bayes' rule is silent here.

Full Consistency (Kreps–Wilson)

(σ, μ) is **consistent** if there exists a sequence of *completely mixed* strategy profiles $\sigma^k \to \sigma$ such that the Bayesian beliefs μ^k computed from σ^k converge to μ.

Sequential Equilibrium

Definition (Kreps and Wilson, 1982)
An assessment (σ^*, μ^*) is a **sequential equilibrium** if:

1. (σ^*, μ^*) is **sequentially rational**
2. (σ^*, μ^*) is **consistent**

- Refines NE: every sequential equilibrium induces a NE
- Refines SPE in games with perfect information
- Handles off-path information sets via the consistency requirement
- **Existence:** every finite extensive-form game has a sequential equilibrium

Perfect Bayesian Equilibrium

Perfect Bayesian Equilibrium (PBE)

Definition

An assessment (σ^*, μ^*) is a **perfect Bayesian equilibrium** if:

1. (σ^*, μ^*) is sequentially rational
2. At every information set reached with positive probability: μ^* is derived from σ^* by Bayes' rule
3. At information sets reached with zero probability: μ^* can be any distribution (no restriction beyond being a valid distribution)

- PBE is *weaker* than sequential equilibrium (less restrictive on off-path beliefs)
- PBE is *stronger* than NE (requires sequential rationality)
- In many applications, PBE and sequential equilibrium coincide
- PBE is the most commonly used solution concept in applied game theory

Relationship Among Solution Concepts

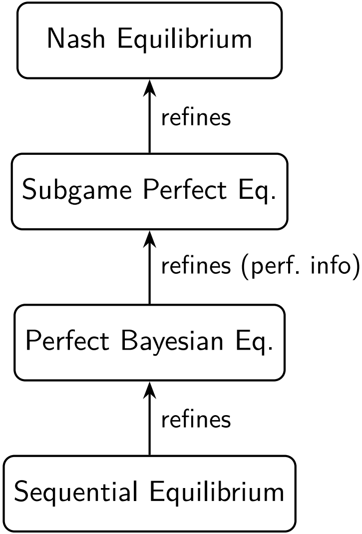

Signaling Games

Signaling Games: Setup

- Two players: **Sender** (S) and **Receiver** (R)
- Nature chooses Sender's type $t \in T$ according to prior $p(t)$
- Sender observes t, sends message $m \in M$
- Receiver observes m (but not t), chooses action $a \in A$
- Payoffs: $u_S(t, m, a)$ and $u_R(t, m, a)$

Information structure:

- Sender's information sets: one per type t (sender knows their type)
- Receiver's information sets: one per message m (receiver sees message but not type)
- Receiver's info set for message m contains $|T|$ nodes

Signaling Game Example: Job Market

Spence's Job Market Signaling (simplified):

- Worker (Sender) is type H (high ability) or L (low), equal prior
- Worker chooses education level: $e \in \{0, 1\}$
- Firm (Receiver) observes e, offers wage $w \in \{w_L, w_H\}$
- Payoffs: Worker gets $w - c(e, t)$ where $c(1, H) < c(1, L)$ (education is cheaper for high-ability workers)

PBE types:

- **Separating:** H gets education, L does not. Firm infers type perfectly from signal.
- **Pooling:** Both types choose the same e. Firm learns nothing from the signal.
- Equilibrium selection depends on cost parameters and wages

Computing PBE in Signaling Games

Step-by-step procedure:

1. **Conjecture** a strategy type (separating, pooling, semi-separating)
2. **Derive beliefs** at each of Receiver's information sets using Bayes' rule (where applicable)
3. **Compute Receiver's best response** given beliefs
4. **Verify Sender's optimality:** given Receiver's response, does each Sender type want to play the conjectured strategy?
5. If any player wants to deviate, the conjectured profile is *not* a PBE

Off-Path Beliefs

For messages sent with zero probability, beliefs are unconstrained. Different off-path beliefs can support different equilibria. Refinements like the "intuitive criterion" (Cho–Kreps) restrict these.

Summary

Summary

- **Subgame perfect equilibrium** eliminates non-credible threats; computed by backward induction in perfect-information games
- **Backward induction:** $O(|V|)$ algorithm for SPE in perfect-information games; always yields a pure-strategy equilibrium
- SPE has limited power in imperfect-information games (few subgames)
- **Beliefs** at information sets formalize what players think when they cannot observe the full history
- **Sequential equilibrium** (Kreps–Wilson): assessment (σ, μ) that is sequentially rational and consistent
- **Perfect Bayesian equilibrium:** practical workhorse; Bayes consistency on-path, unrestricted off-path
- **Signaling games:** canonical application of PBE (separating vs. pooling equilibria)

Next time: Sequence form representation and polynomial time algorithms for