Module 1 — Strategic Form Games

Game Theory

hbar.university

Players, Strategies, and Payoffs

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.1

This is the foundation node of the entire game theory course. Every subsequent structural domain — equilibrium, extensive form, information, repetition, mechanism design, evolution, networks, learning, and strategic systems integration — is constructed on top of this single ontological move: the identification of a strategic situation with a triple **(players, strategies, payoffs)**. Before you can solve a game, you must specify what it *is*, and the normal-form definition is the canonical answer. It is the game-theoretic analogue of "Hilbert space + operators + states" in quantum mechanics, or "sample space + sigma algebra + measure" in probability theory: a minimal ontological skeleton onto which everything else hangs.

Formal Definition

A **strategic form game** (or **normal form game**) is a triple

$$G = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$$

where:

- $N = \{1, 2, \ldots, n\}$ is a finite set of **players**.
- S_i is the set of **pure strategies** available to player i. A pure strategy $s_i \in S_i$ is a complete specification of what player i will do. The **joint strategy space** is the Cartesian product $S = S_1 \times S_2 \times \cdots \times S_n$, and an element $s = (s_1, \ldots, s_n) \in S$ is called a **strategy profile**.
- $u_i : S \to \mathbb{R}$ is the **payoff function** of player i. It assigns a real-valued payoff to every strategy profile. The **payoff vector** at profile s is $u(s) = (u_1(s), \ldots, u_n(s)) \in \mathbb{R}^n$.

A game is **finite** if N is finite and every S_i is finite. A game is **two-player zero-sum** if $|N| = 2$ and $u_1(s) + u_2(s) = 0$ for all $s \in S$. A game is a **constant-sum** game if this sum is a constant for all profiles.

It is standard notation to write $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ for the strategy profile of all players *other* than i, so that $s = (s_i, s_{-i})$. Player i's payoff can then be written $u_i(s_i, s_{-i})$, emphasizing the variable the player actually controls.

The entire subsequent theory — dominance, best response, Nash equilibrium, mixed strategies, evolutionary stability, learning dynamics — takes this triple as its input and asks: *what strategy profiles can we justify as outcomes of rational strategic reasoning?*

Intuition

Three ingredients, each answering one question about the strategic situation:

1. **Who is interacting?** $\to$ Players N.
2. **What can each one do?** $\to$ Strategy sets $\{S_i\}$.
3. **What does each one care about?** $\to$ Payoff functions $\{u_i\}$.

The structural move is that payoffs depend on the *entire* profile, not just on a player's own choice. This is the feature that makes game theory *game* theory rather than *decision* theory. In a decision problem, you choose s_i to maximize $u_i(s_i)$ — a one-player optimization. In a game, you choose s_i to maximize $u_i(s_i, s_{-i})$, and s_{-i} is chosen by other agents whose own choices depend on yours. This mutual dependence is the source of both the mathematical difficulty and the conceptual richness of the field.

The payoff function is a full specification of preferences. It already bakes in risk attitudes, time preferences, altruism, and anything else you want to say about how the player values outcomes. The axiomatic foundation (von Neumann–Morgenstern expected utility, 1944) shows that under mild rationality assumptions, any preference ordering over outcome lotteries is representable by a real-valued utility function, which is what u_i is.

One subtle point: a strategy is a **complete plan of action**, not a single move. In a normal-form game this distinction is trivial (each s_i is just a label), but when we lift to extensive form (Module 3), a strategy will be "what I would do at every information set I could reach" — a function from my information states to actions. The normal form is the flattening of that richer object onto a finite-dimensional table.

Structural Role

1.1 is the minimal ontological commitment on which the entire course rests:

- **1.2 Pure vs mixed strategies** extends S_i from pure points to probability distributions $\Delta(S_i)$ over pure points, allowing randomization.
- **1.3 Dominance** uses u_i to compare strategies pointwise across opponent profiles.
- **1.4 Best response** is the function $B_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i})$, which takes opponent profiles as inputs and returns optimal own-strategies.
- **1.5 Zero-sum and minimax** restricts attention to games where $u_1 + u_2 = 0$, which has special structure (von Neumann's minimax theorem).
- **1.6 Normal form as universal representation** shows that every game, even an extensive-form one with thousands of decision nodes, can be flattened to this triple — possibly at the cost of an enormous strategy space.

From here, the entire equilibrium theory of Module 2 takes off: a **Nash equilibrium** is a profile s^* such that $s_i^* \in B_i(s_{-i}^*)$ for every i, which is a fixed point of the joint best-response correspondence. None of that can even be stated without the definitions of N, S_i, and u_i.

The definition also exports cleanly to every downstream domain. In extensive-form games (Module 3), S_i becomes the set of *plans* — functions from information sets to actions. In Bayesian games (Module 4), it becomes a set of *type-contingent* plans. In mechanism design (Module 6), the designer chooses u_i indirectly through the mechanism. In evolutionary games (Module 7), S_i is a population distribution and u_i is the expected fitness. The triple (N, S, u) is the rigid ontological core that survives every refinement.

Mathematical Development

The payoff matrix

For a two-player finite game, the payoff functions are conveniently displayed as a pair of matrices. Let $S_1 = \{a_1, \ldots, a_m\}$ be Row's strategies and $S_2 = \{b_1, \ldots, b_k\}$ be Column's. Define

$$A_{ij} = u_1(a_i, b_j), \qquad B_{ij} = u_2(a_i, b_j).$$

For a zero-sum game, $B = -A$ and only one matrix is needed. For a general-sum game, the "bimatrix" (A, B) — often written entry by entry as (A_{ij}, B_{ij}) — encodes the whole game.

Example: Matching Pennies. Two players each choose Heads or Tails. Row wins +1 if they match, Column wins +1 if they differ.

$$A = \begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}, \qquad B = -A.$$

Example: Prisoner's Dilemma. Two suspects each choose Cooperate (C) or Defect (D). Payoffs (u_1, u_2):

	C	D
C	$(3,3)$	$(0,5)$
D	$(5,0)$	$(1,1)$

Example: Stag Hunt. Two hunters each choose Stag (large cooperative prize) or Hare (small safe prize):

	Stag	Hare
Stag	$(4,4)$	$(0,3)$
Hare	$(3,0)$	$(3,3)$

These three games — Matching Pennies, Prisoner's Dilemma, Stag Hunt — span the strategic archetypes of the entire theory. Matching Pennies is irreducibly mixed (no pure equilibrium). Prisoner's Dilemma has a dominant strategy that produces a Pareto-inferior outcome. Stag Hunt has multiple equilibria with different risk profiles. Learn to recognize which type you're in.

Strategy profile space

The joint strategy space is the Cartesian product

$$S = \prod_{i \in N} S_i,$$

with cardinality $\prod_i |S_i|$. For an n-player game with k strategies each, $|S| = k^n$ — the strategy space grows exponentially in the number of players. This is the combinatorial cost of representing a game in normal form.

Best-response function (preview of 1.4)

$$B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}.$$

B_i is set-valued (a **correspondence**) because there may be ties. When the maximizer is unique for every s_{-i}, it is a genuine function. The central equilibrium concept of Module 2 — Nash equilibrium — is a fixed point of the joint correspondence $B(s) = (B_1(s_{-1}), \ldots, B_n(s_{-n}))$.

Utility scaling invariance

Payoff functions are only determined up to positive affine transformations: $u_i' = \alpha_i u_i + \beta_i$ with $\alpha_i > 0$ produces a strategically equivalent game. All solution concepts (dominance, best response, Nash equilibrium) are invariant under such transformations. This means the "units" of payoff are irrelevant — only the *ordering* and *ratios of differences* matter. In particular, you can add a constant to any row of Row's payoff matrix without changing the strategic content.

Worked Example

A concrete three-player game

Three competitors — Alice, Bob, Carol — each choose between Invest (I) or Wait (W). The payoffs, (u_A, u_B, u_C):

- If all three Invest: $(2, 2, 2)$. Market saturated.
- If two Invest, one Waits: investors get 1 each, waiter gets 4. Waiter free-rides.
- If one Invests, two Wait: investor gets -1, waiters get 0. Investor bears the cost alone.
- If all three Wait: $(0, 0, 0)$. Nothing happens.

Players: $N = \{A, B, C\}$. **Strategies:** $S_A = S_B = S_C = \{I, W\}$. **Profile space:** $S = \{I, W\}^3$, with 8 elements. **Payoffs:** $u_A, u_B, u_C : S \to \mathbb{R}$ defined by the list above.

Best response of Alice. Condition on $s_{-A} = (s_B, s_C)$:

- (I, I): Alice gets 2 (I) vs 4 (W) → play W.
- (I, W): Alice gets 1 (I) vs 0 (W) → play I.
- (W, I): Alice gets 1 (I) vs 0 (W) → play I.
- (W, W): Alice gets -1 (I) vs 0 (W) → play W.

So Alice's best-response function is "Invest iff exactly one of Bob/Carol Invests." This already hints at the equilibrium structure: the game has pure equilibria of the form "exactly two invest" (who the two are matters but the waiter free-rides) plus the all-Wait equilibrium where no one moves. Module 2 will make this precise.

Strategic archetype. This game is a three-player generalization of a "volunteer's dilemma" with free-riding. The interdependence is obvious: Alice's best action depends on Bob's and Carol's choices, and symmetrically.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Player set N	Seats in the hand (HU, 6-max, full ring — up to 9 seats)
Strategy set S_i	The complete set of decision rules available — your full strategy, not a single decision
Pure strategy s_i	A fully specified plan: "with AA, 3-bet to 9bb; with 72o, fold; with AK on Q-high flop after 3-betting, c-bet 75% pot; etc."
Mixed strategy σ_i	A randomized plan: "with AJs in MP vs UTG open, call 60% / 3-bet 40%"
Payoff $u_i(s)$	Expected winnings (in chips, BB, or \$) of player i given the whole-table strategy profile
Zero-sum structure	Cash poker (minus rake) is zero-sum across the players at the table

The key conceptual move: **a poker strategy is not "what I do with this hand in this spot" — it is the entire decision rule covering every possible hole-card / board / action-history combination.** When a solver outputs "the GTO strategy," it is outputting $s_i \in S_i$ in the game-theoretic sense: a function from every information state to an action distribution.

Concrete situation

Heads-up no-limit hold'em, 100 BB effective. You are the BTN (small blind), opponent is BB. Preflop only. Simplify: each player has either a strong hand (S, top 20% of hands), medium (M, next 30%), or trash (T, bottom 50%). Actions: you can Open (O) or Fold (F). If you Open, BB can Call (C) or Fold (F).

This is a genuine normal-form game once you enumerate types and actions. The strategy set for you is $S_{BTN} = \{O, F\}^3$ (one action per hand class) — **8 pure strategies**, not 2. For example, "Open with S, Open with M, Fold with T" is one pure strategy; "Open with everything" is another; "Fold with everything" is another. The strategy is type-contingent. Similarly BB has $S_{BB} = \{C, F\}^3$ if they commit to call or fold per type (after conditioning on BTN having opened — this is where information sets come in, Module 3).

Payoff. $u_{BTN}(s)$ is the expected chip winnings given the joint strategy, computed over the joint distribution of hand classes and the equity of each class vs opponent's calling range. This is exactly a normal-form payoff function.

What this understanding buys you

Recognizing that your "poker strategy" is the entire conditional plan — not a collection of individual decisions — is the move that unlocks everything. Every discussion of GTO, every solver output, every hand review is working with $s_i \in S_i$ in this sense. When a poker coach says "balance your range," they are saying "construct your s_i so that your action at each information set reveals as little as possible about your type." When they say "your line should be the same with your strong and weak hands in this spot," they are saying your s_i should assign the same action to both types at that information set. The ontology of 1.1 is the ontology of every serious strategic poker discussion.

Cross-Links

- **1.2 Pure vs mixed strategies** — extending S_i to probability distributions.
- **1.4 Best-response correspondences** — the central tool for analyzing strategic situations.
- **2.1 Nash equilibrium** — the fixed-point concept built directly on this triple.
- **3.1 Game trees** — the extensive-form perspective, which flattens back to normal form.
- **Statistics 2.x Decision theory** — single-player special case of this framework.
- **Systems thinking 2.1 State spaces** — S is the state space of the strategic system.
- **Economics** — preferences, choice, utility theory is the single-player foundation.

Structural Compression

Invariant: Every strategic situation reduces to a triple $(N, \{S_i\}, \{u_i\})$ of players, strategy spaces, and payoff functions, with payoffs depending on the entire joint profile.

Mechanism: Cartesian-product construction of the joint strategy space; real-valued payoff functions defined on the joint space; each player controls only their own coordinate.

Cross-System Echo: Compare to the postulates of quantum mechanics (state space + observables + unitary dynamics), the axioms of probability (sample space + sigma-algebra + measure), and the ontology of statistical models (data space + parameter space + likelihood). Each is a minimal ontological skeleton onto which a rich theory is then built. Game theory's skeleton is (N, S, u).

Exercises

1. Write out the payoff matrices (as a bimatrix) for the Battle of the Sexes game: two players choose between Opera and Football; each prefers being together, but Row prefers Opera and Column prefers

Football. Payoff 2 for preferred joint activity, 1 for non-preferred joint activity, 0 for being apart.

2. For the three-player volunteer game in the Worked Example, write out the payoff function u_A as a table indexed by (s_B, s_C) and verify the best-response computation in the text.

3. Show that adding a constant c_i to every payoff of player i (i.e. $u_i \to u_i + c_i$) does not change the set of best responses of any player. Then show that multiplying player i's payoffs by a positive scalar $\alpha_i > 0$ also preserves best responses. Conclude that solution concepts based on best responses are invariant under positive affine transformations.

4. Convert the following poker situation to normal form: two players, each dealt one card from a 3-card deck $\{A, K, Q\}$. Each antes 1 chip. Player 1 acts first and may Bet 1 or Check. If P1 Bets, P2 may Call or Fold. If P1 Checks, the hand goes to showdown. Higher card wins the pot at showdown. Write the strategy sets S_1, S_2 (each player's complete plan as a function of the card they hold) and the payoff function u_1 (average winnings). How many pure strategies does each player have?

5. A finite two-player game has $|S_1| = m$, $|S_2| = k$. How many distinct games are there up to a relabeling of strategies, if payoffs are drawn from $\{0, 1, 2\}$? How does this count grow with m, k?

6. A **dominant strategy** for player i is a strategy s_i^* such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_i' \in S_i$ and all $s_{-i} \in S_{-i}$. Show that in the Prisoner's Dilemma, Defect is a dominant strategy for each player. Then show that in the Stag Hunt, *no* strategy is dominant for either player.

7. **Argue, structurally**, why the normal-form triple (N, S, u) is strictly less expressive than an extensive-form description of the same situation. Identify at least one strategic fact (a feature of the situation that affects how rational players should reason) that the normal form loses compared to the extensive form. (Hint: think about the order in which players move, and what they know when they move.)

Pure vs Mixed Strategies

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.2

This node enlarges the strategy object. Node 1.1 defined a pure strategy as a deterministic choice from S_i. Many strategic situations — any situation whose opponent would exploit a deterministic rule — require randomization. A *mixed strategy* is a probability distribution over pure strategies, and the lift from S_i to the simplex $\Delta(S_i)$ is the single move that makes Nash equilibrium always exist in finite games (Module 2.4). Without this lift, most interesting games have no solution concept at all. It is also the node where expected utility enters game theory as the working notion of payoff.

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite strategic form game. A **mixed strategy** for player i is a probability distribution over S_i:

$$\sigma_i \in \Delta(S_i) = \left\{ \sigma_i : S_i \to [0,1] \;\middle|\; \sum_{s_i \in S_i} \sigma_i(s_i) = 1 \right\}.$$

The **support** of σ_i is the set $\operatorname{supp}(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}$. A **pure strategy** is the special case in which σ_i places unit mass on a single point; equivalently, it is a vertex of $\Delta(S_i)$.

A **mixed strategy profile** is $\sigma = (\sigma_1, \dots, \sigma_n) \in \Delta(S_1) \times \cdots \times \Delta(S_n)$. Under the assumption that players randomize independently, the profile induces a product distribution over $S = \prod_i S_i$:

$$\sigma(s) = \prod_{i \in N} \sigma_i(s_i).$$

Player i's **expected payoff** under the mixed profile is obtained by extending u_i linearly:

$$U_i(\sigma) = \sum_{s \in S} \sigma(s)\, u_i(s) = \sum_{s \in S} \left(\prod_{j \in N} \sigma_j(s_j) \right) u_i(s).$$

The extended payoff $U_i : \prod_j \Delta(S_j) \to \mathbb{R}$ is **multilinear**: linear in each σ_j holding the others fixed. This is the single most important technical fact about mixed strategies.

Intuition

There are three reasons to randomize.

First: to avoid being read. In a zero-sum game with no pure equilibrium — Matching Pennies is the canonical example — any deterministic plan can be exploited by the opponent. Randomization makes your action unpredictable, which is exactly the thing a hidden coin provides. The minimum information you can leak about your action is achieved precisely when you are drawing from a well-chosen distribution.

Second: to make opponents indifferent. In mixed equilibria, your distribution has the subtle property that your opponent is indifferent between the strategies in the support. This is a stability condition: if they were not indifferent, they would concentrate on one strategy, and the game would unravel.

Third: to represent a population. In evolutionary and population games (Module 7), a mixed strategy $\sigma_i(s_i)$ reinterprets as the fraction of the population playing s_i. The "randomization" is not in the head of any individual player; it is in the population distribution.

Multilinearity deserves a moment of its own. Because U_i is linear in each coordinate, an expected-payoff maximizer always has a best response that is pure: you cannot strictly prefer a mixture to every vertex it assigns weight to. This gives the fundamental lemma that we will lean on throughout Module 2: *every best response to σ_{-i} is a mixture whose support lies inside the set of pure best responses.*

Structural Role

Mixed strategies enter every downstream node:

- **1.3 Dominance.** A mixed strategy can strictly dominate a pure strategy that no single pure strategy dominates — so iterated elimination must be stated in terms of mixtures.
- **1.4 Best response.** The best-response correspondence extends from $S_{-i} \to S_i$ to $\Delta(S_{-i}) \to \Delta(S_i)$. The image is always a face of the simplex.
- **1.5 Minimax.** The minimax theorem for zero-sum games requires mixed strategies to hold; it fails in pure strategies.
- **2.3 Mixed-strategy Nash equilibrium.** The central equilibrium concept of the course is a fixed point of the best-response correspondence on $\prod_i \Delta(S_i)$.
- **2.4 Nash existence.** Kakutani's fixed-point theorem requires the domain $\prod_i \Delta(S_i)$ to be compact and convex — which it is, precisely because we allowed mixtures.
- **7 Evolutionary dynamics.** Replicator dynamics live on the mixed-strategy simplex.
- **9 Learning.** No-regret algorithms produce sequences of mixed strategies.

Without the lift from S_i to $\Delta(S_i)$, the entire equilibrium theory of Modules 2 and beyond is either vacuous or ill-defined.

Mathematical Development

The simplex $\Delta(S_i)$

For $|S_i| = m$, $\Delta(S_i)$ is an $(m-1)$-dimensional simplex embedded in $\mathbb{R}^m$:

$$\Delta(S_i) = \{x \in \mathbb{R}^m : x_k \geq 0,\ \textstyle\sum_k x_k = 1\}.$$

This set is **compact** (closed and bounded in $\mathbb{R}^m$) and **convex** (convex combinations of probability distributions are probability distributions). Both properties are essential for existence theorems.

Expected payoff is multilinear

Fix opponent profile σ_{-i}. Then

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot U_i(s_i, \sigma_{-i}),$$

where $U_i(s_i, \sigma_{-i}) = \sum_{s_{-i}} \sigma_{-i}(s_{-i}) u_i(s_i, s_{-i})$ is the expected payoff of the pure strategy s_i against the mixed profile σ_{-i}. This expression is linear in σ_i — each $\sigma_i(s_i)$ appears to first order. Consequently:

Best response to a mixed profile

$$\mathrm{BR}_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i, \sigma_{-i}).$$

By linearity, any maximizer is supported on the set $\arg\max_{s_i \in S_i} U_i(s_i, \sigma_{-i})$. If this argmax is a single pure strategy, it is the unique best response. If it contains several pure strategies, any mixture over them is also a best response. Conversely, no mixed strategy assigning positive weight to a non-best-response pure strategy can be a best response.

The indifference condition

If σ_i^* is a best response and $\mathrm{supp}(\sigma_i^*)$ contains two or more pure strategies, then all pure strategies in the support must yield the same expected payoff against σ_{-i}. This is the **indifference principle** and is the technical backbone of every mixed-equilibrium calculation.

Example: Matching Pennies

$$A = \begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}.$$

Let Column play $(q, 1-q)$. Row's expected payoff from Heads is $q - (1-q) = 2q - 1$; from Tails is $-q + (1-q) = 1 - 2q$. Row is indifferent iff $2q - 1 = 1 - 2q$, i.e. $q = 1/2$. By symmetry Row must play $(1/2, 1/2)$ in equilibrium. This is the unique mixed equilibrium and there is no pure one.

Example: Rock–Paper–Scissors

With payoff $+1$ win, -1 loss, 0 tie, the unique equilibrium mix is $(1/3, 1/3, 1/3)$ for both players, by a symmetric indifference argument.

Kuhn's purification theorem (preview)

Harsanyi (1973) showed that mixed equilibria of a game G can be recovered as limits of *pure* equilibria of a slightly perturbed game with small private noise in the payoffs. So "randomization" is not a philosophical commitment — it is the strategic residue of tiny type uncertainty. Module 4 develops this fully.

Worked Example

Attacker–Defender

A defender chooses which of two targets, $\{T_1, T_2\}$, to protect; an attacker simultaneously chooses one to attack. Target 1 has value $v_1 = 9$, target 2 has value $v_2 = 4$. If the attacker hits an unprotected target, the defender loses its value and the attacker gains it. If the attacker hits a protected target, payoffs are zero. (Zero-sum.)

Payoff matrix for the defender (rows = Defender, columns = Attacker):

$$A = \begin{pmatrix} 0 & -9 \\ -4 & 0 \end{pmatrix}.$$

Let the defender play $(p, 1-p)$ (probability of protecting T_1, then T_2). The attacker's expected gain from attacking T_1 is $9(1-p)$; from attacking T_2 is $4p$. The attacker is indifferent iff

$$9(1-p) = 4p \quad \Longleftrightarrow \quad p = \frac{9}{13} \approx 0.69.$$

Now let the attacker play $(q, 1-q)$. The defender's expected loss from protecting T_1 is $4(1-q)$; from protecting T_2 is $9q$. Indifference gives $4(1-q) = 9q \Rightarrow q = 4/13 \approx 0.31$. The unique mixed equilibrium is

$$\sigma_D^* = (9/13, 4/13), \qquad \sigma_A^* = (4/13, 9/13).$$

Notice the inversion: the defender protects the *more valuable* target with probability 9/13, and the attacker attacks the more valuable target with probability 4/13 — less than half. The high-value target is so well-defended that the attacker is better off going for the soft one more than half the time. This non-obvious inversion is the content of the mixed equilibrium.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Pure strategy s_i	A 100/0 decision: "always 3-bet this hand in this spot"
Mixed strategy σ_i	A frequency: "3-bet this hand 40%, call 60%"
Support $\text{supp}(\sigma_i)$	The set of actions you ever take with this exact hand + spot
Indifference condition	Solver balance: all actions in your mix must have the same EV
Multilinearity of U_i	EV of a mix is the average EV of its pure parts — you cannot mix into something better than your best pure action
Purification	Noise in chip values / rake rounding makes "mixed" play look like tiny private-value tilts

The table reveals a central fact: **every frequency a solver outputs is an indifference frequency.** If the solver tells you to 3-bet AJs 40% and call 60%, it is because at the current game state both 3-betting and calling yield the same EV — you are indifferent, and the 40/60 is chosen to make your *opponent* indifferent at some downstream node.

Concrete situation

BTN vs BB, 100 BB, single-raised pot. Flop comes $K\heartsuit 7\clubsuit 2\diamondsuit$ — a dry, BTN-range-favoring board. You (BTN) are considering your c-bet strategy with $A\clubsuit 5\clubsuit$.

A pure strategy is "always bet" or "always check." A mixed strategy is "bet x, check $1-x$." A modern solver output for this spot is approximately: bet small (25–33% pot) with $x \approx 0.75$, check with $1-x \approx 0.25$. Why a mix? Because with $A5s$, the EV of betting small and the EV of checking back are within a fraction of a BB of each other in equilibrium. The solver assigns exactly the weights that make the **opponent's** best-response correspondence balanced: BB's continue vs fold frequencies, their check-raise frequencies, and their turn play all line up only at this exact mix. Change the frequency, and one of BB's strategies becomes strictly better, and the whole equilibrium collapses.

Concretely: if you always bet with $A5s$, BB can exploit by calling more often with mid-pair holdings because you are bluffing too thin. If you always check, your check-back range becomes too strong, and BB can't profitably float. Mixing is how you keep both of BB's counter-strategies honest.

What this understanding buys you

Two immediate operational lessons.

First, you never need to beat yourself up about "choosing the wrong action" when you were in a mixed spot. By construction, the two actions have the same EV. If the solver says 50/50, then whichever one you took is exactly as good as the other — the only mistake is going 100/0 when the equilibrium calls for a mix.

Second, the practical way to "play a mixed strategy" at the table is to use an external randomizer (the second hand of a watch, the suit of your lowest hole card, an app). Humans are notoriously bad at internally generated randomness — we overproduce alternation. Any pattern in your "random" behavior is exactly what good opponents are reading. The indifference principle is fragile: it only works if your randomization is actually random.

Cross-Links

- **1.1 Players, strategies, payoffs** — the base definition this node extends.
- **1.3 Dominance** — iterated elimination must be stated for mixed strategies.
- **1.4 Best-response correspondences** — now taking values in $\Delta(S_i)$.
- **1.5 Zero-sum and minimax** — minimax only holds after the mixed-strategy lift.
- **2.3 Mixed-strategy Nash equilibrium** — the solution concept this node enables.
- **2.4 Nash existence theorem** — requires $\Delta(S_i)$ compact convex.
- **4 Information and beliefs** — Harsanyi purification relates mixing to type uncertainty.
- **7 Evolutionary dynamics** — mixed strategies as population distributions.
- **Statistics — expected value and linearity of expectation** — the purely technical foundation of U_i.

Structural Compression

Invariant: The strategy set of each player lifts from the finite set S_i to the simplex $\Delta(S_i)$, and the payoff function u_i lifts to its multilinear extension U_i. Pure strategies are the vertices of the simplex; mixed strategies are its interior points and faces.

Mechanism: Independent randomization on each coordinate produces a product distribution on the joint strategy space, and multilinearity of U_i gives an indifference characterization of best responses.

Cross-System Echo: The move from S_i to $\Delta(S_i)$ is the same move as from Boolean logic to probability, from classical states to quantum mixed states (density matrices), and from deterministic MDPs to stochastic policies. In each case, convexification of the decision object rescues existence of a solution.

Exercises

1. Verify that in Matching Pennies the unique Nash equilibrium is $((1/2, 1/2), (1/2, 1/2))$, and that no pure-strategy equilibrium exists. Compute each player's expected payoff at the equilibrium.

2. In the Attacker–Defender game from the worked example, change $v_1 = 16, v_2 = 4$. Recompute the equilibrium. Confirm that the defender still protects the more valuable target with probability strictly greater than 1/2, and that the attacker's attack probability on the high-value target goes *down* as v_1 increases.

3. Prove the indifference principle: if σ_i^* is a best response to σ_{-i} and $s_i, s_i' \in \text{supp}(\sigma_i^*)$, then $U_i(s_i, \sigma_{-i}) = U_i(s_i', \sigma_{-i})$.

4. Show that the expected payoff function $U_i(\sigma_i, \sigma_{-i})$ is continuous in (σ_i, σ_{-i}) as a function on $\prod_j \Delta(S_j)$. (This is used in the Nash existence proof.)
5. Consider a 3-strategy Rock–Paper–Scissors variant where Rock beats Scissors by a margin of 2, and the other wins are by margin 1 (standard). Find the unique mixed equilibrium. Is it still $(1/3, 1/3, 1/3)$? If not, explain what changes.
6. A player's strategy is called **completely mixed** if every pure strategy is in the support. Show that in a completely mixed equilibrium, every pure strategy in S_i gives the same expected payoff against σ^*_{-i}. Conclude that completely mixed equilibria satisfy $|S_i| - 1$ linear equations per player.
7. **Argue, structurally,** why the lift from S_i to $\Delta(S_i)$ is unavoidable if we want a general existence theorem for equilibrium in finite games. Specifically: give a finite game with no pure Nash equilibrium and explain exactly which step of the would-be existence proof fails in pure strategies but succeeds after the lift. (Hint: think about compactness and convexity of the domain of a fixed-point map.)

Dominance and Iterated Elimination

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.3

Dominance is the first solution technique in the course — a way to narrow down rational play without computing an equilibrium. A strategy that is worse than another strategy *regardless of what the opponents do* can be deleted from consideration. Iterating this deletion — on the grounds that rational players know rational players delete such strategies — gives the concept of **iterated elimination of dominated strategies**. It is the weakest-possible form of equilibrium reasoning, but surprisingly powerful: some games collapse to a unique profile under iterated strict dominance alone. It also sets up rationalizability (Node 2.5) and provides the sharpest intuition for Nash equilibrium (Module 2).

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite game, and allow mixed strategies $\Delta(S_i)$ as in Node 1.2.

Strict dominance. A strategy $\sigma_i \in \Delta(S_i)$ **strictly dominates** $s_i \in S_i$ if

$$U_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}.$$

The pure strategy s_i is then said to be **strictly dominated**. Note that the dominating strategy may itself be mixed, and that the inequality is required to hold for every opponent profile.

Weak dominance. σ_i **weakly dominates** s_i if

$$U_i(\sigma_i, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \text{for all } s_{-i}, \text{ with strict inequality for at least one } s_{-i}.$$

Iterated elimination of strictly dominated strategies (IESDS). Define the sequence

$$S_i^{(0)} = S_i, \qquad S_i^{(k+1)} = \{s_i \in S_i^{(k)} : s_i \text{ is not strictly dominated given } S_{-i}^{(k)}\},$$

where the dominance check in the $(k+1)$-th round uses mixtures over $S_i^{(k)}$ and opponent profiles in $S_{-i}^{(k)}$. The set of strategies surviving iterated elimination is

$$S_i^{\infty} = \bigcap_{k \geq 0} S_i^{(k)}.$$

A game is **dominance-solvable** if $|S_i^{\infty}| = 1$ for every i.

Intuition

A strictly dominated strategy is never optimal. Not once. Not against any opponent. A rational player never plays it. So, the reasoning goes, a rational player's choice must lie in $S_i \setminus \{\text{strictly dominated}\}$. A rational player who knows the opponent is rational also knows the opponent's choice lies in the opponent's reduced set, and can therefore delete any of their own strategies that are strictly dominated *on the reduced opponent set*. This recursive structure — "I am rational, I know you are rational, I know you know I am rational" — is **common knowledge of rationality**, and it is exactly what iterated elimination operationalizes.

Two warnings:

Strict vs weak. Iterated elimination of *strictly* dominated strategies is order-independent: whichever order you delete in, you get the same S_i^∞. Iterated elimination of *weakly* dominated strategies is not order-independent — different orders give different surviving sets. This makes weak dominance a delicate solution concept and strict dominance a clean one.

Mixing matters. A pure strategy can be strictly dominated by a mixture without being strictly dominated by any single pure strategy. So iterated elimination *must* allow mixtures on the dominating side, or it will under-delete.

Structural Role

Dominance is upstream of every equilibrium concept. Specifically:

- **Any Nash equilibrium lies in S^∞.** A strictly dominated strategy cannot be a best response, so cannot be in the support of any equilibrium.
- **Rationalizability (2.5)** is the fixed point of a similar elimination procedure that uses best-response correspondences instead of dominance. It is a strictly weaker concept than dominance-solvability for 3+ players and coincides with iterated strict dominance for 2 players.
- **Dominance-solvable games** are exactly the games where common knowledge of rationality pins down a unique profile — no further equilibrium assumption is needed.

Dominance also acts as a *pre-processor*: before computing Nash equilibria of a big game, you can delete all strictly dominated strategies and compute on the reduced game, and every Nash equilibrium of the original is an equilibrium of the reduced game, and vice versa.

Mathematical Development

Order-independence for strict dominance

Theorem. The limit S_i^∞ of iterated strict-dominance elimination is independent of the order in which strategies are deleted.

Sketch. Strict dominance is a monotone operator: deleting a strategy never revives a previously dominated one, because the inequality $U_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ continues to hold when opponent profiles are restricted to a subset. The intersection of all maximal sequences of deletions is therefore the same set.

Non-monotonicity for weak dominance

Consider

$$\begin{array}{c|cc} & L & R \\ \hline T & (1,1) & (0,0) \\ M & (1,1) & (2,1) \\ B & (0,0) & (2,1) \end{array}.$$

T is weakly dominated by M; B is weakly dominated by M. Different orders of elimination lead to different survivors, because weak dominance is lost when opponent strategies supporting the strict inequality are themselves deleted first.

Mixed dominance exceeds pure dominance

$$\begin{array}{c|cc} & L & R \\ \hline T & (3,0) & (0,0) \\ M & (0,0) & (3,0) \\ B & (1,0) & (1,0) \end{array}$$

B is not dominated by T alone (fails against L) or by M alone (fails against R). But the mixture $\frac{1}{2}T + \frac{1}{2}M$ yields 1.5 in both columns, strictly better than B's 1. So B is strictly dominated by a mixture.

Dominance-solvable games

Cournot duopoly with linear demand is dominance-solvable: iteratively eliminating never-best-response quantities converges to the unique Cournot equilibrium. **Iterated beauty contest** (guess a number, the winner is closest to p times the average, with $p < 1$) is dominance-solvable: successive rounds of "no one rational picks above $100p$", "no one rational picks above $100p^2$", and so on, collapse the game to 0. These are the canonical demonstrations that iterated strict dominance has real solving power.

Worked Example

A 3×3 elimination

$$\begin{array}{c|ccc} & L & C & R \\ \hline T & (4,3) & (1,2) & (2,1) \\ M & (3,2) & (2,1) & (1,3) \\ B & (0,0) & (0,0) & (1,1) \end{array}$$

Round 1 (Column). Column's payoffs are $L : (3,2,0)$, $C : (2,1,0)$, $R : (1,3,1)$. Try the mix $(0.6L, 0.4R)$: $(0.6 \cdot 3 + 0.4 \cdot 1, 0.6 \cdot 2 + 0.4 \cdot 3, 0.6 \cdot 0 + 0.4 \cdot 1) = (2.2, 2.4, 0.4)$. Compared to $C = (2,1,0)$, the mixture strictly dominates C. Delete C.

Round 2 (Row, on reduced game).

$$\begin{array}{c|cc} & L & R \\ \hline T & (4,3) & (2,1) \\ M & (3,2) & (1,3) \\ B & (0,0) & (1,1) \end{array}.$$

Row payoffs: $T : (4,2)$, $M : (3,1)$, $B : (0,1)$. T strictly dominates M $(4 > 3, 2 > 1)$. Delete M. Does T strictly dominate B? $4 > 0$ and $2 > 1$ — yes. Delete B.

Round 3. Row is pinned to T. Column chooses between $(4,3)$ and $(2,1)$; Column prefers L. The game is dominance-solvable with outcome (T, L), payoff $(4,3)$.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Strictly dominated strategy	A line that is EV-worse than another line for *every* opponent response (open-limping 72o UTG is dominated by folding)
Mixed dominance	A mixture of two lines beats a pure line — e.g., balanced bet/check strictly dominates "always min-raise" with a polarized holding
IESDS	Preflop range construction: "don't open into 4 limpers with $K6o$" is iterated dominance
Dominance-solvable subgame	Face-up river decisions in a solver — the best action is often pinable by pointwise comparison

Formal object	Poker instantiation
Rationalizability preview	The set of “defensible plays” — the GTO range of openable hands

Concrete situation

UTG in a 9-handed $1/$2 live game, stack 100 BB. You are dealt 7♠2♡.

Pure-strategy analysis. Your actions are Fold, Limp, Open-raise. Against any reasonable opponent strategy profile, Fold strictly dominates Limp: your EV from limping $72o$ UTG is negative (you will rarely flop anything playable, and when you do the effective stack is deep enough that you lose more than you win). Fold also strictly dominates Open-raise: raising $72o$ from UTG is burning chips. Hence both Limp and Open-raise are strictly dominated by Fold, and an IESDS pre-processing eliminates them.

Repeat this for every garbage hand in UTG. After one round of IESDS you have the UTG opening range: roughly the top 15% of hands open, everything else folds.

Now move to MP. MP’s decision depends on what UTG did. If UTG opened, MP knows UTG has a strong range (top 15%), so MP’s strategy is computed in the reduced game. Some MP hands that would have been opens in isolation (say, ATo) become dominated by fold *in the reduced game* because they fare badly against UTG’s concentrated strong range. This is IESDS iterating: round 2 eliminates strategies that were only rational under round-0 opponent behavior.

What this understanding buys you

The preflop ranges you memorize are the output of IESDS. The charts you study are not arbitrary — they are the dominance-reduced action set. When you see a hand and cannot recall the chart, you can often reconstruct it by asking: “Is there any opponent response under which opening this hand is better than folding?” If the answer is no under every reasonable opponent range (the set that *they* have not had eliminated), then folding is the rational play.

Second lesson: the hands you agonize over are exactly the ones *not* eliminated by dominance — the rationalizable but non-dominance-pinned hands. The AJo from MP, the suited connectors UTG+1. Those are the hands where equilibrium mixing matters. For all the obviously foldable and obviously raiseable hands, dominance already tells you what to do, and you should spend no thought at all.

Cross-Links

- **1.1 Players, strategies, payoffs** — the objects dominance operates on.
- **1.2 Mixed strategies** — required for mixed dominance.
- **1.4 Best response** — the functional companion to dominance.
- **2.1 Nash equilibrium** — every Nash profile survives IESDS.
- **2.5 Rationalizability** — the best-response analogue of iterated dominance.
- **Economics — Cournot and Bertrand** — classic dominance-solvable models.
- **AI Systems — minimax pruning** — dominance as tree-pruning in game search.

Structural Compression

Invariant: A strictly dominated strategy can never be a best response, for any belief about opponents; consequently no rational player ever plays one, and no equilibrium assigns it positive weight.

Mechanism: Pointwise comparison of payoff functions across opponent profiles, extended to mixtures to capture all dominance relations, iterated over rounds that model increasing depth of common knowledge of rationality.

Cross-System Echo: Dominance plays the role of "obviously irrational" pruning in many settings: dead-code elimination in compilers, constraint propagation in CSPs, alpha-beta pruning in game tree search, strictly-dominated action elimination in multi-armed bandits. In each case, a monotone pre-processing step removes options that can be proved suboptimal without a full solution.

Exercises

1. Apply IESDS to the Prisoner's Dilemma. Show that Defect strictly dominates Cooperate for each player, and conclude the game is dominance-solvable with outcome (D, D).
2. In the Stag Hunt of Node 1.1, show that *no* strategy is strictly dominated. (Dominance gives no information — equilibrium refinement is needed.)
3. Construct a 2×3 game in which a pure strategy of Row is strictly dominated by a mixture of two other pure strategies but by no single pure strategy.
4. Prove order-independence of IESDS: if s_i is strictly dominated at some round of elimination, and s_i' is eliminated in that round, then s_i remains strictly dominated in the reduced game.
5. In the p-beauty contest with $p = 2/3$ on the interval $[0, 100]$, show that after k rounds of IESDS, every surviving strategy is in $[0, 100(2/3)^k]$. Conclude the game is dominance-solvable with unique outcome 0.
6. Show that if a strategy profile s^* is a (pure) Nash equilibrium, then each $s_i^* \in S_i^\infty$. (Contrapositive: a dominated strategy cannot be part of any equilibrium.)
7. **Argue, structurally,** why iterated elimination of *weakly* dominated strategies is order-dependent while iterated elimination of *strictly* dominated strategies is not. Identify the precise step in the order-independence argument that uses strictness of the inequality.

Best Response Correspondences

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.4

Dominance (Node 1.3) prunes strategies that are bad against every opponent profile. The **best-response correspondence** is the positive dual: it tells you exactly which strategies are optimal *against a given opponent profile*. Nash equilibrium will be defined as a fixed point of the joint best-response correspondence (Module 2), so this node is the direct precursor to the central solution concept of the course. It is also the object whose continuity, convex-valuedness, and upper-hemicontinuity power the Nash existence proof via Kakutani's theorem (Node 2.4).

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite game with mixed extension $\Delta(S_i)$ and multilinear expected-payoff U_i.

Best response (pure version). Given $s_{-i} \in S_{-i}$,

$$\mathrm{BR}_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}) \subseteq S_i.$$

Best response (mixed version). Given an opponent mixed profile $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$,

$$\mathrm{BR}_i(\sigma_{-i}) = \arg\max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i, \sigma_{-i}) \subseteq \Delta(S_i).$$

Joint best-response correspondence. Define

$$\mathrm{BR} : \prod_i \Delta(S_i) \rightrightarrows \prod_i \Delta(S_i), \qquad \mathrm{BR}(\sigma) = \mathrm{BR}_1(\sigma_{-1}) \times \cdots \times \mathrm{BR}_n(\sigma_{-n}).$$

A **Nash equilibrium** is a profile σ^* with $\sigma^* \in \mathrm{BR}(\sigma^*)$ — a fixed point of the joint correspondence.

Because BR_i may return multiple optima (ties are common when opponents mix), it is generally a *correspondence* (set-valued map), not a function. Writing $\mathrm{BR}_i : \cdot \rightrightarrows \cdot$ is standard notation for a correspondence.

Intuition

Three ways to think about best response.

As an oracle. If a small demon tells you exactly what your opponents will do, the best response is simply what you should do given that information. It is the single-player decision problem that emerges inside a game once the opponents' play is fixed.

As a reaction curve. In economic applications (e.g., Cournot duopoly), the best response is a function mapping the opponent's quantity to your optimal quantity. Plotting it gives the "reaction function"; the intersection of reaction functions is the Nash equilibrium.

As a fixed-point operator. The joint best response BR is a map from profiles to sets of profiles. An equilibrium is a profile that is returned by the map when fed into the map. This lifts the problem from "solve for equilibrium" to "find a fixed point of an operator," which is a *much* more powerful framing because we have centuries of fixed-point theorems to apply.

The set-valued nature is essential: when two pure strategies tie in expected payoff against σ_{-i}, both are best responses, and so is any convex combination of them. Ignoring this would miss most mixed Nash equilibria.

Structural Role

Best response is the joint pivot between dominance and equilibrium:

- **1.3 ⇒ 1.4.** A strictly dominated strategy is never a best response. The best-response graph lives inside the dominance-reduced game.
- **1.4 ⇒ 2.1.** Nash equilibrium is defined as a fixed point of BR.
- **1.4 ⇒ 2.4.** The existence proof (Kakutani) checks that BR is upper-hemicontinuous with non-empty compact convex values — each of these hypotheses is a property of BR, not of the payoffs directly.
- **1.4 ⇒ 2.5.** Rationalizability is the fixed point of iterated best-response elimination.
- **1.4 ⇒ 9.1.** Fictitious play iterates best responses against empirical frequencies.
- **1.4 ⇒ 9.3.** Counterfactual regret minimization is a weakening of best-response dynamics that guarantees convergence of time averages.

Best response is the functional heart of the course: almost every theorem and almost every algorithm in the remaining modules is phrased as a statement about BR.

Mathematical Development

Characterization of mixed best response

By linearity of U_i in σ_i, a mixed strategy σ_i is in $\mathrm{BR}_i(\sigma_{-i})$ if and only if

$$\mathrm{supp}(\sigma_i) \subseteq \arg\max_{s_i \in S_i} U_i(s_i, \sigma_{-i}).$$

Equivalently, the pure strategies in $\mathrm{supp}(\sigma_i)$ must all tie for first place in expected payoff, and every pure strategy outside the support must do no better. This is the **support characterization** of best response.

Properties of BR_i

Non-empty: $\Delta(S_{-i})$ is compact and U_i is continuous, so the argmax is non-empty. **Compact:** the argmax is a closed subset of the compact simplex. **Convex:** any convex combination of two best-response mixtures has support inside the same argmax set, hence is also a best response. (This uses multilinearity crucially.) **Upper hemicontinuous:** if $\sigma_{-i}^{(k)} \to \sigma_{-i}$ and $\sigma_i^{(k)} \in \mathrm{BR}_i(\sigma_{-i}^{(k)})$ with $\sigma_i^{(k)} \to \sigma_i$, then $\sigma_i \in \mathrm{BR}_i(\sigma_{-i})$. This is the Berge Maximum Theorem applied to a continuous objective on a compact domain.

These four properties — non-empty, compact, convex, upper-hemicontinuous — are exactly the hypotheses of Kakutani's fixed-point theorem.

Reaction functions in continuous games

For games with continuous action spaces and strictly concave payoff in own action (e.g., Cournot with quadratic costs and linear demand), the argmax is a singleton, and the best-response correspondence reduces to a genuine function $r_i(s_{-i})$. The reaction function is differentiable under standard regularity conditions, and its slope at equilibrium carries information about stability of best-response dynamics.

Cournot example. Two firms produce quantities $q_1, q_2 \geq 0$; price is $P(Q) = a - Q$ with $Q = q_1 + q_2$; marginal cost c. Firm 1's payoff is $\pi_1 = (a - q_1 - q_2 - c)q_1$. First-order condition:

$$\frac{\partial \pi_1}{\partial q_1} = a - 2q_1 - q_2 - c = 0 \quad \Longrightarrow \quad r_1(q_2) = \frac{a - c - q_2}{2}.$$

Symmetric for firm 2. The Nash equilibrium is the intersection $q^* = (a - c)/3$ each.

Fixed-point formulation

A profile $\sigma^* \in \prod_i \Delta(S_i)$ is a Nash equilibrium iff σ^* is a fixed point of the self-map $f(\sigma) := \mathrm{BR}(\sigma)$ — meaning $\sigma^* \in \mathrm{BR}(\sigma^*)$ in the set-valued sense. All of Module 2 builds on this formulation.

Worked Example

Cournot with three firms

Three symmetric firms produce $q_i \geq 0$; demand $P = 12 - Q$ with $Q = q_1 + q_2 + q_3$; marginal cost $c = 0$. Firm i's profit:

$$\pi_i(q_i, q_{-i}) = (12 - q_i - q_{-i,\mathrm{sum}})q_i, \qquad q_{-i,\mathrm{sum}} = \sum_{j \neq i} q_j.$$

FOC: $12 - 2q_i - q_{-i,\mathrm{sum}} = 0 \Rightarrow r_i(q_{-i,\mathrm{sum}}) = (12 - q_{-i,\mathrm{sum}})/2$.

At a symmetric equilibrium $q_1 = q_2 = q_3 = q^*$, the opponent sum is $2q^*$. The fixed-point equation becomes

$$q^* = (12 - 2q^*)/2 = 6 - q^* \quad \Longrightarrow \quad q^* = 3.$$

Each firm produces 3 units, total output is 9, price is 3, and each firm earns 9. This is the unique symmetric Nash equilibrium, obtained entirely by setting $\sigma^* = \mathrm{BR}(\sigma^*)$.

Discrete analogue. Consider Matching Pennies. Row's best response to Column playing $(q, 1-q)$:

$$\mathrm{BR}_1(q) = \begin{cases} \{\text{Heads}\} & q > 1/2 \\ \Delta(\{H, T\}) & q = 1/2 \\ \{\text{Tails}\} & q < 1/2. \end{cases}$$

The fixed point is $q = 1/2$, where both players mix 50/50. Notice BR_1 is convex-valued at $q = 1/2$ — the entire simplex is the best-response set — which is what allows the fixed point to land there.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
$\mathrm{BR}_i(\sigma_{-i})$	The maximally exploitative strategy against a fixed opponent strategy
$\mathrm{BR}(\sigma^*) = \sigma^*$	GTO — every player is already playing maximally against every other
Best response as exploit	"Nodelocking" a solver to opponent tendencies and reading the exploitative strategy off the solution
Upper hemicontinuity of BR	Small changes in opponent frequencies lead to small changes in the exploit — until a threshold is crossed, then the exploit jumps
Convex-valued	Any mixture of two tied best responses is also a best response — "mix at indifferences"

Concrete situation

100 BB deep, BTN vs BB. BTN opens 2.5x, BB calls. Flop $K\spadesuit 9\heartsuit 4\diamondsuit$. BTN bets 33% pot. BB's raise frequency in equilibrium is around 8% on this board with a specific range.

Exploit-by-best-response. Suppose you (BTN) are facing a fish whose raise frequency on this board texture is empirically around 2% — they almost never raise. What is your best response? Given the lower raise frequency, your c-bet EV is higher because you face less resistance; in particular, you can c-bet **wider** than solver recommends because the punishment for c-betting a weak range (getting raised off it) is reduced. Concretely, solver says c-bet 62% of range; exploit says c-bet $\approx$ 80%.

Similarly, against a fish whose c-bet defense is too folding-heavy, your best response is to c-bet 100% of range at a smaller size. Against a fish whose c-bet defense is too calling-heavy, your best response is to reduce bluff frequency and value-bet thinner.

These are literally best-response computations in the sense of this node. You fix σ_{-i} (the fish's strategy) and compute $\mathrm{BR}_i(\sigma_{-i})$ — the strategy that maximizes your EV against the observed opponent. Solvers support this via "nodelocking": you freeze the opponent's strategy at some node and re-solve your own play.

What this understanding buys you

Exploit = deviate from GTO toward the best response against observed opponent. GTO is the Nash equilibrium — the fixed point of the joint best-response map. If the opponent is not at their own GTO (they are off-equilibrium), then your GTO is not a best response to them — it is only guaranteed not to *lose*. The best response to their specific mistake is strictly more profitable than GTO.

The art of live poker is **when** to deviate toward the best response and **how far**. Deviate too little, and you leave money on the table; deviate too much, and you become exploitable in return if the opponent adjusts. The best-response correspondence is the map you are walking along when you move from GTO toward pure exploit — and the upper-hemicontinuity of BR is what guarantees small reads translate into small exploits, not wild swings.

Cross-Links

- **1.1 Players, strategies, payoffs** — inputs to the BR operator.
- **1.2 Mixed strategies** — the domain on which BR lives.
- **1.3 Dominance** — dominated strategies lie outside the image of BR.
- **2.1 Nash equilibrium** — fixed point of BR.
- **2.4 Nash existence** — BR's properties power the Kakutani proof.
- **9.1 Fictitious play** — iterated best response to empirical frequencies.
- **Economics — reaction functions** — the differentiable case.
- **Systems thinking — fixed points** — the unifying structural concept.

Structural Compression

Invariant: For every opponent profile, there is a non-empty, convex, compact set of own-strategies that maximize expected payoff, and this set varies upper-hemicontinuously with the opponent profile.

Mechanism: The Berge Maximum Theorem applied to the multilinear objective U_i on the compact convex domain $\Delta(S_i)$ — continuity plus compactness of the parameter space gives the well-behavedness of the correspondence.

Cross-System Echo: Best-response correspondences are the multi-agent analogue of the gradient in single-agent optimization: they tell you which direction to move given the current state of the environment. The fixed-point formulation of equilibrium is the multi-agent analogue of the first-order condition in calculus.

Exercises

1. Compute $\mathrm{BR}_1(q)$ for Matching Pennies as a function of $q = \Pr[\text{Column plays H}]$. Verify it is convex-valued only at $q = 1/2$.
2. For symmetric Cournot with n firms, demand $P = a - Q$, marginal cost c, derive the best-response function of firm i and the symmetric Nash equilibrium $q^* = (a - c)/(n + 1)$.
3. Prove that $\mathrm{BR}_i(\sigma_{-i})$ is always convex: if $\sigma_i, \sigma_i' \in \mathrm{BR}_i(\sigma_{-i})$, then any convex combination is also in $\mathrm{BR}_i(\sigma_{-i})$.
4. Prove that BR_i is upper-hemicontinuous. (You may use the Berge Maximum Theorem.)
5. Give an example of a game in which BR_i is not *lower*-hemicontinuous at some point. (Lower-hemicontinuity is the non-trivial condition — it does not follow from continuity of payoffs.)
6. Verify the support characterization of best response: $\sigma_i \in \mathrm{BR}_i(\sigma_{-i})$ iff every pure strategy in $\mathrm{supp}(\sigma_i)$ yields the maximal expected payoff.
7. **Argue, structurally,** why Nash equilibrium is naturally defined as a fixed point of the joint best-response correspondence BR rather than as a solution of some system of payoff equations directly. What is the mathematical advantage of phrasing equilibrium as a fixed-point problem on a correspondence?

Zero-Sum Games and the Minimax Theorem

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.5

Zero-sum games are the original setting of game theory. Von Neumann's 1928 minimax theorem is the first major existence result in the field and is historically the moment the subject was born. It shows that every finite two-player zero-sum game has a *value* — a unique number that the row player can guarantee from below and the column player can guarantee from above — and that this value is achieved by a pair of mixed strategies that are mutual best responses. This node is the specialized prelude to the general Nash existence theorem (2.4): zero-sum minimax is the case where the equilibrium computation reduces to a linear program. It is also the mathematical home of modern poker solving, because the two-player heads-up subgame is a zero-sum game.

Formal Statement

Let $G = (\{1, 2\}, S_1, S_2, u_1, u_2)$ be a finite two-player game with $u_1 + u_2 = 0$. Writing $A_{ij} = u_1(s_i, s_j)$ for the payoff matrix, player 1 maximizes A and player 2 minimizes it.

Define the **maximin value**

$$\underline{v} = \max_{\sigma_1 \in \Delta(S_1)} \min_{\sigma_2 \in \Delta(S_2)} \sigma_1^\top A \sigma_2,$$

and the **minimax value**

$$\overline{v} = \min_{\sigma_2 \in \Delta(S_2)} \max_{\sigma_1 \in \Delta(S_1)} \sigma_1^\top A \sigma_2.$$

Minimax Theorem (von Neumann, 1928). For every finite zero-sum game,

$$\underline{v} = \overline{v}.$$

This common value is the **value of the game** $v(G)$, and the maximin–minimax strategies form a Nash equilibrium.

Two crucial facts:

1. The inequality $\underline{v} \le \overline{v}$ is trivial (swapping max and min never improves the objective).
2. The reverse inequality $\overline{v} \le \underline{v}$ is the non-trivial theorem and requires the lift from S_i to $\Delta(S_i)$. In pure strategies, $\underline{v} < \overline{v}$ is possible (a "saddle point" need not exist).

Intuition

In a zero-sum game, one player's gain is the other's loss. There is no "win–win"; strategic interaction is pure conflict. In that setting, the question "what should I play?" has a particularly clean answer: play the strategy that maximizes the *worst* outcome your opponent can impose on you.

The maximin strategy is the "paranoid" strategy: assume the opponent will play to hurt you as much as possible, and choose accordingly. In a zero-sum game, that is also the "rational" strategy, because the opponent really *is* trying to minimize your payoff — that is literally what they maximize.

The minimax theorem says this paranoid reasoning has a fixed point: there is a value that both players can simultaneously guarantee against the other's worst behavior. This value is not too low (row can force it) and not too high (column can force it against row). It is unique.

The linear-programming reformulation (below) is both the fastest proof and the fastest computational method: finding the value of a zero-sum game is equivalent to solving a linear program.

Structural Role

Zero-sum minimax is the simplest non-trivial case of Nash existence:

- **1.4 ⇒ 1.5.** Maximin strategies are best-response-optimal; minimax-optimal strategies are mutual best responses in the zero-sum context.
- **1.5 ⇒ 1.6.** Minimax provides a worst-case lower bound for the normal-form representation of any game.
- **1.5 ⇒ 2.1.** The unique minimax pair is a Nash equilibrium of the zero-sum game.
- **1.5 ⇒ 2.4.** The Nash existence theorem generalizes the zero-sum case to general-sum games, replacing linear programming duality with a fixed-point theorem.
- **1.5 ⇒ 9.2.** No-regret learning on zero-sum games converges to the value; this is a dynamic realization of the static minimax result.

Zero-sum games are also the unique setting in which "equilibrium" and "security" coincide: the minimax strategy is both a Nash equilibrium component and the solution to the paranoid decision problem. In general-sum games, these two objects diverge, and the strategic question "what should I play?" becomes ambiguous.

Mathematical Development

Proof via LP duality

Player 1 wants to maximize $\min_j(\sigma_1^\top A)_j$. Introducing v as the minimum, this becomes

$$\max_{\sigma_1, v} \; v \quad \text{s.t.} \quad (\sigma_1^\top A)_j \geq v \text{ for all } j, \; \sum_i \sigma_1(i) = 1, \; \sigma_1 \geq 0.$$

This is a linear program. Its dual is

$$\min_{\sigma_2, w} \; w \quad \text{s.t.} \quad (A\sigma_2)_i \leq w \text{ for all } i, \; \sum_j \sigma_2(j) = 1, \; \sigma_2 \geq 0.$$

This is exactly the minimax LP for player 2. LP duality (strong duality holds under feasibility) gives $v^* = w^*$, which is the statement of the theorem.

The saddle-point characterization

A pair (σ_1^*, σ_2^*) is a **saddle point** of A if

$$\sigma_1^\top A \sigma_2^* \leq \sigma_1^{*\top} A \sigma_2^* \leq \sigma_1^{*\top} A \sigma_2 \quad \text{for all } \sigma_1, \sigma_2.$$

Minimax equilibria are exactly saddle points. The saddle-point structure is what makes zero-sum games tractable: the interests of the two players are perfectly anti-aligned, so a single function's saddle encodes both best-response conditions.

Security level and value

For player 1, the **security level** is $\max_{\sigma_1} \min_{\sigma_2} \sigma_1^\top A \sigma_2 = \underline{v}$. The minimax theorem says the security level is attained and equals $\overline{v}$. In pure strategies:

$$\underline{v}_{\text{pure}} = \max_i \min_j A_{ij}, \qquad \overline{v}_{\text{pure}} = \min_j \max_i A_{ij}.$$

These are the "pure-strategy security levels." Without mixing, $\underline{v}_{\text{pure}} < \overline{v}_{\text{pure}}$ is possible. Matching Pennies is the canonical example: $\underline{v}_{\text{pure}} = -1, \overline{v}_{\text{pure}} = +1$, but $v = 0$ in mixed strategies.

Continuity and convexity arguments

A second (more general) proof: the function $f(\sigma_1, \sigma_2) = \sigma_1^\top A\sigma_2$ is bilinear, hence continuous, concave in σ_1, and convex in σ_2 on the compact convex domains. Sion's minimax theorem then gives $\max_{\sigma_1} \min_{\sigma_2} f = \min_{\sigma_2} \max_{\sigma_1} f$. This is the "convex-analytic" form of the theorem and lifts directly to infinite action spaces under appropriate continuity assumptions.

Computing the value

For a 2×2 zero-sum game with no pure saddle, the value is found by the indifference principle. Let $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, with the interior mixed equilibrium. Row's equilibrium mix $(p, 1-p)$ satisfies

$$pa + (1-p)c = pb + (1-p)d, \quad p = \frac{d-c}{a-b-c+d}.$$

The value is $v = pa + (1-p)c = \frac{ad-bc}{a-b-c+d}$.

Worked Example

A 3×3 zero-sum game

$$A = \begin{pmatrix} 2 & -1 & 0 \\ -1 & 3 & -2 \\ 0 & -2 & 4 \end{pmatrix}.$$

Pure-strategy analysis. $\underline{v}_{\text{pure}} = \max_i \min_j A_{ij} = \max(-1, -2, -2) = -1$. $\overline{v}_{\text{pure}} = \min_j \max_i A_{ij} = \min(2, 3, 4) = 2$. Since $-1 < 2$, no pure saddle exists.

LP for Row. Maximize v subject to $2p_1 - p_2 + 0p_3 \geq v$, $-p_1 + 3p_2 - 2p_3 \geq v$, $0p_1 - 2p_2 + 4p_3 \geq v$, $p_1 + p_2 + p_3 = 1$, $p_i \geq 0$. (Solving this LP is a computer exercise; the value of the game comes out to approximately $v \approx 0.769$.)

Interpretive takeaway. Row has a "strong" row 1 (positive against col 2 is its worst: -1), and a "strong" row 3 (positive against col 3 with $+4$ but weak against col 2 with -2). The equilibrium balances these to achieve a guaranteed payoff that neither row alone can secure.

A clean analytical case

Consider

$$A = \begin{pmatrix} 0 & -1 & 3 \\ 2 & 3 & -1 \\ -3 & 1 & 2 \end{pmatrix}.$$

This can be solved via the LP or by inspection of support cardinality. In practice, for anything bigger than 2×2 by hand, the LP is the way.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Zero-sum assumption	Heads-up cash poker (minus rake) is zero-sum
Minimax value $v(G)$	The expected chip EV of GTO play against GTO opponent — often exactly 0 modulo position
Maximin strategy	The "unexploitable" strategy — guarantees at least the game value regardless of opponent
LP formulation	Linear-programming solutions of HU limit poker — CFR is an iterative alternative
Saddle point	HU GTO pair — mutual best responses where neither player can improve unilaterally

Concrete situation

Heads-up limit hold'em (HULHE) was solved to optimality in 2015 by the Cepheus program (Bowling et al., 2015) via counterfactual regret minimization, a no-regret algorithm that provably converges to the minimax value of a zero-sum game. The solution is a set of mixed strategies — one per information set — that together constitute a near-optimal pair of saddle-point strategies.

Concrete number: the HULHE value for the small-blind (dealer) is approximately +0.088 BB per hand, rather than exactly 0. Why? Because the ante/blind structure is not perfectly symmetric between the two players; the SB posts half a bet and acts last postflop, and these positional asymmetries combine to produce a small positive value for the button. In the strict zero-sum sense, BB's value is −0.088 BB per hand. The game is solved: if you play the Cepheus strategy, you cannot lose more than ϵ per hand, no matter who your opponent is.

For heads-up **no-limit** hold'em, the game is too big to solve exactly, but Libratus (2017) and DeepStack (2017) both used CFR variants and beat top human professionals. The game value is approximately 0 for SB (slightly positive) in no-limit as well; the exact number is not known because an exact solution is infeasible.

What this understanding buys you

Heads-up poker has a well-defined value. Whatever you believe about your edge at the table, the ceiling is the difference between your win rate and the game value. If HULHE has a value of +0.088 BB per hand for the button, no button player can ever sustainably win more than that per hand against equally strong competition. The value of the game is the upper bound on your edge at infinite skill.

Unexploitable = at the minimax value. When you hear "play GTO," what is being asked is that you play a strategy achieving the minimax value — a strategy against which no opponent can do better than the zero-sum value of the game. You give up exploitation potential against weak players, but you also eliminate every losing counter-strategy. This is the exact tradeoff the minimax theorem makes precise: the value is what you can guarantee, and no more.

Cross-Links

- **1.1 Players, strategies, payoffs** — the underlying objects.
- **1.2 Mixed strategies** — the lift required for the theorem.
- **1.4 Best response** — minimax strategies are mutual best responses.
- **2.1 Nash equilibrium** — minimax is the zero-sum case.
- **2.4 Nash existence theorem** — generalization to general-sum.
- **9.2 No-regret learning** — time-averaged dynamics converge to minimax value.
- **9.3 CFR** — the algorithm that solved HU limit poker.
- **Linear programming / convex optimization** — the computational backbone.

Structural Compression

Invariant: In every finite two-player zero-sum game, the saddle-point value exists, is unique, and equals both the maximin and minimax value; it is attained by a pair of mixed strategies that form mutual best responses.

Mechanism: Linear-programming duality (or Sion's minimax theorem) applied to the bilinear payoff function on the product of simplices; strong duality equates the primal optimum (maximin) to the dual optimum (minimax).

Cross-System Echo: Minimax is the mathematical skeleton of adversarial settings throughout computer science: GANs train at a saddle point of generator vs discriminator; robust optimization is a minimax over worst-case parameter realizations; adversarial examples in machine learning are the inner maximization problem. All of these are lifts of the zero-sum minimax theorem to richer function spaces.

Exercises

1. Verify that Matching Pennies has pure-strategy security levels $\underline{v}_{\text{pure}} = -1$ and $\overline{v}_{\text{pure}} = +1$, and mixed-strategy value $v = 0$.
2. Solve the 2×2 zero-sum game $A = \begin{pmatrix} 3 & -1 \\ -2 & 4 \end{pmatrix}$ for the value and optimal mixed strategies via the indifference principle.
3. Write out the primal LP and the dual LP for a general $m \times n$ zero-sum game and verify that they are duals of each other.
4. Prove the trivial inequality $\underline{v} \leq \overline{v}$ from scratch, without using LP duality.
5. Show that in a zero-sum game, any maximin strategy of Row and any minimax strategy of Column together form a Nash equilibrium.
6. Construct a 3×3 zero-sum game with a pure saddle point. Show that any iterated strict-dominance argument would find the saddle directly.
7. **Argue, structurally,** why the minimax theorem requires mixed strategies in general but a pure saddle suffices in games with a dominant strategy. Specifically: identify which hypothesis of LP strong duality fails in the pure-strategy formulation of Matching Pennies, and explain how the lift to mixed strategies restores it.

Normal Form as Universal Representation

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Strategic Form Games **Node:** 1.6

The normal form (N, S, u) from Node 1.1 is the minimal specification of a strategic situation. This node makes the stronger claim: *every* game, including sequential and information-rich ones introduced in Modules 3 and 4, can be flattened back to normal form — at the cost of an exponentially larger strategy space. This universality is what lets Module 2's equilibrium theory apply to the whole subject; it is also a warning, because the flattened form loses information that is visible in extensive form. The node is the formal interface between strategic form (Module 1) and extensive form (Module 3), and it motivates the equilibrium refinements of subgame perfection (3.5) and sequential equilibrium (4.6).

Formal Definition

Let Γ be an extensive-form game with player set N, a rooted game tree T, information partitions $\{\mathcal{I}_i\}_{i \in N}$, action assignments, chance moves with known probabilities, and outcome payoffs. (Module 3 develops this structure.)

A **behavioral strategy** for player i is a map $\beta_i : \mathcal{I}_i \to \Delta(A)$, assigning to each of i's information sets a distribution over the actions available at that set.

A **pure strategy** s_i in extensive form is a map $s_i : \mathcal{I}_i \to A$, assigning to each information set a single action. The set of pure strategies of player i is

$$S_i^{\text{ext}} = \prod_{I \in \mathcal{I}_i} A(I).$$

The **normal form of** Γ is the strategic form game $G(\Gamma) = (N, \{S_i^{\text{ext}}\}, \{u_i^{\text{ext}}\})$, where the payoff $u_i^{\text{ext}}(s_1, \dots, s_n)$ is the expected payoff to player i along the stochastic path induced by $(s_1, \dots, s_n)$ through the tree T, averaged over chance moves.

Universal representation claim. Every finite extensive-form game Γ has a well-defined normal form $G(\Gamma)$, and every Nash equilibrium of $G(\Gamma)$ corresponds to a strategy profile of Γ. Conversely, every Nash equilibrium of Γ (interpreted as a behavioral profile) induces a Nash equilibrium of $G(\Gamma)$ via Kuhn's theorem on the equivalence of mixed and behavioral strategies under perfect recall.

Combinatorial cost. If player i has k_I actions at information set I, then $|S_i^{\text{ext}}| = \prod_{I \in \mathcal{I}_i} k_I$, which is typically exponential in the number of information sets.

Intuition

A pure strategy in the normal form is a **complete contingent plan** — a specification of what you would do at every decision point you could reach. In a game with a single decision, the plan is just your action. In a game with many decisions and branches, the plan must cover every branch, including ones you would never actually visit under rational play.

Flattening an extensive-form game to normal form means enumerating all such plans. For chess (if it were finite and deterministic), a pure strategy would be a function specifying your move in every possible board position. The strategy space is astronomical.

But the normal form is still *uniquely* defined, and the abstract theory of Module 2 — existence, mixed equilibria, rationalizability — applies to it without modification. What's lost is the *temporal structure*: the normal form treats your entire contingent plan as a single choice, ignoring the fact that in the real game you'd choose step-by-step and update your plan as you learn new information. Module 3 will recover this structure via subgame perfection, and Module 4 via sequential equilibrium.

A key principle: **Kuhn's theorem** (1953) says that under *perfect recall* (a player never forgets their own past actions or observations), mixed strategies in the normal form are strategically equivalent to behavioral strategies in the extensive form. Both representations yield the same distribution over terminal nodes and the same expected payoffs.

Structural Role

This node closes Module 1 and opens the door to Modules 3 and 4:

- **1.1–1.5 $\Rightarrow$ 1.6.** The entire machinery of strategic-form analysis (dominance, best response, minimax, mixed equilibria) applies to any extensive-form game via its normal form.
- **1.6 $\Rightarrow$ 2.4.** Nash existence in extensive form is an immediate corollary of Nash existence in normal form plus this universality.
- **1.6 $\Rightarrow$ 3.3.** Extensive-form strategies are formally defined in Module 3; this node previews the enumeration.
- **1.6 $\Rightarrow$ 3.5.** Subgame perfection is a refinement that excludes normal-form Nash equilibria that rely on non-credible off-path behavior.
- **1.6 $\Rightarrow$ 4.6.** Sequential equilibrium refines further using consistent beliefs at out-of-equilibrium information sets.
- **Practical corollary.** "Solving" a sequential game can be done either on the normal form (LP for zero-sum, fixed point for general-sum) or directly on the extensive form (backward induction, CFR, tree search). The existence theorems are about the normal form; the efficient algorithms are about the extensive form.

Mathematical Development

From tree to table

Given extensive-form game Γ, enumerate all pure strategies for each player and compute payoffs for each strategy profile. This is the bijective construction of $G(\Gamma)$.

Example. Consider the game: player 1 moves first, choosing L or R. If L, payoff $(2, 0)$. If R, player 2 chooses l or r. If r, payoff $(0, 1)$; if l, payoff $(1, 2)$.

- Player 1 has 2 pure strategies: $\{L, R\}$.
- Player 2 has 2 pure strategies: $\{l, r\}$ — even though player 2 only moves in the R subtree, the pure strategy is defined for all information sets; here there is only one information set for player 2.
- Payoff matrix (rows = player 1, columns = player 2):

	l	r
L	$(2, 0)$	$(2, 0)$
R	$(1, 2)$	$(0, 1)$

Notice player 1's L row is constant in player 2's action: L "closes" the game before player 2 moves. This is an artifact of the flattening: in the tree, player 2 doesn't act at all after L.

Exponential blowup

For a depth-d binary tree where each player moves on alternating levels, each player has $2^{d/2}$ information sets and $2^{2^{d/2}}$ pure strategies. For $d = 20$, that is already 2^{1024} — vastly more than the number of atoms in the observable universe.

Kuhn's theorem on behavioral strategies

Theorem (Kuhn, 1953). In a finite extensive-form game with perfect recall, for every mixed strategy $\sigma_i \in \Delta(S_i^{\text{ext}})$ there exists a behavioral strategy β_i such that σ_i and β_i induce the same distribution over terminal nodes, and conversely.

The significance: you never need to reason about $\Delta(S_i^{\text{ext}})$, which is doubly exponential; you can instead work with behavioral strategies β_i, which live in the product of simplices $\prod_{I \in \mathcal{I}_i} \Delta(A(I))$, which is only linear in the number of information sets. Without Kuhn's theorem, computational game theory would be infeasible.

Caveat. Kuhn's theorem fails under **imperfect recall** — when a player forgets their own past actions or observations. In such games, mixed and behavioral strategies are not equivalent. This matters in abstraction-based poker solving, where information-set merging introduces imperfect-recall artifacts.

Payoff-equivalence of representations

Two games G and G' are **strategically equivalent** if they have the same set of Nash equilibria (identified via any bijection on strategies). The normal form of an extensive-form game is strategically equivalent to the extensive form in this sense, but it is not *refinement-equivalent*: some Nash equilibria of the normal form fail subgame perfection in the extensive form. Module 3 will develop this.

Worked Example

Subgame-perfect failure in the normal form

Consider the entry deterrence game: an Entrant (player 1) decides whether to enter a market (In) or stay out (Out). If In, the Incumbent (player 2) decides to Fight (F) or Accommodate (A).

Extensive form payoffs: Out $\to (0,2)$; In, A $\to (1,1)$; In, F $\to (-1,-1)$.

Normal form:

	F	A
In	$(-1,-1)$	$(1,1)$
Out	$(0,2)$	$(0,2)$

Nash equilibria of the normal form. Two pure ones: 1. (In, A): payoff $(1,1)$. Both best-respond. 2. (Out, F): payoff $(0,2)$. Entrant best-responds given Incumbent threatens to Fight; Incumbent best-responds given Entrant is Out (Incumbent never has to actually fight, so any action is a best response on the unreached subgame).

What's lost in the flattening. In the extensive form, equilibrium #2 relies on an *incredible* threat: if the Entrant actually entered, the Incumbent's payoff from Fight is -1 vs Accommodate's $+1$. The threat to fight is not credible — it would be suboptimal to execute. In the normal form, this incredibility is invisible because "F" on the off-path branch never lowers Incumbent's realized payoff. This is the motivating example for subgame perfection (Node 3.5).

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Extensive form Γ	The full betting tree of a hand, including decision points, action branches, and chance (dealt cards)

Formal object	Poker instantiation
Information set	A set of game histories indistinguishable to the acting player (unknown hole cards, unknown previous mixed-action resolutions)
Pure strategy in extensive form	A complete contingent plan across every hole-card/board/history combination
Normal-form enumeration	Astronomical — HUNL has ~10^{161} information sets, making full normal form infeasible
Behavioral strategy	The object a solver actually stores — frequency distributions over actions at each information set
Kuhn's theorem	Why CFR can output behavioral frequencies and still be "solving" the mixed-strategy game

Concrete situation

HUNL flop-and-beyond decision tree. After the flop, each player can check, bet (multiple sizes), call, fold, or raise (multiple sizes). A reasonable abstraction uses 3 bet sizes and 2 raise sizes, producing ~7 actions per decision. A typical flop–turn–river tree has 20–40 decision points per hand, and with 1326 hole-card combinations and 22100 boards, the number of information sets is in the hundreds of millions even after aggressive abstraction. The pure-strategy count is 7^{millions} — a number whose *description* is longer than this lesson.

Nobody enumerates the normal form. Instead, modern solvers (PioSolver, MonkerSolver, GTO+) store **behavioral strategies**: frequency distributions at each information set. Kuhn's theorem guarantees this is strategically equivalent to solving the normal form. CFR iterates on behavioral strategies directly, updating frequencies by counterfactual regret, and converges to a near-Nash equilibrium of the normal form *as measured by exploitability.*

What this understanding buys you

You never have to think about the normal form of a poker game. The strategy you are studying, the frequencies the solver outputs, the "GTO mix" you are trying to memorize — these are behavioral strategies in the extensive-form tree. Kuhn's theorem guarantees they are strategically equivalent to mixed strategies in the normal form, which is the object the existence theorems apply to. You get the abstract theory's guarantees (Nash equilibrium exists, zero-sum games have a value, etc.) without ever having to enumerate 7^{millions} contingent plans.

The warning is about abstraction. When a solver abstracts the game (e.g., by bucketing similar hands together), the abstraction may introduce imperfect recall, and Kuhn's theorem may fail. This means the solution on the abstracted tree is not guaranteed to be a Nash equilibrium of the original game. Libratus and DeepStack handled this with careful subgame re-solving and abstraction-refinement techniques to bound the error.

Cross-Links

- **1.1 Players, strategies, payoffs** — the base objects.
- **1.3 Dominance** — dominated strategies in the normal form of an extensive-form game.
- **3.1 Game trees** — the extensive form itself.
- **3.3 Strategies in extensive form** — the formal definition of s_i as a function on information sets.
- **3.5 Subgame perfect equilibrium** — the refinement that distinguishes normal-form Nash equilibria on the basis of off-path credibility.
- **4.6 Sequential equilibrium** — further refinement with consistent beliefs.

- **9.3 Counterfactual regret minimization** — the algorithm that exploits behavioral-strategy representation.
- **Statistics — conditional expectation** — how payoffs are computed over randomness.

Structural Compression

Invariant: Every finite extensive-form game has a unique normal-form representation obtained by enumerating complete contingent plans; under perfect recall, mixed and behavioral strategies yield equivalent solution concepts.

Mechanism: Contingent-plan enumeration plus payoff computation along all induced paths through the tree; Kuhn's theorem converts the doubly-exponential mixed-strategy space to the linear-in-information-sets behavioral-strategy space.

Cross-System Echo: The reduction of sequential/stateful problems to flat, stateless representations is a recurring move in computer science: POMDPs reduce to belief MDPs, Markov games reduce to matrix games at each state, policy spaces in reinforcement learning are the behavioral-strategy analogue of the normal-form flattening. The computational-versus-representational tradeoff is universal.

Exercises

1. Write out the full normal form of the entry deterrence game above. Identify both pure Nash equilibria and verify which is subgame perfect (preview of 3.5).
2. Consider a two-stage game: player 1 chooses $\{L, R\}$; player 2 observes player 1's choice and then chooses $\{l, r\}$. Payoffs: $LL \to (1,1), LR \to (0,2), RL \to (2,0), RR \to (-1,-1)$. Write the normal-form payoff matrix. Player 2 has how many pure strategies? Find all Nash equilibria.
3. For a perfect-information game of depth d with branching factor 2, count the pure strategies of the first player assuming they move at every other depth.
4. State Kuhn's theorem precisely and sketch why it requires perfect recall.
5. Give an example of a game with imperfect recall in which a mixed strategy in the normal form achieves a payoff unattainable by any behavioral strategy.
6. Verify that if a normal-form equilibrium assigns positive probability to a strategy that prescribes a strictly dominated action at an off-path information set, the equilibrium fails subgame perfection.
7. **Argue, structurally,** why normal-form equivalence is necessary but not sufficient for a satisfying solution concept in extensive-form games. Identify one concrete piece of strategic information that the normal form preserves and one it loses.