Subgame Perfect Equilibrium

Game Theory · Module 3 — Extensive Form Games

Node 3.5

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Extensive Form Games **Node:** 3.5

The entry deterrence game (Node 1.6) revealed that the normal form of an extensive-form game can contain Nash equilibria supported by **incredible threats** — off-path behavior that is not individually rational but that equilibrium conditions cannot punish because it is never reached. Subgame perfect equilibrium (Selten, 1965) is the refinement that eliminates such equilibria. It requires that the strategy profile induce a Nash equilibrium not only in the whole game but in every **subgame** — every part of the tree rooted at a singleton information set. It is the extensive-form refinement of Nash equilibrium, and it is what backward induction computes in perfect-information games. For imperfect-information games, the refinement extends to perfect Bayesian equilibrium (Node 4.4) and sequential equilibrium (Node 4.6).

Formal Definition

Let Γ be an extensive-form game. A **subgame** of Γ is a subtree Γ' rooted at some decision node t such that:

1. The root t is a **singleton information set** (the acting player knows they are at t and nowhere else).
2. For every information set I that meets the subtree rooted at t, either $I \subseteq \Gamma'$ or $I \cap \Gamma' = \emptyset$. (Information sets cannot be "split" by the subgame.)
3. The subgame inherits players, actions, chance probabilities, and terminal payoffs from Γ.

A **subgame perfect equilibrium (SPE)** of Γ is a strategy profile σ^* such that the restriction of σ^* to every subgame is a Nash equilibrium of that subgame.

Theorem (Selten). In every finite perfect-information game, backward induction produces a subgame perfect equilibrium. Every perfect-information game has at least one SPE; in generic games it is unique.

Nash $\supseteq$ SPE. Every SPE is a Nash equilibrium (take the subgame to be Γ itself). Not every Nash equilibrium is SPE; the non-SPE Nash equilibria are precisely those relying on incredible off-path behavior.

Intuition

Nash equilibrium requires "nobody has a profitable deviation at the root of the game." SPE requires "nobody has a profitable deviation at the root of *any subgame*." The extra condition is that, even at nodes that would not be reached in equilibrium, the strategy prescribes optimal play conditional on being there.

The motivation is **credibility**. If a player commits to a threat ("I will fight if you enter") that would be suboptimal to execute if the threat were tested, the threat is not credible. A rational opponent should not fear it, and equilibrium reasoning that relies on fear of the threat is not robust. SPE formalizes this by requiring optimality everywhere in the tree, not just on the equilibrium path.

A second way to see it: SPE is "Nash equilibrium in every subtree." The normal form loses subtree-level information (Node 1.6), and SPE is the refinement that recovers it.

The canonical SPE computation algorithm is backward induction: at each subtree, the optimal action is chosen given optimal continuations. For perfect-information games this is exact; for imperfect-information games, subgames only exist at singleton information sets, which may be rare, and the refinement must be further strengthened (Node 4.4, 4.6).

Structural Role

SPE is the extensive-form strengthening of Nash equilibrium:

- **2.1 Nash ⇒ 3.5 SPE.** SPE is a strict refinement: fewer equilibria, eliminating incredible threats.
- **3.4 Backward induction ⇒ 3.5.** Backward induction produces SPE in perfect-info games.
- **3.5 ⇒ 3.6 Centipede.** SPE is the solution concept that predicts "Take immediately" in Centipede, setting up the empirical puzzle.
- **3.5 ⇒ 5.x Repeated games.** In finitely repeated games, SPE unraveling forces the unique stage-Nash profile; in infinitely repeated games, the Folk Theorem characterizes the SPE payoff set.
- **3.5 ⇒ 4.4 PBE.** For imperfect-information games, SPE is too weak; PBE extends it.
- **3.5 ⇒ 10.4 Commitment.** Credibility of commitment is directly about whether a threat survives SPE refinement.

Mathematical Development

Subgame identification

A subgame is rooted at a singleton information set. In perfect-information games, every decision node is a singleton, so every subtree is a subgame. In games with non-singleton information sets, most subtrees are *not* subgames — their root is inside an information set with other nodes, so the acting player doesn't know they are at the root of the subtree.

Example: AKQ poker. After the chance deals cards, P1 has 3 singleton info sets (one per card). Each P1-level decision is the root of a subgame (the subtree starting from P1's decision). After P1's action, P2's info set (one per P2's card, conditioned on P1 betting) is also a singleton. So there are subgames rooted at each P2 decision. But the *whole game* — starting from chance — is the only global subgame; the chance node is the root, and it has multiple "positions" (for each (P1 card, P2 card) pair) but chance is distinct from players so the subgame definition still works at the chance-root level.

Imperfect-information games have fewer subgames than decision nodes, and SPE is consequently weaker.

One-shot deviation principle

Theorem. In a finite extensive-form game, a strategy profile is SPE iff no player can profitably deviate at any single information set holding all other play (their own and opponents') fixed.

This reduces the SPE check from "consider all possible deviations" to "check one-shot deviations at each information set." It is the discrete-time analogue of the Bellman principle: optimality is preserved under single-step deviations. Proof is by induction on the tree.

Selten's example

Selten (1965) introduced the chain store paradox to motivate SPE. A chain store operates in many small towns; in each, an entrant may enter or stay out. The chain store prefers to "fight" (drive out entrants) to build a reputation, even at a one-town loss, so that future entrants fear entry. But at the last town, the store has no future reputation to protect, so fighting is irrational. Backward induction unravels: at town k, fighting is irrational; knowing this, town $k-1$ entrant enters; knowing this, fighting at $k-1$ is irrational too; iterating, the store never fights, and every entrant enters. The "reputation strategy" fails SPE at every subgame.

The paradox is that this backward induction prediction contradicts intuition and experiment. It is resolved only by introducing incomplete information about the store's type (Node 5.5).

SPE in stochastic games and bargaining

Rubinstein's alternating-offer bargaining game (1982) has a unique SPE computable by backward induction on the continuation values: with discount factor δ, the proposer offers $\frac{1-\delta}{1-\delta^2} = \frac{1}{1+\delta}$ to themselves. This uniqueness is the famous Rubinstein result and shows how SPE can pin down a single outcome in games with continuous action spaces.

Worked Example

Entry deterrence with SPE check

Entry deterrence (Node 1.6): Entrant Out $\to (0,2)$; Entrant In, Incumbent Accommodate $\to (1,1)$; Entrant In, Incumbent Fight $\to (-1,-1)$.

Nash equilibria of normal form. Two pure: (In, Accommodate) and (Out, Fight).

SPE check. The whole game is one subgame. Is (In, Accommodate) SPE? The only non-root decision is the Incumbent's — and in this subgame (rooted at Incumbent's node), Accommodate gives 1 vs Fight giving -1. Accommodate is the optimal action. ✓

Is (Out, Fight) SPE? The subgame rooted at the Incumbent's node has Fight as the strategy. But Fight is suboptimal in that subgame: Accommodate gives 1 vs Fight giving -1. Not a Nash of the subgame. Not SPE. ×

Conclusion. Unique SPE: (In, Accommodate). The equilibrium (Out, Fight) is eliminated.

Rubinstein bargaining (brief)

Two players alternate making offers to split 1 unit. Each period without agreement discounts by $\delta \in (0,1)$. P1 offers, P2 accepts/rejects; if rejects, P2 offers, and so on.

Backward induction via stationary SPE. Let x^* be P1's equilibrium offer to themselves. In the SPE, P1's offer gives P2 exactly the discounted continuation value $\delta(1-x^*)$. P2 accepts iff $1-x^* \geq \delta(1-x^*)$, which holds. P2's counteroffer (if reached) would give themselves x^* (by symmetry), so P2's continuation value is δx^*. P1's indifference condition: $x^* = 1 - \delta \cdot (\delta x^*) \Rightarrow x^* = \frac{1}{1+\delta}$.

Unique SPE: P1 offers $(x^*, 1-x^*) = (1/(1+\delta), \delta/(1+\delta))$ and P2 accepts immediately.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Subgame	A subtree of a hand — e.g., "turn onward given this board and betting history"
SPE	Every subgame's play is Nash of that subgame — i.e., the strategy is credible at every decision point

Formal object	Poker instantiation
Non-SPE Nash	A Nash equilibrium supported by incredible off-path threats — rare in solver output because solvers enforce SPE-like optimality
One-shot deviation principle	Check your play at each info set: would a single deviation, holding everything else fixed, improve EV? If no for every info set, it's SPE.
Credible commitment	"I will always 3-bet aggressively" — not credible without a mechanism enforcing it

Concrete situation

Subgame decomposition for study. When you load a flop spot into a solver and "solve" it, you are solving a subgame: the subtree rooted at the flop decision, given fixed preflop ranges (the input). The solver output is an SPE of this subgame — it is Nash within the subgame, and because the subtree is perfectly-information-set at its root (the acting player knows their own hand and the board), the subgame root is a valid singleton info set in the usual sense.

Why this decomposition is legitimate. SPE guarantees that if you know the correct input ranges at the root of a subgame, the optimal play *within* the subgame can be computed independently of the larger game. This is exactly what "solving a flop in isolation" relies on. Libratus and DeepStack used this to do "re-solving" during play: when a new state was reached, they solved the subgame rooted there, using the opponent's range as previously computed.

Incredible threats in poker. A pure Nash equilibrium where a player commits to "I will always 3-bet shove with 72o preflop, never bluff with suited connectors" can exist in the normal form if the opponent's belief is consistent with it. But it fails SPE: at the 3-betting information set, the optimal action is not to shove 72o — it is to fold. Any realistic solver refuses such an equilibrium because it is not SPE.

What this understanding buys you

"Solving a subgame" is rigorous because of SPE. When you freeze the preflop ranges and solve the flop in isolation, you are relying on the SPE property: optimal play in a subtree depends only on the subtree plus the input ranges at its root. This is why subgame solvers work — and also why getting the input ranges wrong is catastrophic (the "subgame" you're solving isn't actually a subgame of the right game).

Every solver output is SPE, by construction. CFR iterates toward an SPE by ensuring that no one-shot deviation at any information set is profitable. The output is not just Nash; it is subgame perfect. Any claim "this is the GTO play at this info set" is implicitly a claim that the full strategy is SPE.

Incredible commitments do not work in poker. If you try to adopt a strategy like "I always shove with this" that would be suboptimal at the info set in question, the solver (and a good human opponent) will exploit you. There is no hidden commitment mechanism; your strategy is always executing at the info set you are at, and it must be optimal there.

Cross-Links

- **1.6 Normal form as universal representation** — the flattening that loses SPE information.
- **2.1 Nash equilibrium** — SPE is a refinement.
- **3.1–3.4** — the extensive-form machinery.
- **3.6 Centipede game** — puzzle of SPE empirical predictions.
- **4.4 Perfect Bayesian equilibrium** — extension to imperfect information.
- **5.x Repeated games** — finitely repeated unraveling via SPE.
- **10.4 Commitment and credibility** — SPE formalizes credibility.

Structural Compression

Invariant: A subgame perfect equilibrium is a strategy profile whose restriction to every subgame is a Nash equilibrium of that subgame — equivalently, no one-shot deviation at any information set is profitable.

Mechanism: Recursive optimality on the tree: backward induction or its generalization to imperfect-information games, enforced by the one-shot deviation principle.

Cross-System Echo: Subgame perfection is the game-theoretic analogue of the principle of optimality in dynamic programming (Bellman) and the markov property in stochastic processes. In all three, local optimality at every state/time implies global optimality, and vice versa.

Exercises

1. Verify that the Stackelberg duopoly equilibrium is SPE.
2. Construct a 2-player extensive-form game with two Nash equilibria, exactly one of which is SPE. Identify the incredible threat.
3. Apply the one-shot deviation principle to verify that the backward induction equilibrium of a 3-stage perfect-info game is SPE.
4. In Rubinstein bargaining, derive the unique SPE by writing the two indifference conditions at the proposer's info set and the responder's.
5. Show that in a perfect-information game, every SPE is a backward induction equilibrium.
6. Construct an imperfect-information extensive-form game where SPE fails to pin down a unique equilibrium. What refinement would be needed?
7. **Argue, structurally,** why subgame perfection is the "right" refinement of Nash in extensive-form games with perfect information. What specific strategic pathology does it rule out, and why is that pathology a problem?

Game Theory · Module 3 — Extensive Form Games · Node 3.5
Concepts: nash-equilibrium, fixed-point
Prerequisites: 2.1, 3.1, 3.3, 3.4
Enables: 3.6, 4.4, 5.3

hbar.university — own.your.cognition