Strategies in Extensive Form

Game Theory · Module 3 — Extensive Form Games

Node 3.3

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Extensive Form Games **Node:** 3.3

In the normal form, a strategy is a single choice from S_i. In the extensive form, it is a *function* mapping information sets to actions — a complete contingent plan of action that specifies what the player would do at every decision point. This node makes the strategy object formal and introduces the three representations — **pure**, **mixed**, and **behavioral** — that will be used throughout dynamic game theory. Kuhn's theorem ties them together under perfect recall. This is the node that lets us formally write down "what a poker solver stores," which is a behavioral strategy.

Formal Definition

Let Γ be an extensive-form game with information partitions $\mathcal{I}_i$ and action set $A(I)$ at each information set I.

Pure strategy. A function

$$s_i : \mathcal{I}_i \to \bigcup_{I \in \mathcal{I}_i} A(I) \quad \text{with } s_i(I) \in A(I).$$

The set of pure strategies is the Cartesian product $S_i = \prod_{I \in \mathcal{I}_i} A(I)$. Its cardinality is $\prod_{I \in \mathcal{I}_i} |A(I)|$, which is generally exponential in the number of information sets.

Mixed strategy. A probability distribution over pure strategies:

$$\sigma_i \in \Delta(S_i).$$

A mixed strategy can be implemented by randomizing once, at the start of the game, over a complete contingent plan.

Behavioral strategy. A collection of independent action distributions, one per information set:

$$\beta_i = (\beta_i(I))_{I \in \mathcal{I}_i}, \quad \beta_i(I) \in \Delta(A(I)).$$

A behavioral strategy can be implemented by randomizing independently at each information set. The space of behavioral strategies is $\prod_{I \in \mathcal{I}_i} \Delta(A(I))$, which has dimension only linear in the number of information sets.

Kuhn's theorem (perfect recall). Under perfect recall, every mixed strategy is equivalent (in outcome distribution) to some behavioral strategy, and vice versa. The two representations yield the same expected payoff against any opponent strategy.

Intuition

A pure strategy in extensive form is *exhaustive*: it tells you what you would do in every possible situation you could be in, including situations you will never actually reach in any specific play. A pure strategy is a complete contingency plan.

A mixed strategy randomizes *once* at the start of the game over complete contingency plans. You commit to a single plan but the plan itself is random.

A behavioral strategy randomizes *at each information set*, independently. When you reach an information set, you flip a coin (possibly biased) and choose an action based on the outcome.

These feel different but, under perfect recall, they are equivalent. Kuhn's theorem says: for any mixed strategy, you can construct a behavioral strategy with the same outcome, and vice versa. The proof uses perfect recall essentially: since you remember your own past actions, the behavioral decomposition can be done information-set-by-information-set without losing strategic content.

The practical consequence: solvers store behavioral strategies (linear space) rather than mixed strategies (exponential space). Everything we call "the GTO frequency at this info set" is the behavioral strategy's conditional distribution at that info set. Without Kuhn's theorem, we would be stuck storing mixed strategies, and the whole computation would be infeasible.

Structural Role

Extensive-form strategies underpin every downstream concept:

- **3.1–3.2 ⇒ 3.3.** The tree and information sets set the domain of strategies.
- **3.3 ⇒ 3.4.** Backward induction computes optimal actions at each information set.
- **3.3 ⇒ 3.5.** Subgame perfection requires optimality at every information set, not just on the equilibrium path.
- **3.3 ⇒ 4.4.** Perfect Bayesian equilibrium adds beliefs at information sets.
- **3.3 ⇒ 9.3.** CFR iterates over information sets, updating the behavioral strategy based on counterfactual regret at each.

Kuhn's theorem is also the reason extensive-form solutions are computationally feasible at all: the behavioral-strategy space is $O(|\mathcal{I}_i| \cdot |A|)$, whereas the mixed-strategy space is $O(|A|^{|\mathcal{I}_i|})$ — an exponential difference.

Mathematical Development

Equivalence of mixed and behavioral under perfect recall

Let σ_i be a mixed strategy. Define the **behavioral representation** β_i^σ as follows: at information set I, let $\Pr[\text{reach } I \mid \sigma_i]$ be the probability of reaching I when player i plays σ_i and the other players move to allow I to be reached. Set

$$\beta_i^\sigma(I)(a) = \frac{\Pr[\text{reach } I \text{ and play } a \text{ at } I \mid \sigma_i]}{\Pr[\text{reach } I \mid \sigma_i]} \quad \text{if the denominator is positive.}$$

Under perfect recall, this conditional distribution is well-defined (independent of opponent play through I), and the induced outcome distribution of β_i^σ against any opponent profile matches that of σ_i. This is the content of Kuhn's theorem.

Pure, mixed, behavioral in a small tree

Consider a tree where player 1 has 2 information sets, each with 2 actions.

- **Pure strategies:** $2 \times 2 = 4$. $((A, A), (A, B), (B, A), (B, B))$
- **Mixed strategies:** $\Delta(\{(A, A), (A, B), (B, A), (B, B)\})$, a 3-dimensional simplex.
- **Behavioral strategies:** $\Delta(\{A, B\}) \times \Delta(\{A, B\})$, a $1 \times 1 = 2$-dimensional rectangle (parametrized by two probabilities).

Under perfect recall, any mixed strategy σ_1 can be represented by a behavioral strategy β_1 with the same outcome distribution, even though σ_1 has 3 parameters and β_1 has only 2. The "missing" parameter is the correlation between information sets — which mixed can capture but behavioral cannot — but perfect recall makes such correlation strategically irrelevant.

Imperfect recall counterexample

Consider a game where player 1 moves at two information sets in sequence, but at the second information set player 1 has forgotten the first action. This is absent memory. A mixed strategy can correlate the two actions (e.g., always play AA or always play BB, 50/50 each). A behavioral strategy cannot: it mixes independently at each information set, so it would produce AA, AB, BA, BB each with probability 1/4.

The outcome distributions differ. Kuhn's theorem fails. This is why perfect recall is essential.

Reach probabilities

For a behavioral strategy β_i, the probability of reaching information set I is

$$\pi_i(I \mid \beta_i) = \prod_{(I',a) \preceq I} \beta_i(I')(a),$$

where the product is over i's own information sets on the path to I and the actions taken at them. The total reach probability of I, factoring in opponents and chance, is $\pi_i(I \mid \beta_i) \cdot \pi_{-i}(I)$, where $\pi_{-i}(I)$ is independent of β_i under perfect recall. This factorization is what makes CFR work.

Worked Example

AKQ game strategies

From the prior node's AKQ game. P1 has 3 information sets (one per P1's card), each with 2 actions (B, C). P2 has 3 information sets (one per P2's card, reached after P1 bets), each with 2 actions (C, F).

P1 pure strategies: $2^3 = 8$. Example: $s_1 = (B, C, B)$ means "bet A, check K, bet Q."

P1 behavioral strategy: 3 numbers in $[0, 1]$, one per information set, giving the bet probability.

P1 mixed strategy: distribution over 8 pure strategies — 7-dimensional simplex.

Equivalence check. A mixed strategy $\sigma_1 = 0.5 \cdot (B, C, B) + 0.5 \cdot (B, B, C)$ randomizes over two plans, each chosen with probability 0.5. Its behavioral representation: at the A info set, both plans play B, so $\beta_1(A)(B) = 1$. At the K info set, one plan plays C, the other plays B, each with prob 0.5, so $\beta_1(K)(B) = 0.5, \beta_1(K)(C) = 0.5$. At the Q info set, one plan plays B, the other plays C, each with prob 0.5, so $\beta_1(Q)(B) = 0.5, \beta_1(Q)(C) = 0.5$.

Behavioral strategy: $\beta_1 = (B \text{ prob } 1, B \text{ prob } 0.5, B \text{ prob } 0.5)$. Outcome distribution: same as σ_1 (under perfect recall).

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Pure strategy s_i	A complete deterministic plan: "I will do X with every hand at every spot in every betting history"
Mixed strategy σ_i	A randomization over complete plans — a distribution over "whole strategies"
Behavioral strategy β_i	The object the solver outputs: a distribution over actions at *each* info set, independent per info set
Reach probability of an info set	How often, given your whole strategy, you end up at this exact decision point
Perfect recall assumption	"I remember my own cards and my own past actions" — always true in real poker

Kuhn's theorem is why you can memorize solver output. Solver tables look like "at this info set, bet 62% / check 38%" — that is a behavioral strategy (local probabilities per info set). Without Kuhn's theorem, you would need to memorize a joint distribution over the astronomical pure-strategy space. With Kuhn's theorem, you only need one distribution per info set. This reduces the memorization problem from "impossible" to "merely huge."

Concrete situation

Solver output interpretation. You are drilling a flop spot with $AhKh$ on $Kd7s2c$ in a BTN vs BB single-raised pot. PioSolver tells you: - bet 33% pot: 73% - check: 27%

This is a **behavioral strategy** at your information set. The interpretation is "independently randomize at this spot with these frequencies." To implement, you need a real random device (second hand of a watch, suit of a lower hole card in a different range, an app).

Notice the number 73% is *not* a probability over complete strategies. It is the conditional probability of betting 33% *at this specific info set.* The solver also outputs independent mixes at every other info set. Your complete strategy is the product of all these conditional mixes, and — by Kuhn's theorem — this product is strategically equivalent to some mixed strategy over complete plans.

What this understanding buys you

Solver outputs are per-info-set, not joint. You do not need to correlate your flop check frequency with your turn mixing. Each info set is independent. This is the crucial simplification that makes memorization possible.

Range is a consequence of behavioral strategy + reach probabilities. Your c-bet range on a given board is the set of your hole cards for which the behavioral strategy at the corresponding info set puts positive probability on betting. This is a derived quantity, not an independent object. When a coach says "your betting range should be X," they are saying "your behavioral strategy at this info set should put weight on betting with these hole cards."

Imperfect recall is a real risk in abstractions. Solvers that bucket hands together (treating AKs and AQs identically at some info sets, e.g.) can introduce imperfect recall. When this happens, Kuhn's theorem fails, and the behavioral output is no longer equivalent to a mixed strategy, which means the "solution" may not be a Nash equilibrium of the original game. Libratus and DeepStack handled this with careful re-solving at abstraction boundaries.

Cross-Links

- **1.6 Normal form as universal representation** — extensive strategies flatten to normal-form pure strategies.
- **3.1 Game trees** — the domain.
- **3.2 Information sets** — the argument domain of strategies.
- **3.4 Backward induction** — computes optimal pure strategies.
- **3.5 Subgame perfect equilibrium** — requires optimality at every info set.
- **4.4 Perfect Bayesian equilibrium** — adds belief systems.
- **9.3 CFR** — iterates on behavioral strategies directly.
- **POMDPs** — the computer-science analogue.

Structural Compression

Invariant: A strategy in an extensive-form game is a function from information sets to actions (pure), a distribution over such functions (mixed), or a collection of independent conditional distributions at each information set (behavioral). Under perfect recall, all three representations yield the same outcome distributions.

Mechanism: Factorization of strategic randomization across information sets; Kuhn's theorem establishes equivalence between local (behavioral) and global (mixed) randomization under memory-preservation constraints.

Cross-System Echo: The equivalence of mixed and behavioral strategies under perfect recall mirrors the equivalence of Markov policies and history-dependent policies in reinforcement learning under appropriate informational conditions. In both cases, a "local" rule (behavioral, Markov) suffices when the agent has adequate memory.

Exercises

1. Count the pure strategies for each player in the two-stage entry game (Node 3.1).
2. For the AKQ game, write out a pure strategy for P1 that dominates a specific alternative pure strategy. Verify the dominance.
3. Prove Kuhn's theorem in the special case of a game with two information sets.
4. Construct a game with imperfect recall and exhibit a mixed strategy whose outcome distribution is not achievable by any behavioral strategy.
5. For a given behavioral strategy on a tree of depth 4, compute the reach probability of a specific leaf.
6. Show that the behavioral-strategy space has dimension linear in the number of information sets while the mixed-strategy space has dimension exponential in it.
7. **Argue, structurally,** why perfect recall is the minimal assumption that equates mixed and behavioral strategies, and why this equivalence is the single most important reason extensive-form game theory is computationally feasible.

Game Theory · Module 3 — Extensive Form Games · Node 3.3
Concepts: probability-space
Prerequisites: 3.1, 3.2
Enables: 3.4, 3.5, 4.4, 9.3

hbar.university — own.your.cognition