Nash Equilibrium — Definition

Game Theory · Module 2 — Nash Equilibrium

Node 2.1

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).

Game Theory · Module 2 — Nash Equilibrium · Node 2.1
Concepts: nash-equilibrium, fixed-point
Prerequisites: 1.1, 1.2, 1.4
Enables: 2.2, 2.3, 2.4, 2.5, 2.6, 3.5

hbar.university — own.your.cognition