Module 2 — Nash Equilibrium

Game Theory

hbar.university

Nash Equilibrium — Definition

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.1

Nash equilibrium is the central solution concept of the course. It is the configuration of strategies from which no player has a unilateral incentive to deviate — a mutual-best-response fixed point. Once defined, it becomes the target of every subsequent chapter: we will prove it exists (2.4), explore when it is pure (2.2) or mixed (2.3), refine it by credibility (3.5), belief consistency (4.4, 4.6), cooperation (5.4), mechanism design (6.x), evolutionary stability (7.2), and learning (9.x). The entire course is, in a sense, a study of the ways Nash equilibrium can be characterized, computed, selected, refined, and approximated.

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite strategic-form game with mixed extension $\Delta(S_i)$.

A **Nash equilibrium** is a strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*) \in \prod_{i \in N} \Delta(S_i)$ such that for every player i and every alternative strategy $\sigma_i \in \Delta(S_i)$,

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i, \sigma_{-i}^*).$$

Equivalently, $\sigma_i^* \in \mathrm{BR}_i(\sigma_{-i}^*)$ for every i, i.e. σ^* is a fixed point of the joint best-response correspondence BR.

Pure Nash equilibrium. The special case in which every σ_i^* is a pure strategy. Pure Nash equilibria may or may not exist; mixed Nash equilibria always do in finite games (2.4).

Strict Nash equilibrium. A profile $s^* \in S$ such that $u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*)$ for every i and every $s_i \neq s_i^*$. Strict equilibria are necessarily pure, and they are robust to small payoff perturbations.

The Nash condition has two equivalent informal phrasings:

1. **No-profitable-deviation:** no single player can improve their expected payoff by unilaterally changing their strategy.
2. **Self-enforcing belief:** if every player expects every other to play σ_{-i}^*, their own best response is σ_i^*, so the expectations are self-consistent.

Intuition

Nash equilibrium captures the simplest possible stability criterion: *everyone's strategy is a best response to everyone else's*. It is not a claim that the equilibrium is good, fair, Pareto-efficient, or even unique. It is only a statement that it is a configuration nobody wants to deviate from, given what everybody else is doing.

Several common misunderstandings are worth clearing up:

Nash equilibrium is not cooperative. The Prisoner's Dilemma has a Nash equilibrium at (D, D) even though both players would prefer (C, C). Nash equilibrium does not respect Pareto improvements; it only respects unilateral deviations.

Nash equilibrium is not unique. Most games have multiple equilibria, sometimes infinitely many. Equilibrium selection is a separate problem.

Nash equilibrium does not require explicit communication. Players do not need to coordinate to play an equilibrium; they only need to have beliefs about each other that are *consistent* with the equilibrium, and the equilibrium is a configuration where such beliefs are self-fulfilling.

Nash equilibrium can be mixed. In a mixed equilibrium, each player is indifferent between the pure strategies in their support, and the mix is chosen so that the *opponent* is indifferent at the margin.

The reason Nash equilibrium is central is that almost every other solution concept — from subgame perfection to correlated equilibrium to rationalizability to evolutionarily stable strategies — is either a *refinement* (narrower) or a *coarsening* (broader) of Nash equilibrium.

Structural Role

Every subsequent node refers back to Nash equilibrium:

- **2.2 Pure Nash.** When is the fixed point a pure profile?
- **2.3 Mixed Nash.** How do we compute mixed equilibria via the indifference principle?
- **2.4 Existence.** Every finite game has at least one Nash equilibrium (possibly mixed).
- **2.5 Rationalizability.** A broader concept — every Nash profile is rationalizable, but not conversely.
- **2.6 Correlated equilibrium.** A broader concept — every Nash profile is a correlated equilibrium, but not conversely.
- **3.5 Subgame perfection.** A refinement for extensive-form games.
- **4.3 Bayesian Nash.** Extension to incomplete information.
- **4.4 Perfect Bayesian equilibrium.** Refinement with consistent beliefs.
- **5.4 Folk theorem.** In repeated games, the set of Nash-supportable payoffs is much larger than in the one-shot game.
- **7.2 ESS.** A refinement motivated by evolutionary dynamics.
- **9.1–9.3 Learning.** Convergence of dynamic play to Nash equilibrium.

Mathematical Development

Fixed-point characterization

A profile σ^* is a Nash equilibrium iff

$$\sigma^* \in \mathrm{BR}(\sigma^*) \quad \text{where} \quad \mathrm{BR}(\sigma) = \mathrm{BR}_1(\sigma_{-1}) \times \cdots \times \mathrm{BR}_n(\sigma_{-n}).$$

This is a fixed-point equation on the compact convex set $\prod_i \Delta(S_i)$. Existence follows from Kakutani's theorem (proved in 2.4).

Support characterization of mixed equilibria

If σ^* is a Nash equilibrium and $s_i, s_i' \in \mathrm{supp}(\sigma_i^*)$, then $U_i(s_i, \sigma_{-i}^*) = U_i(s_i', \sigma_{-i}^*)$ — the **indifference principle**. Any pure strategy not in the support yields payoff no higher than the pure strategies in the support.

This gives a linear-algebraic way to compute mixed equilibria: fix a candidate support, write the indifference equations, solve for the mixing probabilities, and verify (a) the probabilities are in $[0, 1]$, (b) the out-of-support strategies do not do better.

Two-player bimatrix games

For a bimatrix game (A, B) with row player 1 and column player 2: - A pure profile (i, j) is a Nash equilibrium iff $A_{ij} \geq A_{kj}$ for all k (Row best-responds) and $B_{ij} \geq B_{il}$ for all l (Column best-responds). - A mixed profile (p, q) is a Nash equilibrium iff p is a best response to q (support supported by Row's indifference) and q is a best response to p (support supported by Column's indifference).

The n-player case

For $n > 2$, the equations are multilinear rather than linear, and the computation is PPAD-complete in general. Even verifying that a given profile is a Nash equilibrium is easy (just check best-response conditions), but finding one is computationally hard.

Relation to minimax

In a two-player zero-sum game, Nash equilibrium coincides with the minimax saddle point. The set of Nash equilibria is convex (it is the product of the row player's optimal strategies and the column player's optimal strategies). In general-sum games, the set of Nash equilibria need not be convex.

Worked Example

Stag Hunt revisited

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

Pure equilibria. $(\text{Stag}, \text{Stag})$: Row gets 4 vs 3 if deviates, Column gets 4 vs 3 — both best-respond. Strict equilibrium. $(\text{Hare}, \text{Hare})$: Row gets 3 vs 0 if deviates, Column gets 3 vs 0 — both best-respond. Strict equilibrium.

Mixed equilibrium. Let Row play $(p, 1-p)$, Column play $(q, 1-q)$. Row's expected payoff from Stag is $4q$; from Hare is $3q + 3(1-q) = 3$. Row is indifferent iff $4q = 3 \Rightarrow q = 3/4$. Symmetrically $p = 3/4$. Mixed equilibrium: both play Stag with probability 3/4, Hare with probability 1/4. Expected payoff 3 each — worse than the Stag equilibrium and the same as the Hare equilibrium, but the mixed profile is still a Nash equilibrium because nobody can profitably deviate.

Three equilibria in total. Two pure (one Pareto-efficient, one risk-dominant), one mixed. This is a game-selection problem: Nash equilibrium by itself does not tell you which of the three to expect. Later nodes (risk dominance, global games, evolutionary selection) will.

Battle of the Sexes

	Opera	Football
Opera	$(2, 1)$	$(0, 0)$
Football	$(0, 0)$	$(1, 2)$

Pure equilibria: $(\text{Opera}, \text{Opera})$ and $(\text{Football}, \text{Football})$. Mixed equilibrium: Row plays Opera with probability 2/3, Football with probability 1/3; Column plays Opera with probability 1/3, Football with probability 2/3. Check: Row's indifference between Opera and Football given Column's mix: $2 \cdot 1/3 + 0 \cdot 2/3 = 2/3$; $0 \cdot 1/3 + 1 \cdot 2/3 = 2/3$. Correct.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Strategy profile σ^*	A complete strategy assignment for all seats — every range, every action frequency, every bet size
Nash equilibrium	**GTO** — Game-Theoretic Optimal play
Mutual best response	Every player's strategy is already maximally exploitative against every other player's strategy
Indifference principle	At every mixed action, all supported options have the same EV
Set of Nash equilibria	The "GTO space" — in zero-sum HU poker, convex; in 3+ player games, can be non-convex
Mixed equilibrium	Solver's frequency-based recommendations (e.g. bet 62%, check 38%)

The central identification: GTO is Nash equilibrium of poker. When a poker coach or a solver outputs "the GTO strategy," what is being computed is literally a Nash equilibrium of the game. Every frequency, every bet size, every action recommendation is one coordinate of a fixed point of the joint best-response map on the space of strategy profiles. This is not a metaphor — it is an equality of formal objects.

When someone says "this is a GTO decision," they mean "this is the best-response component of a Nash equilibrium profile." When someone says "I'm playing GTO," they mean "I'm playing my component σ_i^* of a computed Nash equilibrium profile σ^*." When someone says "GTO doesn't care about your opponent," they mean "the Nash equilibrium is defined relative to an assumed equilibrium opponent, not the actual opponent you are facing" — which is correct, and also the precise statement of what makes GTO *unexploitable* rather than *optimal*.

Concrete situation

6-max cash, HJ open 2.5x, CO 3-bet to 8 BB, folds to HJ. HJ is holding *AQo* with 100 BB effective. Should HJ fold, call, or 4-bet?

GTO analysis. In the Nash equilibrium of 6-max no-limit poker, HJ's strategy at this information set is a mix: - Fold *AQo* ~40% of the time - Call ~30% - 4-bet ~30%

The mix exists because at the current state of the game tree, the three actions yield *nearly equal* EV for HJ against CO's equilibrium 3-betting range. The indifference principle says this must be so in any mixed Nash equilibrium. The 40/30/30 split is not "HJ's preference" — it is HJ's equilibrium mix *given CO's equilibrium 3-bet frequency.* Change CO's range (e.g., if CO is tighter), and HJ's mix shifts.

Why it's a fixed point. If HJ deviates to, say, 100% 4-bet with *AQo*, then CO's equilibrium 5-bet and call ranges are no longer best responses — CO can counter-adjust, which in turn makes HJ's adjustment worse. Only at the mutual-best-response fixed point does the strategy profile stabilize.

What this understanding buys you

GTO is not magic. It is a solution to a fixed-point equation. Every frequency the solver outputs can be traced back to "this is the value at which my opponent is indifferent, so they can't exploit me by deviating, and at which I am indifferent, so they can't exploit me by shifting either." Once you see it this way, you stop memorizing charts and start doing indifference computations at the table. This is the single largest conceptual shift between a "chart player" and a "GTO player" — the chart player has the answer, the GTO player has the equation that generates the answer.

Implication for study. When you are reviewing a hand, the right question is never "what is the GTO play?" but "what are the indifference conditions that make the mix the solver is showing me a fixed point?" Answering this question teaches you the game; answering the first question teaches you a lookup table.

Cross-Links

- **1.4 Best-response correspondences** — the operator whose fixed points are Nash equilibria.
- **1.5 Zero-sum minimax** — the special case where Nash coincides with saddle points.
- **2.2, 2.3 Pure and mixed Nash equilibria** — structural refinements.
- **2.4 Nash existence theorem** — guarantees fixed points exist.
- **2.5 Rationalizability** — broader concept based on best-response iterations.
- **2.6 Correlated equilibrium** — broader concept with shared signals.
- **3.5 Subgame perfect equilibrium** — refinement of Nash for extensive-form games.
- **7.2 Evolutionarily stable strategy** — evolutionary refinement.
- **9 Learning** — dynamic processes that converge to Nash.

Structural Compression

Invariant: A Nash equilibrium is a strategy profile σ^* such that $\sigma^* \in \mathrm{BR}(\sigma^*)$: each player's strategy is simultaneously a best response to the others'.

Mechanism: Fixed-point condition on the joint best-response correspondence defined over the compact convex set $\prod_i \Delta(S_i)$.

Cross-System Echo: Fixed-point characterizations of equilibrium appear throughout science: chemical equilibria (reaction rates balance), physical equilibria (forces balance), economic equilibria (supply equals demand at the price-taker's best response), dynamical attractors (state is fixed under the dynamics). Nash equilibrium is the strategic analogue: beliefs are fixed under the best-response map.

Exercises

1. Verify directly that (D, D) is the unique pure Nash equilibrium of the Prisoner's Dilemma, and that no mixed equilibrium exists beyond this pure profile.

2. Show that the mixed equilibrium of Stag Hunt computed above is indeed a Nash equilibrium by verifying the indifference condition and confirming no profitable deviations.

3. In Battle of the Sexes, compute the expected payoff of each player in the mixed equilibrium. Compare to the pure equilibria. Is the mixed equilibrium Pareto-dominated?

4. Prove: every strict Nash equilibrium is pure.

5. Construct a 2×2 game in which the set of Nash equilibria has a continuum of elements. (Hint: use a game with identical payoffs in one row or column.)

6. Prove: the set of Nash equilibria of a zero-sum game is convex in the profile space. Give a counterexample for general-sum games.

7. **Argue, structurally,** why "best response to itself" is the minimal possible stability criterion for rational play in a strategic game, and why Nash equilibrium is therefore both universal (every rational outcome is Nash) and weak (Nash-ness alone does not pick a single outcome or guarantee desirable properties like Pareto efficiency).

Pure-Strategy Nash Equilibrium

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.2

This node specializes Nash equilibrium (2.1) to the case where each player plays a deterministic strategy. Pure Nash equilibria are the easiest to find, the easiest to interpret, and the ones with the strongest stability properties, but they are not guaranteed to exist in every game. Recognizing which games have pure equilibria — and which classes of games are *guaranteed* to have them (supermodular games, potential games, finite games with strategic complementarities) — is one of the main applied-theory skills in game theory. This node also establishes pure-strategy analysis via best-response matrix inspection, the "find all pairs of marked cells" technique that is used throughout the course.

Formal Definition

A **pure-strategy Nash equilibrium** of a game $G = (N, \{S_i\}, \{u_i\})$ is a profile $s^* = (s_1^*, \ldots, s_n^*) \in S = \prod_i S_i$ such that for every player i and every $s_i \in S_i$,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*).$$

Equivalently, $s_i^* \in \mathrm{BR}_i(s_{-i}^*)$ for every i, where best responses are taken over pure strategies only.

Strict pure Nash. $u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*)$ for all $s_i \neq s_i^*$. Every strict Nash equilibrium is pure (you cannot strictly improve on a mixture over ties).

Pure Nash need not exist. Matching Pennies has no pure Nash equilibrium. Rock–Paper–Scissors has no pure Nash. Any game in which the best-response correspondence cycles without a fixed point in pure strategies lacks one.

Games that guarantee pure Nash. Several classes: - **Finite dominance-solvable games** (the surviving profile is a pure Nash). - **Finite two-player games with an ordinal potential function** (Monderer–Shapley). - **Supermodular games on lattices** (Topkis, Milgrom–Roberts): always have smallest and largest pure equilibria. - **Games with strategic complementarities** (continuous-action version of the above). - **Weakly acyclic games**: every better-response path ends at a pure equilibrium.

Intuition

A pure Nash equilibrium is a "frozen" configuration: each player is playing a single strategy, and knows the others are too, and would not want to move. It is the cleanest form of stability.

Two practical diagnostic techniques for finding pure Nash equilibria by hand:

Double-underline method. In each column, underline Row's best payoff. In each row, underline Column's best payoff. A cell with *both* numbers underlined is a pure Nash equilibrium. This works immediately for 2×2 up to 5×5 bimatrix games.

Best-response diagram. Draw Row's best-response set as a function of Column's strategy, and vice versa. Intersections are Nash equilibria. This works for continuous-action games.

The existence of pure equilibria in specific game classes is the content of several classical theorems: Topkis (1978) for supermodular games, Monderer–Shapley (1996) for potential games, Debreu–Fan–Glicksberg for concave games. Each exploits a structural feature of the payoff functions to reduce pure-Nash existence to a fixed-point result.

Structural Role

Pure Nash equilibria are the conceptually clean subset of Nash equilibria:

- **2.1 ⇒ 2.2.** Pure Nash are the special case of Nash where all strategies are deterministic.
- **2.2 ⇒ 2.3.** If a game has no pure Nash, you must look for mixed ones.
- **2.2 ⇒ 3.5.** Subgame perfection often picks out a specific pure Nash in extensive-form games with sequential structure.
- **2.2 ⇒ 5.4.** Nash reversion strategies in repeated games can be "stay at pure Nash forever" — pure is easiest to coordinate on.
- **2.2 ⇒ 7.2.** An evolutionarily stable strategy (ESS) must be a strict Nash, which is necessarily pure.

Pure Nash is also the computationally easiest class: checking if a given profile is pure Nash takes time linear in the game description. Finding one is NP-hard in general, but polynomial for supermodular, potential, and zero-sum games.

Mathematical Development

Potential games

A game G is a **potential game** (Monderer–Shapley, 1996) if there exists a function $\Phi : S \to \mathbb{R}$ such that for every player i, every s_{-i}, and every s_i, s_i',

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) = \Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}).$$

That is, any player's unilateral change in payoff equals the change in Φ. Every maximizer of Φ over S is a pure Nash equilibrium. Since Φ is a real function on a finite set, it has a maximum, so every finite potential game has a pure Nash.

Examples. Congestion games (Rosenthal, 1973). Coordination games with symmetric additive payoffs. Quadratic network games.

Supermodular games

A game is **supermodular** if the strategy sets are lattices and the payoff functions have **increasing differences** in own strategy and opponent profile: for $s_i > s_i'$ and $s_{-i} > s_{-i}'$,

$$u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i}') - u_i(s_i', s_{-i}').$$

This says: the marginal benefit of raising my strategy is higher when my opponents' strategies are higher. Equivalently, best responses are weakly increasing in opponents' strategies (strategic complementarities).

Topkis's theorem. In a supermodular game, the set of pure Nash equilibria is non-empty and has a smallest and largest element with respect to the lattice order.

Best-response dynamics and cycles

Consider Matching Pennies. Row's best response to Column playing Heads is Heads. Column's best response to Row playing Heads is Tails. Row's best response to Column playing Tails is Tails. Column's best response to Row playing Tails is Heads. The best-response dynamics cycles: $HH \to HT \to TT \to TH \to HH$. No pure fixed point exists.

Weakly acyclic games are those where from every profile, there exists a finite sequence of better-response moves ending at a pure Nash. Matching Pennies is not weakly acyclic.

Counting pure equilibria

For a random game with iid uniform payoffs on $[0,1]$, the expected number of pure Nash equilibria in an $m \times m$ bimatrix game is approximately 1 for large m. This is an interesting fact: "typical" finite games have roughly one pure equilibrium on average, but the variance is large — some have several, many have none.

Worked Example

A 3×3 bimatrix with multiple pure equilibria

$$\begin{array}{c|ccc} & L & C & R \\ \hline T & (3,2) & (1,3) & (2,1) \\ M & (4,0) & (2,2) & (3,1) \\ B & (1,4) & (0,1) & (5,3) \end{array}$$

Underline Row's best in each column. Column L: $\underline{4}$ (M). Column C: $\underline{2}$ (M). Column R: $\underline{5}$ (B).

Underline Column's best in each row. Row T: $\underline{3}$ (C). Row M: $\underline{2}$ (C). Row B: $\underline{4}$ (L).

Double-underlined cells. Cell (M,C): Row best in column C (4 underlined… wait, let me recheck. Row payoffs in column C: $1,2,0$, best is M with 2 — correct. Column payoffs in row M: $0,2,1$, best is C with 2 — correct.) So (M,C) is a pure Nash equilibrium with payoff $(2,2)$.

Check other double-underlined cells: (B,R)? Row best in R is B with 5 ✓. Column best in B is L with 4; C gives 1, R gives 3. Column's best in row B is L, not R. So R is not best for Column in row B. (B,R) is not a Nash equilibrium.

Unique pure Nash: (M,C) with payoff $(2,2)$.

A game with no pure Nash

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

This is Matching Pennies (relabeled). Best responses cycle. No pure Nash. Module 2.3 will find the mixed equilibrium $((1/2,1/2),(1/2,1/2))$.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Pure strategy profile	A 100/0 assignment for every hand in every spot — no mixing
Pure Nash equilibrium	A game configuration where every player's pure strategy is a best response to every other's
Non-existence of pure Nash	Most interesting poker spots — bluff frequencies, bluff-catcher frequencies, polarized river bets — have no pure equilibrium
Pure Nash in simple subgames	Face-up river spots where betting is unambiguously best — no mixing required
Strict pure Nash	A "locked" decision: say, folding 72o to a 4-bet — nothing ambiguous, no indifference

Concrete situation

A pocket pair on the button in a heads-up spot where villain is known to be going all-in with their entire range. You hold QQ with 100 BB effective. Villain shoves. Do you call?

Pure-strategy equilibrium subgame. At this exact information set, calling and folding are the only actions. Calling with QQ against a shove-with-any-two-cards range has ~85% equity, a clear +EV call. Folding is strictly worse. The best response is pure: Call. Villain's best response, given the all-in shoving strategy is fixed, is also pure (by construction). So in this specific subgame the pure strategy profile "(Hero calls with QQ, Villain jams any two)" is a pure Nash equilibrium of the restricted game.

In a more realistic equilibrium, Villain would not shove any two cards, and the game would have a mixed component — but the present conditional subgame is pure.

Contrast: bluff-catching. On a river with a polarized betting range from opponent, your J-high bluff-catcher typically has *mixed* best response: call with some frequency, fold with the rest. There is no pure Nash at that information set because pure "always call" is exploitable by more bluffs, and pure "always fold" is exploitable by more value. Indifference forces mixing.

What this understanding buys you

When you are in a pure-Nash spot, don't mix. If the subgame has a dominant strategy (say, calling AA vs a shove), no mixing is required; the equilibrium is deterministic. Mixing in a pure-Nash spot is just random noise.

When you are in a mixed-Nash spot, mix. If the EV of two actions is within 0.1 BB/100 of each other in equilibrium, the spot is a mixing spot, and you need an external randomizer to hit the right frequency.

The practical diagnostic. Before agonizing over a decision, ask "is this a pure-Nash spot or a mixed-Nash spot?" If the answer is pure, stop thinking and take the action. If the answer is mixed, you should have *already* decided the frequency and just execute. No pure-Nash spot should take more than a second; no mixed-Nash spot should take more than a second *either*, because the EV difference is negligible.

Cross-Links

- **1.4 Best-response correspondences** — the underlying operator.
- **2.1 Nash equilibrium definition** — the parent concept.
- **2.3 Mixed Nash** — the case that arises when no pure Nash exists.
- **2.4 Existence** — pure Nash not guaranteed; mixed Nash is.
- **3.5 Subgame perfection** — often selects a pure subgame-perfect Nash.
- **Economics — potential games, supermodular games** — classes guaranteeing pure Nash.
- **Network economics — congestion games** — a classical potential-game family.

Structural Compression

Invariant: A pure Nash equilibrium is a profile s^* of deterministic strategies such that every player's action is a best response to the others' actions.

Mechanism: Pointwise verification of the best-response condition on every player. For structured classes (potential, supermodular), existence follows from maximizing a single real-valued function or from lattice-theoretic fixed-point theorems.

Cross-System Echo: Pure equilibria correspond to "deterministic attractors" in dynamical systems — configurations where no variable has a reason to change. Contrast mixed equilibria, which correspond to "stable mixing distributions" in ergodic dynamical systems.

Exercises

1. Use the double-underline method to find all pure Nash equilibria of the 3×3 game

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

2. Show that the Prisoner's Dilemma is a potential game by exhibiting a potential function.
3. Prove that every strict Nash equilibrium is pure.
4. Show that a congestion game with cost functions $c_e(x_e)$ on edges has the Rosenthal potential $\Phi(s) = \sum_e \sum_{k=1}^{x_e(s)} c_e(k)$ and hence has a pure Nash equilibrium.
5. For the Stag Hunt, compute both pure Nash equilibria and identify which is Pareto-efficient and which is risk-dominant.
6. Construct a 2×2 game with no pure Nash but with continuous payoffs, demonstrating that the lattice hypothesis of Topkis's theorem is necessary for pure existence in supermodular games.
7. **Argue, structurally,** why the existence of a pure Nash equilibrium is a *fragile* property — specifically, why small perturbations of payoffs can destroy pure equilibria but not mixed ones. What does this suggest about the relative robustness of pure-strategy and mixed-strategy analysis?

Mixed-Strategy Nash Equilibrium

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.3

Mixed Nash is where the deep machinery of game theory lives. Most games of strategic interest have no pure Nash equilibrium, and those that do usually have mixed equilibria as well. Mixed Nash is the object that Nash existence (2.4) provides unconditionally for finite games, and it is the object that modern solvers (CFR, fictitious play, Lemke–Howson) actually compute. This node develops the support characterization, the indifference principle, and the practical technique of guessing supports, writing indifference equations, and solving them — the computational rhythm of mixed-equilibrium analysis.

Formal Definition

A **mixed-strategy Nash equilibrium** of a finite game $G = (N, \{S_i\}, \{u_i\})$ is a profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \prod_{i \in N} \Delta(S_i)$ such that for every player i and every $\sigma_i \in \Delta(S_i)$,

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i, \sigma_{-i}^*).$$

Support characterization. σ^* is a mixed Nash equilibrium iff for every i, every pure strategy in $\operatorname{supp}(\sigma_i^*)$ yields the same expected payoff against σ_{-i}^*, and this common payoff is at least the payoff of any pure strategy outside the support:

$$U_i(s_i, \sigma_{-i}^*) = v_i^*,\ s_i \in \operatorname{supp}(\sigma_i^*); \qquad U_i(s_i, \sigma_{-i}^*) \leq v_i^*,\ s_i \notin \operatorname{supp}(\sigma_i^*).$$

Here $v_i^* = U_i(\sigma_i^*, \sigma_{-i}^*)$ is the equilibrium payoff of player i.

Completely mixed Nash. An equilibrium in which every $s_i \in S_i$ is in the support of σ_i^*. This requires $|S_i| - 1$ equality constraints per player.

Indifference principle. If any mixed strategy is a best response and its support has more than one element, those elements must all yield the same expected payoff.

Intuition

A mixed equilibrium has a subtle self-referential quality. The player who is mixing is *indifferent* between the pure strategies in their mix — they do not personally prefer one over the other. So why do they mix? Not to optimize their own payoff (any pure strategy in the support is equally good), but to keep the *opponent's* best response correct. If the mixing player deviated to a pure strategy, the opponent would have a new best response, and the whole equilibrium would unravel.

The mathematical content of "mixing to keep the opponent honest" is exactly the support characterization: *my* mix solves an indifference equation for *the opponent*, and vice versa.

This makes mixed equilibria counterintuitive. You are not mixing because you want to; you are mixing because not mixing would reveal information that the opponent could use to improve their own play.

The practical technique for finding mixed Nash:

1. Guess a support for each player.
2. Write indifference equations: each pure in the support yields the same expected payoff against the opponent's mix.
3. Solve the linear system for the mixing probabilities.

4. Check that the resulting probabilities are in $[0,1]$ and sum to 1.
5. Check that out-of-support pure strategies do not exceed the equilibrium payoff.
6. If any check fails, try a different support.

This loop is the mixed-equilibrium solver for small games by hand. For large games it is infeasible (exponential in number of supports), and algorithms like Lemke–Howson or CFR are used instead.

Structural Role

Mixed Nash generalizes pure Nash and is preserved under existence theorems:

- **2.2 ⇒ 2.3.** Pure is a special case of mixed.
- **2.3 ⇒ 2.4.** Nash existence applies to the mixed-strategy simplex.
- **2.3 ⇒ 2.5.** Mixed Nash implies rationalizability (support of any equilibrium mix survives iterated best response).
- **2.3 ⇒ 2.6.** Mixed Nash is a special case of correlated equilibrium (with the correlating device being trivial).
- **2.3 ⇒ 9.1.** Fictitious play converges to mixed Nash in zero-sum and some potential games.
- **2.3 ⇒ 9.3.** CFR converges to mixed Nash of the extensive-form game.

Mixed Nash is also the refinement-resistant solution concept: many equilibrium refinements in Module 3 and 4 preserve mixed equilibria while eliminating some pure equilibria that rely on incredible threats.

Mathematical Development

Support enumeration

For a 2×2 bimatrix game (A, B):

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \quad B = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}.$$

Assume a completely mixed equilibrium $(p, 1-p)$ for Row, $(q, 1-q)$ for Column.

Row's indifference (makes Column willing to mix): Column's payoff from L is $b_{11}p + b_{21}(1-p)$; from R is $b_{12}p + b_{22}(1-p)$. Setting equal:

$$p = \frac{b_{22} - b_{21}}{b_{11} - b_{12} - b_{21} + b_{22}}.$$

Column's indifference (makes Row willing to mix):

$$q = \frac{a_{22} - a_{12}}{a_{11} - a_{21} - a_{12} + a_{22}}.$$

These formulas hold whenever the denominators are non-zero and the resulting p, q lie in $(0, 1)$.

Support guessing for larger games

For $m \times n$ games, one enumerates possible supports. The total number of supports is $(2^m - 1)(2^n - 1)$, exponential in m, n. Most supports do not yield equilibria. Heuristics like Lemke–Howson pivoting systematically explore the support graph.

Oddness theorem

Theorem (Wilson, 1971). Almost every finite game has an odd number of Nash equilibria. "Almost every" means that the set of games with even or infinite equilibria has measure zero in the space of payoff matrices.

This is a topological counting result; it says that Nash equilibria come in odd numbers generically because they are intersections of best-response curves with generic transversality.

Continuity of equilibrium correspondence

The Nash equilibrium correspondence $G \mapsto \{\sigma^* : \sigma^* \text{ is Nash of } G\}$ is upper-hemicontinuous in G on the space of bimatrices: small changes in payoffs produce small changes in the equilibrium set. This is the basis of the robustness and purification arguments of Harsanyi.

Computing mixed Nash for specific game classes

- **Zero-sum games.** LP; polynomial time.
- **Potential games.** Maximize the potential; polynomial in the number of profiles.
- **Two-player general-sum.** Lemke–Howson; exponential worst case, polynomial in practice.
- **n-player general-sum.** PPAD-complete; no known polynomial algorithm.

Worked Example

Battle of the Sexes (mixed)

	O	F
O	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

Pure Nash: (O,O) and (F,F).

Mixed Nash: Guess support $= \{O,F\}$ for both. Row plays $(p, 1-p)$, Column plays $(q, 1-q)$.

Column's indifference (Row's mix makes Column's O vs F equal): $1 \cdot p + 0 \cdot (1-p) = 0 \cdot p + 2 \cdot (1-p) \Rightarrow p = 2(1-p) \Rightarrow p = 2/3$.

Row's indifference: $2q + 0 = 0 + 1(1-q) \Rightarrow q = 1/3$.

Equilibrium: Row plays $(2/3, 1/3)$, Column plays $(1/3, 2/3)$.

Equilibrium payoff. Row: $2 \cdot 1/3 + 0 \cdot 2/3 = 2/3$. Column: $1 \cdot 1/3 + 0 = 1/3 \ldots$ wait, let me redo. Row's expected payoff: $pq \cdot 2 + p(1-q) \cdot 0 + (1-p)q \cdot 0 + (1-p)(1-q) \cdot 1 = 2 \cdot (2/3)(1/3) + (1/3)(2/3) = 4/9 + 2/9 = 6/9 = 2/3$. Column's: $1 \cdot (2/3)(1/3) + 2 \cdot (1/3)(2/3) = 2/9 + 4/9 = 6/9 = 2/3$. Both get $2/3$, same payoff, much worse than either pure equilibrium. Mixed Nash is Pareto-dominated.

A 3-strategy game

	L	C	R
T	$(0,0)$	$(2,3)$	$(3,2)$
M	$(3,2)$	$(0,0)$	$(2,3)$
B	$(2,3)$	$(3,2)$	$(0,0)$

This is a 3-strategy "cyclic" game. No pure Nash (check best responses cycle). Completely mixed symmetric equilibrium: guess all three actions in support.

Row's indifference: $U_R(T, \sigma_C) = U_R(M, \sigma_C) = U_R(B, \sigma_C)$. Let $\sigma_C = (q_L, q_C, q_R)$. Row T: $0 \cdot q_L + 2q_C + 3q_R$. Row M: $3q_L + 0 + 2q_R$. Row B: $2q_L + 3q_C + 0$. Setting equal with $q_L + q_C + q_R = 1$ gives a linear system. By symmetry, $q_L = q_C = q_R = 1/3$, and each row's expected payoff is $(0+2+3)/3 = 5/3$ for T, which should match $(3+0+2)/3 = 5/3$ for M, and $(2+3+0)/3 = 5/3$ for B. All equal. Confirmed.

Symmetric completely mixed equilibrium: $\sigma^* = ((1/3, 1/3, 1/3), (1/3, 1/3, 1/3))$, each with expected payoff $5/3$.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Mixed equilibrium σ^*	Frequency-based GTO solutions — "bet 62%, check 38%" on every mixed information set
Support of σ_i^*	The set of actions the solver ever takes at a given hand in a given spot
Indifference condition	All actions in a mix must have the same EV against opponent's equilibrium response
Solving for mixed probabilities	Exactly what a solver does iteratively: update frequencies to enforce opponent indifference
Out-of-support pure strategies	Actions the solver never recommends — dominated or dominated-in-context
Guessing supports	What a human analyst does when trying to "reconstruct" a solver's output by hand

Concrete situation

River decision, single-raised heads-up pot. Board: $A♠K♣7\diamondsuit3\heartsuit2♣$. BTN has bet flop and turn, BB has called both. On the river, BTN faces a 2/3-pot bet from BB.

What is the solver doing? With each hand in BTN's range, the solver is computing the call/fold mix that makes BB indifferent between value-betting with their best hands and bluffing with their worst. The indifference condition for BB is:

$$\alpha \cdot (\text{EV of value bet}) + (1-\alpha) \cdot (\text{EV of bluff}) = \text{EV of checking},$$

where α is the fraction of BB's river betting range that is value. The mixed strategy BTN plays with bluff-catchers must make this equation hold; otherwise BB could profitably shift toward more bluffs or more value.

Concrete mix. Suppose BTN has a bluff-catcher like $K\heartsuit T♣$. The solver outputs: call 65%, fold 35%. These are *not* BTN's preferences — BTN is indifferent between the two (both yield the same EV against BB's equilibrium betting frequency). The 65/35 is the unique mix at which BB's best response is neither "bet more bluffs" nor "bet fewer bluffs" — it is exactly the frequency that makes BB's own indifference hold.

This is the support characterization in its most operational form: the support is $\{\text{call}, \text{fold}\}$, the indifference condition on *BTN's* side is $\text{EV(call)} = \text{EV(fold)}$, and the mix is chosen to balance *BB's* incentives.

What this understanding buys you

Mixing is for the opponent, not for you. You mix at a river spot because not mixing would let the opponent shift their own range to exploit you. You do not mix because you are "uncertain" — you mix because the fixed-point condition demands it.

Any deviation from the mix is exploitable — but only if the opponent notices. Against a strong opponent, stick to the mix. Against a weak opponent who will not adjust, deviate toward the exploit. The mixed equilibrium is the unexploitable floor; the best response is the exploit.

When solver output confuses you, run the indifference check yourself. Write down the EV of each action in the mix against the opponent's equilibrium response. They should be equal (up to numerical precision of the solver). If they are, the mix is right. If they are not, the solver has not converged yet.

Cross-Links

- **1.2 Mixed strategies** — the lift that makes mixed equilibria well-defined.
- **1.4 Best response** — mixed best response = support indifference.
- **2.1 Nash equilibrium** — the parent definition.
- **2.2 Pure Nash** — the special case with singleton supports.
- **2.4 Nash existence** — guarantees at least one mixed Nash.
- **2.5 Rationalizability** — broader solution concept.
- **9.1 Fictitious play** — dynamic convergence to mixed Nash.
- **9.3 CFR** — iterative construction of mixed equilibria in extensive form.

Structural Compression

Invariant: In a mixed Nash equilibrium, each player's mixing distribution is supported on pure strategies that are all optimal against the others' equilibrium profile, and these supports satisfy an indifference principle that makes every opponent's equilibrium response a best response in turn.

Mechanism: Linear indifference equations over mixing probabilities, supplemented by inequality constraints verifying out-of-support optimality; existence is guaranteed by convexity of the mixed-strategy simplex and the fixed-point theorem (2.4).

Cross-System Echo: Mixed strategies at equilibrium are the strategic analogue of *stationary distributions* in ergodic Markov chains: the fixed distribution that the system occupies in the long run. The indifference condition is the analogue of detailed balance.

Exercises

1. Compute the mixed Nash of Matching Pennies and verify directly that both players are indifferent at the mix.

2. For the game

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

   find the pure and mixed equilibria.

3. Apply the indifference principle to compute the mixed equilibrium of a generic 2×2 zero-sum game $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $a + d \neq b + c$. Derive the value of the game.

4. Prove: in any completely mixed Nash equilibrium of a finite game, the number of equality constraints per player equals $|S_i| - 1$, giving a system that may or may not have solutions with non-negative probabilities.

5. State the oddness theorem and show by example that a 2×2 game with infinitely many Nash equilibria is non-generic.
6. Given a bimatrix game with a mixed Nash (p^*, q^*), show that the equilibrium payoff to Row equals $p^{*\top} A q^*$, and that this equals $A_{i,j}$ for any $i \in \mathrm{supp}(p^*)$ and $j \in \mathrm{supp}(q^*)$ when payoff entries coincide.
7. **Argue, structurally,** why the indifference principle — the fact that a mixing player is indifferent between supported pure strategies — is not a contradiction with rationality, even though the player has "no reason" to pick one pure strategy over another. Frame your answer in terms of the self-referential fixed-point nature of Nash equilibrium.

Nash's Existence Theorem

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.4

Nash's existence theorem (1950) is the foundational result of the entire field. It says: every finite strategic-form game has at least one Nash equilibrium in mixed strategies. Without this theorem, equilibrium reasoning would be a guess-and-check exercise with no guarantee of finding anything. With it, the solution concept of Nash equilibrium is *universal* — applicable to every game you might write down — and all subsequent developments in the course proceed knowing that the object they are studying exists. The proof is a beautiful application of Kakutani's fixed-point theorem, generalizing the minimax theorem (1.5) from zero-sum to general-sum games and from bilinear to multilinear payoffs.

Formal Statement

Nash's Theorem (1950). Every finite strategic-form game $G = (N, \{S_i\}, \{u_i\})$ has at least one Nash equilibrium in mixed strategies.

More generally, the theorem holds for any game where:

1. The strategy spaces $\Delta(S_i)$ (or, in continuous games, X_i) are non-empty, compact, and convex subsets of Euclidean space.
2. The payoff functions U_i are continuous in the full profile σ and quasi-concave in own strategy σ_i (holding others fixed).

Under these conditions, a Nash equilibrium exists. For finite games, the multilinear extension automatically satisfies continuity and quasi-concavity (linearity is the special case).

Corollary. Every finite two-player zero-sum game has a value (minimax theorem, Node 1.5) and this value is attained by a pair of Nash equilibrium strategies.

Intuition

Nash's theorem is a "fixed point exists" result. Every player has a best-response correspondence BR_i from the opponent's profile to the set of optimal own-strategies. The joint best-response $\mathrm{BR}(\sigma) = \mathrm{BR}_1(\sigma_{-1}) \times \cdots \times \mathrm{BR}_n(\sigma_{-n})$ is a correspondence from the product of simplices to itself. A fixed point $\sigma^* \in \mathrm{BR}(\sigma^*)$ is a Nash equilibrium.

Kakutani's theorem says: a correspondence from a non-empty, compact, convex set to itself, with non-empty, convex, compact values and upper-hemicontinuous graph, has a fixed point. Each of these properties of BR is derivable from the multilinearity and continuity of U_i, so the theorem applies, and a fixed point exists. That fixed point is a Nash equilibrium.

The conceptual work Nash did was *identifying* the correct setup: map best-response to a correspondence on the simplex, verify the Kakutani hypotheses, conclude existence. The technical work is mostly packaged inside Kakutani. The result is the single most-cited theorem in game theory and the reason Nash won the Nobel Prize in Economics in 1994.

Structural Role

Nash existence is the *launching pad* for every subsequent chapter. Without it:

- **2.5 Rationalizability** would have no guaranteed survivors.

- **3.5 Subgame perfection** would be a refinement of a potentially empty set.
- **4.3 Bayesian Nash** would need its own existence theorem.
- **4.4 Perfect Bayesian equilibrium** would not be guaranteed.
- **7.2 ESS** as a Nash refinement would have no base object.
- **9 Learning** algorithms would have no target fixed point.

Every result that characterizes or refines Nash equilibrium assumes existence. Every computational method for "finding" Nash equilibrium proceeds as if such an object exists. Nash's theorem is what makes that assumption legitimate.

The theorem also generalizes in several directions:

- **Glicksberg (1952):** continuous games on compact metric spaces with continuous payoffs.
- **Debreu (1952), Fan (1952):** games with concave payoffs on compact convex sets.
- **Mertens (1986):** universal belief spaces, Bayesian games with incomplete information.

These generalizations preserve the structural template: non-empty compact convex domain + continuous quasi-concave payoffs + fixed-point theorem.

Mathematical Development

Kakutani's fixed-point theorem

Theorem (Kakutani, 1941). Let $X \subseteq \mathbb{R}^n$ be non-empty, compact, and convex. Let $\phi : X \rightrightarrows X$ be a correspondence with: - $\phi(x)$ non-empty, compact, and convex for every $x \in X$. - Graph $\{(x, y) : y \in \phi(x)\}$ closed (equivalently, ϕ is upper-hemicontinuous).

Then ϕ has a fixed point: there exists $x^* \in X$ with $x^* \in \phi(x^*)$.

Kakutani's theorem generalizes Brouwer's fixed-point theorem (for continuous functions) to correspondences. It is itself a corollary of Brouwer via an approximation argument.

Proof of Nash's theorem

Let $X = \prod_i \Delta(S_i)$. Check: - X **is non-empty, compact, convex.** Each simplex $\Delta(S_i)$ is; products preserve these properties. - BR **is non-empty-valued.** U_i is continuous on the compact domain; the Weierstrass theorem gives a maximizer. - BR **is convex-valued.** If $\sigma_i, \sigma_i' \in \mathrm{BR}_i(\sigma_{-i})$, then by multilinearity of U_i, any convex combination is also in $\mathrm{BR}_i(\sigma_{-i})$. - BR **has closed graph.** By the Berge Maximum Theorem (or direct verification), BR_i is upper-hemicontinuous because the objective is continuous in both arguments and the constraint set is constant.

All Kakutani hypotheses hold. Apply Kakutani: there exists $\sigma^* \in X$ with $\sigma^* \in \mathrm{BR}(\sigma^*)$. By definition, σ^* is a Nash equilibrium. ■

Alternate proof via Brouwer

Define a continuous self-map $f : X \to X$ by

$$f_i(\sigma)(s_i) = \frac{\sigma_i(s_i) + \max(0, U_i(s_i, \sigma_{-i}) - U_i(\sigma))}{1 + \sum_{s_i' \in S_i} \max(0, U_i(s_i', \sigma_{-i}) - U_i(\sigma))}.$$

This map "adjusts" each player's strategy toward pure actions whose payoffs exceed the current expected payoff — a "gain-function" increment. f is continuous on the compact convex X and so has a Brouwer fixed point. At a fixed point, the adjustment terms must vanish for every s_i, which forces $U_i(s_i, \sigma_{-i}) \le U_i(\sigma)$ for all s_i — exactly the Nash condition.

This is Nash's original (1951) proof, simpler than Kakutani but more ad hoc.

Why multilinearity matters

The proof uses multilinearity of U_i in two places: 1. To show BR_i is convex-valued (combinations of best-responses are best-responses). 2. To show U_i is continuous on the simplex (multilinear is automatically continuous).

Without multilinearity, one would need alternative convexity/concavity hypotheses.

Why compactness matters

Compactness gives the Weierstrass theorem (continuous function on compact set attains its max). Without compactness, even single-player decision problems may fail to have optimal strategies. In continuous games on non-compact spaces, one compactifies or truncates before applying a fixed-point theorem.

Worked Example

Verifying all Kakutani hypotheses on Matching Pennies

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

Strategy spaces: Row plays $\sigma_1 = (p, 1-p)$, Column plays $\sigma_2 = (q, 1-q)$, $p, q \in [0,1]$.

- $X = [0,1]^2$**.** Non-empty, compact, convex. ✓
- $U_1(p,q) = p(2q-1) + (1-p)(1-2q) = (2q-1)(2p-1)$**.** Multilinear in p and q. ✓
- $\mathrm{BR}_1(q)$**:** $\{1\}$ if $q > 1/2$; $[0,1]$ if $q = 1/2$; $\{0\}$ if $q < 1/2$. Non-empty, compact, convex. ✓
- **Graph of** BR_1 **is closed:** yes (the jump at $q = 1/2$ is filled in by $[0,1]$). ✓

By Kakutani, a fixed point exists. Indeed it does: $(1/2, 1/2)$, which is a fixed point because at $q = 1/2$, $\mathrm{BR}_1 = [0,1]$ contains $1/2$, and symmetrically. Nash equilibrium confirmed.

A three-player game

Three players each choose $\{A, B\}$. Payoff equals $+1$ if all three agree, -1 otherwise. This is a coordination game with two pure Nash equilibria (AAA and BBB) and one symmetric mixed Nash at $(1/2, 1/2, 1/2)$ — three equilibria total. Kakutani guarantees at least one; we have found three. The oddness theorem (2.3) suggests generically there are an odd number.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Kakutani fixed point	GTO equilibrium of poker — guaranteed to exist in every finite abstraction
Compactness of $\Delta(S_i)$	Finite abstraction of a poker information set — a simplex over a bounded number of actions
Convexity of BR_i	At an indifference point, any mix over tied actions is a best response
Upper hemicontinuity	Small changes in opponent frequencies produce small changes in best responses
Nash's theorem's role	Justifies the existence of "the GTO solution" for any finite poker abstraction before we try to compute it

Concrete situation

Why Cepheus could claim to "solve" HU limit hold'em. Before the 2015 CFR run, it was mathematically certain that a Nash equilibrium of HULHE exists, because HULHE is a finite game (though very large: ~10^14 information sets) and Nash's theorem applies. The existence was not in doubt; the question was computing it.

CFR is an algorithm that converges to an ϵ-Nash equilibrium as the number of iterations grows. Cepheus ran for several months on a cluster and converged to an ϵ so small (one-millibet per game) that the resulting strategy is "essentially unbeatable" in the statistical sense: the entire game value achievable by any opponent is bounded above by ϵ.

Without Nash's theorem, the claim "there is a GTO strategy" would have been conjectural. With Nash's theorem, it was a *theorem* — Cepheus was computing an object whose existence was established 65 years earlier.

What this understanding buys you

GTO exists. Before any computation, the theorem says so. This is not a trivial point. Early poker AI research briefly entertained the notion that the game might be too large to have a "solution" in any meaningful sense. Nash's theorem rules this out: the space of mixed strategies on any finite abstraction is a product of simplices, compact and convex, and the best-response correspondence on it has a fixed point. The fixed point is the GTO strategy.

Finite abstraction matters. Nash's theorem applies to finite games. When a solver abstracts the infinite HUNL tree (e.g., by bucketing bet sizes), the abstraction is finite, and Nash's theorem guarantees a fixed point of the abstracted game. The abstraction error — how much the abstracted Nash differs from the true HUNL Nash — is a separate (and hard) problem, but existence of the abstracted solution is not in doubt.

Continuous-action poker variants. Some toy analyses (e.g., AKQ games, Kuhn poker) use continuous bet sizes. Nash existence extends to such games via Glicksberg's theorem, provided payoffs are continuous and strategy spaces are compact. The theoretical guarantee carries over.

Cross-Links

- **1.2 Mixed strategies** — the simplex domain.
- **1.4 Best-response correspondences** — the operator.
- **1.5 Zero-sum minimax** — the two-player zero-sum case with a simpler LP proof.
- **2.1 Nash equilibrium** — the object the theorem guarantees.
- **2.3 Mixed strategy Nash** — the form the equilibrium takes.
- **Glicksberg's theorem** — continuous-action extension.
- **Real analysis — compactness, continuity, Brouwer/Kakutani** — the mathematical backbone.
- **Systems thinking — fixed points** — the unifying concept.

Structural Compression

Invariant: Every finite game has at least one Nash equilibrium in mixed strategies, guaranteed by the existence of a fixed point of the joint best-response correspondence on the compact convex mixed-strategy product space.

Mechanism: Kakutani's fixed-point theorem applied to a correspondence with the required regularity properties (non-empty, convex, compact values; closed graph), derived from continuity and multilinearity of expected payoffs.

Cross-System Echo: Fixed-point theorems are the universal technique for proving existence of equilibria across mathematical sciences: Brouwer for maps, Kakutani for correspondences, Banach for contractions, Lefschetz for topological equilibria, Tarski for order-preserving maps on lattices. Nash existence is the strategic instance of a pattern that pervades mathematics.

Exercises

1. Verify the four Kakutani hypotheses on BR directly for the Prisoner's Dilemma (mixed-strategy version).
2. Prove that the best-response correspondence in a finite game is upper-hemicontinuous. (Use the Berge Maximum Theorem.)
3. Show that multilinearity of U_i in σ_i implies convexity of $\mathrm{BR}_i(\sigma_{-i})$.
4. Sketch the construction of Nash's original Brouwer-based proof, filling in the definition of the map f and verifying it is continuous.
5. Give a counterexample to Nash existence in an infinite game with non-compact strategy spaces. (E.g., players choose a real number; the payoff depends on the difference.)
6. Show that Glicksberg's theorem extends Nash existence to continuous games on compact metric spaces with continuous payoffs. What replaces multilinearity?
7. **Argue, structurally,** why Nash's theorem is a consequence of topology rather than algebra — that is, why the hypotheses are topological (compactness, continuity) rather than arithmetic (specific functional forms). What would be the analog of Nash's theorem for infinite-player games or games with non-metrizable strategy spaces?

Rationalizability

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.5

Rationalizability (Bernheim 1984, Pearce 1984) is a solution concept that sits between IESDS and Nash equilibrium. It asks: which strategies can be justified by some coherent belief that the opponent is also rational? The answer is the set of strategies surviving iterated deletion of *never-best-response* strategies — strategies that cannot be rationalized as a best response to any opponent belief. Rationalizability relaxes the Nash equilibrium requirement that beliefs be *correct*; it only requires that they be *rationalizable*. It is therefore weaker than Nash (every Nash strategy is rationalizable, but not conversely), and broader than IESDS (for 3+ players). It is the solution concept most closely aligned with pure common knowledge of rationality.

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite game with mixed extension.

A strategy $s_i \in S_i$ is a **best response to some belief** if there exists a probability distribution μ over S_{-i} such that $s_i \in \arg\max_{s_i' \in S_i} \sum_{s_{-i}} \mu(s_{-i}) u_i(s_i', s_{-i})$.

Define iteratively:

$$R_i^{(0)} = S_i, \qquad R_i^{(k+1)} = \{s_i \in R_i^{(k)} : s_i \text{ is a best response to some belief over } R_{-i}^{(k)}\}.$$

The **rationalizable strategies** of player i are

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

Here $R_{-i}^{(k)}$ denotes the product $\prod_{j \neq i} R_j^{(k)}$, and beliefs over it are (in the standard Bernheim–Pearce formulation) required to be **independent** across opponents — each $j \neq i$ draws from their own $R_j^{(k)}$ independently. Some treatments allow correlated beliefs (Aumann), in which case rationalizability coincides with IESDS for all n.

A strategy s_i is **rationalizable** if $s_i \in R_i^{\infty}$.

Intuition

The idea is common knowledge of rationality, stripped of the Nash requirement that beliefs be correct.

A rational player cannot play a strategy that is never a best response: there would be no belief under which playing it is optimal. So rationality alone eliminates "never-best-response" strategies.

A rational player who knows the opponent is rational cannot justify a belief that assigns positive probability to the opponent's never-best-response strategies. So round-2 elimination throws out strategies that require irrationality of the opponent.

Iterating, at round k, a surviving strategy is one that can be justified by a chain of beliefs: "I believe the opponent believes I believe ... (k levels deep) ... that everyone is rational." Rationalizability is the fixed point of this recursion.

In Nash equilibrium, the player's belief about the opponent is *correct*: μ equals the opponent's actual σ_{-i}^*. In rationalizability, the player's belief is *coherent with common rationality* but need not be correct. This makes rationalizability strictly weaker.

Structural Role

Rationalizability bridges dominance and equilibrium:

- **IESDS ⇒ rationalizability (for 2 players).** In two-player games, iterated strict dominance and rationalizability coincide (Pearce 1984).
- **For 3+ players:** rationalizability is weaker than IESDS under independent beliefs. Correlated-belief rationalizability coincides with IESDS for any n.
- **Rationalizability ⇒ Nash existence.** If R_i^∞ is a singleton for each i, the unique profile is a Nash equilibrium.
- **Rationalizability ⊆ Nash strategies.** Every strategy that is played with positive probability in some Nash equilibrium is rationalizable. The converse fails: some rationalizable strategies are not Nash.

Rationalizability also connects to interactive epistemology (Module 4): it is the solution concept for games with *common knowledge of rationality* and no equilibrium assumption. Nash equilibrium additionally requires common knowledge of the equilibrium profile itself, which is a strong assumption that rationalizability dispenses with.

Mathematical Development

Best-response sets

For a finite game, the set of strategies that are best responses to *some* mixed belief is well-defined:

$$\mathrm{BR}_i^{\mathrm{any}}(S_{-i}) = \bigcup_{\mu \in \Delta(S_{-i})} \mathrm{BR}_i(\mu).$$

The iteration $R_i^{(k+1)} = \mathrm{BR}_i^{\mathrm{any}}(R_{-i}^{(k)}) \cap R_i^{(k)}$ converges in finitely many steps for a finite game.

Pearce's Lemma

Lemma (Pearce, 1984). In a finite two-player game, a pure strategy s_i is a best response to some mixed opponent belief iff it is not strictly dominated (possibly by a mixed strategy).

This gives the equivalence between IESDS and rationalizability for two-player games: rationalizability-surviving = IESDS-surviving.

For more than two players, if beliefs are required to be independent (the standard definition), then there can be a strategy that is strictly dominated (even by a mixture) but *becomes* rationalizable because the dominating mixture requires correlation among opponents. The Bernheim–Pearce formulation therefore says rationalizability is slightly coarser than IESDS for $n \geq 3$ unless correlated beliefs are allowed.

Cournot revisited

Iteratively eliminating never-best-responses in Cournot duopoly converges to the unique Cournot–Nash equilibrium. The rationalizable set collapses to a single profile. In contrast, Bertrand competition with capacity constraints can have a large rationalizable set (many strategies are best responses to some belief).

Example: 3 strategies with no dominance but non-trivial rationalizability

	L	C	R
T	$(4,0)$	$(0,4)$	$(2,2)$
M	$(0,4)$	$(4,0)$	$(2,2)$
B	$(1,1)$	$(1,1)$	$(1,1)$

No pure strategy is strictly dominated. Is B rationalizable? Check: is there any belief under which B is a best response? Row's payoff from B is 1 regardless. Row's payoff from T is $4q_L + 0q_C + 2q_R \geq 2$ when $q_L \geq 1/2$ or $q_R \geq 1/2$. If $q_L = q_C = 0, q_R = 1$, Row gets 2 from T, 2 from M, 1 from B — B is not a best response. If $q_L = 0, q_C = 1$, Row gets 0, 4, 1 — B not best. It turns out B is never a best response: the mix $0.5T + 0.5M$ always yields ≥ 2, which beats 1. So B is dominated by $(1/2, 1/2, 0)$ — strictly dominated by a mixed strategy. Rationalizability (= IESDS here) deletes B. Remaining game has only T, M vs L, C, R.

Iterated best response as fictitious-play approximation

Rationalizability is the "limit" of infinitely iterated best response starting from the full strategy space. Fictitious play (Node 9.1) is a dynamic that realizes this iteratively, updating beliefs based on empirical opponent frequencies. In zero-sum and some potential games, fictitious play converges to a Nash equilibrium within the rationalizable set.

Worked Example

Computing rationalizable strategies

	L	C	R
T	$(5,2)$	$(1,3)$	$(0,0)$
M	$(3,4)$	$(0,0)$	$(2,1)$
B	$(0,0)$	$(4,1)$	$(3,2)$

Round 1. Any pure dominated strategy? Check Row: T = (5, 1, 0), M = (3, 0, 2), B = (0, 4, 3). No pure domination. Check mixed: $0.5M + 0.5B$ gives (1.5, 2, 2.5) — does this dominate T's (5, 1, 0)? No, T beats it in column L. No domination.

Column: L = (2, 4, 0), C = (3, 0, 1), R = (0, 1, 2). No domination.

Try: is there a strategy for either player that is not a best response to any belief? Row's T is a best response when Column's belief concentrates on L (Row gets 5). Row's M is a best response to L? $3q_L + 0q_C + 2q_R$ vs T's $5q_L + q_C + 0q_R$. For M to beat T, need $3q_L + 2q_R > 5q_L + q_C \Rightarrow -2q_L - q_C + 2q_R > 0$. With $q_L = 0, q_C = 0, q_R = 1$, this is $2 > 0$ ✓, but then we should also check if M beats B: 2 vs 3 — B is better. So M is not a best response to $(0, 0, 1)$. Try $(0, 0.5, 0.5)$: M gives 1, T gives 0.5, B gives 3.5 — M not best. Try $(0.5, 0, 0.5)$: M gives $1.5 + 1 = 2.5$, T gives 2.5, B gives 1.5. M ties with T, so is a best response. ✓

So M is rationalizable. By similar analysis, T and B are best responses to some beliefs, and all Column strategies are too. The rationalizable set is the entire game in this case — no elimination.

Nash equilibria lie in this set. Computing them is the next step (indifference equations).

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Rationalizable strategy	A line that can be justified under *some* plausible read of opponent play
Never-best-response	A line that cannot be defended regardless of opponent tendencies — "just don't do this"
Iterated rationalizability	Successive rounds of "well, if they're rational, they wouldn't do X, and knowing that, I shouldn't do Y"

Formal object	Poker instantiation
Rationalizable but not Nash	An exploit that works against some plausible opponent but is not GTO
Nash $\subset$ rationalizable	Every GTO action is a best response to GTO, hence rationalizable

Concrete situation

Multi-way pot on the river. You are in a 4-way pot on the river with a medium strength hand. UTG bets large, MP folds, CO calls, and action is on you. The decision space is "fold / call / raise."

Rationalizability analysis. Is "raise" a best response to *any* reasonable read? You would need to believe (a) UTG's betting range is weighted toward thin value and bluffs that will fold to a raise, (b) CO's calling range is weighted toward weak hands that will also fold. If you believe *both*, a raise is a best response. If you believe either is strong, a raise is catastrophic. So a raise is *rationalizable* if you can defend it with a coherent read — say, you've seen UTG bluff this exact spot before and CO is a known caller who folds to raises — but it is *not Nash* unless this belief is actually correct in the equilibrium population.

In Nash (GTO) analysis, the raise frequency in this spot is close to 0 — raising is rarely optimal in 4-way river pots because the ranges are too condensed. But the raise remains *rationalizable* because it is a best response to some belief about opponents' ranges.

Training implication. When a coach tells you "that play is defensible against this opponent type, but it's not GTO," they are saying the play is rationalizable but not Nash — it lies in $R_i^\infty \setminus \{\text{Nash supports}\}$. It is a play that can be justified, but only under specific beliefs, not under the common-rationality fixed point.

What this understanding buys you

Rationalizability is the right concept for exploitation. When you deviate from GTO to exploit a specific opponent, you are selecting a rationalizable strategy that is a best response to your read of them. As long as your read is coherent — i.e., the opponent's strategy is itself rationalizable given *their* reads of *you* — the play is defensible.

Don't confuse "rationalizable" with "good." Many rationalizable plays are losing plays if the belief they rely on is wrong. The discipline is to make the belief explicit, verify the belief is plausible given your information, and only then take the rationalizable-but-not-Nash line. The failure mode is taking a rationalizable line with a sloppy belief.

Nash is the unexploitable floor; rationalizability is the exploit ceiling. GTO is guaranteed not to lose to any opponent. The best rationalizable exploit is guaranteed to win the *most* against a specific opponent, *if* your read is correct. Live poker is the art of navigating between these two.

Cross-Links

- **1.3 Dominance and IESDS** — the upstream concept.
- **1.4 Best response** — the defining operator for rationalizability.
- **2.1 Nash equilibrium** — the special case with correct beliefs.
- **2.6 Correlated equilibrium** — broader than Nash, incomparable to rationalizability.
- **4.3 Bayesian Nash equilibrium** — where beliefs over types become explicit.
- **Epistemic game theory** — formal treatment of common knowledge of rationality.

Structural Compression

Invariant: Rationalizable strategies are exactly those surviving iterated best-response elimination: a strategy survives at round $k + 1$ iff it is a best response to some belief over opponents' round-k surviving strategies.

Mechanism: Iterative best-response filtering on the joint strategy space; for 2 players, equivalent to iterated strict dominance; for $n \geq 3$, equivalent under correlated beliefs.

Cross-System Echo: Rationalizability is the strategic analogue of "provable in at most k steps of a formal system": at round k, a strategy is rationalizable iff it is justifiable by a proof tree of depth k in the logic of rational best response. The limit is the set of strategies provable under unbounded common knowledge of rationality.

Exercises

1. In the Prisoner's Dilemma, show that both Defect profiles are rationalizable and Cooperate is not. Conclude that rationalizability gives the same answer as Nash in dominance-solvable games.
2. In Rock–Paper–Scissors, show that every pure strategy is rationalizable.
3. Prove Pearce's Lemma: a pure strategy s_i in a two-player finite game is never a best response iff it is strictly dominated by a mixed strategy.
4. Construct a 3-player game in which a strategy is rationalizable under correlated beliefs but not under independent beliefs.
5. Show that every strategy in the support of a Nash equilibrium is rationalizable.
6. Give an example of a rationalizable strategy that is not in the support of any Nash equilibrium.
7. **Argue, structurally,** why rationalizability is a weaker solution concept than Nash equilibrium in general but a stronger one than IESDS in the multi-player case (under independent beliefs). Identify which epistemic assumption — correctness of beliefs, independence of beliefs, or common knowledge of rationality — distinguishes each concept.

Correlated Equilibrium

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Nash Equilibrium **Node:** 2.6

Correlated equilibrium (Aumann, 1974) is a generalization of Nash equilibrium that allows players to condition their actions on a shared public signal. It is a strictly broader concept than Nash: every Nash equilibrium is a correlated equilibrium (take the signal to be degenerate), but there are correlated equilibria that are not the mixture of any Nash. It often supports *better* outcomes than any Nash — higher social welfare, more cooperation — and is the solution concept that naturally arises in no-regret learning (Hart–Mas-Colell), making it the central target of learning dynamics (Node 9.2). It is also computationally tractable: correlated equilibria can be found in polynomial time via linear programming, unlike Nash equilibria.

Formal Definition

Let $G = (N, \{S_i\}, \{u_i\})$ be a finite game.

A **correlated equilibrium** is a probability distribution $\mu \in \Delta(S)$ over strategy profiles such that for every player i, every $s_i \in S_i$ with $\mu(s_i) > 0$, and every alternative action $s_i' \in S_i$:

$$\sum_{s_{-i} \in S_{-i}} \mu(s_{-i} \mid s_i)\, u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \mu(s_{-i} \mid s_i)\, u_i(s_i', s_{-i}).$$

Interpretation. A trusted external mediator draws a profile $s \sim \mu$ and tells each player i only their own component s_i. Given this “recommendation,” player i updates their belief about the others’ actions to $\mu(\cdot \mid s_i)$, and the condition says that following the recommendation s_i is a best response under this updated belief.

Equivalent obedience formulation: for every s_i and every function $f : S_i \to S_i$,

$$\sum_s \mu(s) u_i(s) \geq \sum_s \mu(s) u_i(f(s_i), s_{-i}).$$

That is, no deviation rule f is profitable — players have no incentive to deviate from any recommended action.

Every Nash equilibrium is a correlated equilibrium. Take $\mu = \sigma_1^* \otimes \cdots \otimes \sigma_n^*$ (product distribution).

Intuition

Imagine a trusted mediator draws a profile from a distribution μ and whispers to each player what they should play. You only hear your own recommendation. Given what you heard, you form a conditional belief about what the others were told, and ask: is my recommendation a best response under this belief?

If the answer is yes for every player and every possible recommendation they could receive, the distribution μ is a correlated equilibrium. No player benefits from disobeying their recommendation.

The correlation is the key: the mediator’s draws can be *dependent* across players. A standard traffic light is a correlated equilibrium in a “chicken” game: the mediator (the light) tells one driver to go and the other to stop, with perfect negative correlation. Obeying the light is a best response given what you see. This cannot be achieved by a product distribution (= Nash in mixed strategies), because independent mixing would sometimes tell both drivers to go.

Correlated equilibrium expands the solution space by introducing a **correlation device**. It can strictly improve social welfare over the best Nash equilibrium. And it is achievable in practice by any mechanism that can broadcast a public signal (a traffic light, a lottery, a shared randomness source).

Structural Role

Correlated equilibrium is the natural solution concept for settings with shared randomness:

- **2.1 Nash ⇒ 2.6 Correlated equilibrium.** Every Nash is correlated.
- **2.6 ⇒ 9.2 No-regret learning.** Time-averaged play under no-regret algorithms converges to the set of correlated equilibria, not necessarily to Nash.
- **2.6 ⇒ Hart–Mas-Colell theorem.** Individual no-internal-regret dynamics converge to correlated equilibrium.
- **Computational tractability.** Finding a correlated equilibrium is a polynomial-time LP, whereas Nash is PPAD-complete. This makes correlated equilibrium the "computer-feasible" solution concept.
- **Welfare.** Correlated equilibria can strictly Pareto-dominate all Nash equilibria in a game (example: Chicken).

In mechanism design (Module 6), correlated equilibrium arises when the designer provides a correlation device (a lottery or a shared random message). In repeated games (Module 5), the correlation device emerges endogenously from past history.

Mathematical Development

LP formulation

For a finite game with n players and strategy profiles in S, a correlated equilibrium is a distribution $\mu \in \Delta(S)$ satisfying, for every player i and every pair (s_i, s_i') with $s_i \neq s_i'$:

$$\sum_{s_{-i}} \mu(s_i, s_{-i}) \left[u_i(s_i, s_{-i}) - u_i(s_i', s_{-i})\right] \geq 0.$$

This is a linear constraint in μ. Together with $\mu \geq 0$ and $\sum_s \mu(s) = 1$, it defines a polytope. Correlated equilibria are exactly the points in this polytope. Finding one is a feasibility LP — solvable in polynomial time.

Chicken example

	D	C
D	$(0,0)$	$(4,1)$
C	$(1,4)$	$(3,3)$

Nash equilibria: (D, C) and (C, D) pure, plus a mixed one with payoff 8/3 each.

Correlated equilibrium via public signal. Mediator draws one of three profiles uniformly: $(D, C), (C, D), (C, C)$, each with probability 1/3. Tell each player only their own action.

Check player 1's obedience. If told "D," player 1 knows the profile is (D, C) (the only one with D for player 1 — wait, we should check). Actually the profiles are $(D, C), (C, D), (C, C)$. Player 1's action is D only in (D, C). If player 1 is told D, they infer the profile is (D, C), and their expected payoff from obeying is 4 (vs deviation to C giving 3). Obey.

If player 1 is told C, they infer the profile is either (C, D) or (C, C) with equal probability (each conditional probability 1/2). Expected payoff from obeying: $0.5 \cdot 1 + 0.5 \cdot 3 = 2$. From deviating to D: $0.5 \cdot 0 + 0.5 \cdot 4 = 2$. Tie — obeying is weakly optimal. ✓ (In a strict sense, need $\geq$, which holds.)

Player 2's obedience: symmetric. ✓

Welfare comparison. Expected payoff: player 1 gets $1/3(4+1+3) = 8/3$, player 2 gets $1/3(1+4+3) = 8/3$, same as mixed Nash. Player 1+2 sum: $16/3 \approx 5.33$.

Better correlated equilibrium. Try $(D, C), (C, D), (C, C), (C, C)$ with probabilities $1/4, 1/4, 1/2$ — check. Expected payoff: player 1 gets $0.25 \cdot 4 + 0.25 \cdot 1 + 0.5 \cdot 3 = 2.75$. Better than $8/3 \approx 2.67$. Verify obedience: conditioning on D, inference is (D, C), expected payoff 4 vs deviation 3 — obey. Conditioning on C, inference is (C, D) with probability $1/3$, (C, C) with probability $2/3$. Expected payoff from obeying: $1/3 \cdot 1 + 2/3 \cdot 3 = 7/3$. From deviating to D: $1/3 \cdot 0 + 2/3 \cdot 4 = 8/3$. $8/3 > 7/3$ — NOT obedient. This distribution is not a correlated equilibrium.

Correlated equilibrium has structural constraints; not every seemingly good distribution works.

Hart–Mas-Colell theorem

Theorem (Hart–Mas-Colell, 2000). If each player in a repeated game uses a regret-matching algorithm, then the empirical distribution of play converges almost surely to the set of correlated equilibria.

Regret-matching is a simple adaptive rule: the probability of playing action s_i tomorrow is proportional to the total regret (if positive) of not having played s_i in the past. No knowledge of the game or opponents' payoffs is required, only own-payoff feedback.

Polynomial-time computation

Correlated equilibrium is an LP; Nash is PPAD-complete. The gap between these two computational classes is one of the major structural differences between the two solution concepts. In large games, "find a correlated equilibrium" is feasible; "find a Nash equilibrium" may not be.

Worked Example

Coordination game with a traffic signal

Two drivers meet at an intersection. Actions: Go, Stop.

	Go	Stop
Go	$(-10, -10)$	$(1, 0)$
Stop	$(0, 1)$	$(-1, -1)$

Two pure Nash equilibria: $(\text{Go}, \text{Stop}), (\text{Stop}, \text{Go})$. Mixed Nash: symmetric, with probability of $\text{Go} = 1/10$ (indifference calculation).

Correlated equilibrium via traffic light. Mediator tosses a fair coin: heads $\to$ (Go, Stop), tails $\to$ (Stop, Go). Each player is told only their own recommendation. Player 1's expected payoff from obeying "Go" is 1 vs deviation to Stop giving -1 — obey. From obeying "Stop": 0 vs deviation to Go giving -10 — obey. Similarly for player 2. Correlated equilibrium verified.

Expected payoff: $1/2 \cdot 1 + 1/2 \cdot 0 = 1/2$ each. Vastly better than the mixed Nash (which has frequent (Go, Go) crashes) and equal to the average of the two pure Nash.

This is precisely what traffic lights do. The light is the mediator; the "red means stop, green means go" convention is the obedience rule; the independence of the draw from driver identity makes the mechanism fair.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Correlation device	Rare in heads-up cash poker; more relevant in staking networks and multi-way tournaments
Correlated equilibrium	Outcome distribution where obeying the mediator is a best response
Correlation via community cards	Not a correlation device in the technical sense — the cards are public information
Poker bot collusion	A genuine correlation device: two bots sharing information is a correlated mechanism, and the resulting play may be a correlated equilibrium of the multi-way game
No-regret $\to$ correlated equilibrium	CFR and related algorithms converge to correlated equilibria in multi-player games, not to Nash

Concrete situation

Pluribus and 6-max. When Pluribus (Facebook/CMU 2019) solved 6-max no-limit hold'em, it did not, strictly speaking, compute a Nash equilibrium. For $n \geq 3$, Nash equilibrium in extensive-form games is both hard to define cleanly (multiple equilibria with welfare tradeoffs) and computationally PPAD-complete. What Pluribus computed was a strategy that performs well in a no-regret sense — and the Hart–Mas-Colell theorem tells us such strategies converge to *correlated* equilibria of the game, not Nash.

In practice, Pluribus's play was close enough to "nice" correlated equilibria that it dominated top human professionals in play — but there is no theorem saying it plays Nash in the 6-max game. There is a theorem saying it plays a correlated equilibrium asymptotically (modulo subtleties of the abstraction and online blueprint).

This is not a defect; it is the natural solution concept for multi-player games with limited computational resources. Correlated equilibria have higher welfare potential and are computationally easier. For multi-way poker, correlated equilibrium is the right object.

What this understanding buys you

For 2-player poker (HU), GTO = Nash = correlated equilibrium. The simplex of correlated equilibria in a zero-sum two-player game coincides with the Nash set (all live at the minimax value). You do not need correlated equilibrium as a distinct concept for HU.

For 3+ player poker, GTO as Nash becomes ambiguous. There are multiple Nash equilibria in 6-max, with different welfare properties, and the strategies your opponents actually play may not coincide with any specific Nash. Correlated equilibrium is broader and captures more of what "good play" actually looks like in practice. The equilibria Pluribus converges to are correlated, not Nash.

Collusion in poker is formal collusion on a correlation device. Two colluders sharing information via a side channel are implementing a correlation device; the resulting play is a correlated equilibrium of the game *with the side-channel mediator.* This is why it is illegal: correlated equilibrium of the mediated game gives strictly higher welfare to the colluders than Nash of the unmediated game.

Cross-Links

- **2.1 Nash equilibrium** — correlated generalizes Nash.

- **2.3 Mixed Nash** — Nash is the product-distribution special case.
- **2.5 Rationalizability** — correlated CE is incomparable to rationalizability.
- **5 Repeated games** — history serves as an endogenous correlation device.
- **6 Mechanism design** — external correlation devices as mechanisms.
- **9.2 No-regret learning** — dynamics converging to correlated equilibrium.
- **Linear programming** — the algorithmic backbone.

Structural Compression

Invariant: A correlated equilibrium is a joint distribution over profiles such that, conditional on one's own recommendation, obeying the mediator is a best response. The set of correlated equilibria is convex, contains all Nash equilibria, and can strictly improve welfare.

Mechanism: Conditional-expectation best-response constraints, encoded as linear inequalities on the distribution $\mu \in \Delta(S)$; feasibility is polynomial-time.

Cross-System Echo: Correlated equilibrium captures the structural role of *signaling* and *institutional coordination* across systems: traffic lights, auctions with reserve mechanisms, coordinating protocols in distributed systems, shared-randomness communication channels. Whenever a system introduces a common signal, it lifts the equilibrium concept from Nash to correlated.

Exercises

1. For the Chicken game, find a correlated equilibrium with strictly higher expected welfare than the mixed Nash.
2. Write out the full LP for computing correlated equilibrium of a 3×3 coordination game.
3. Prove that every Nash equilibrium is a correlated equilibrium (product-distribution case).
4. Show that the set of correlated equilibria is convex in $\Delta(S)$.
5. Construct a 2×2 game in which the set of correlated equilibrium payoffs strictly contains the set of Nash equilibrium payoffs.
6. State the Hart–Mas-Colell theorem and explain why regret-matching is only guaranteed to converge to correlated, not Nash, equilibria.
7. **Argue, structurally,** why the computational gap between Nash (PPAD) and correlated (LP) suggests that correlated equilibrium is a "more natural" target for learning algorithms in large games. What does this imply about the practical meaning of "solving" a game?