Bayesian Nash Equilibrium

Game Theory · Module 4 — Information And Beliefs

Node 4.3

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Information and Beliefs **Node:** 4.3

Bayesian Nash equilibrium (BNE) is the solution concept for Bayesian games: a strategy profile in which every type of every player is playing a best response given their interim belief. It extends Nash equilibrium (2.1) to incomplete-information settings by making strategies type-contingent and best-response conditions interim. Harsanyi (1967–68) introduced BNE alongside the Bayesian game framework; it underwrites auction theory, mechanism design, and every incomplete-information application in game theory. For poker, Bayesian Nash equilibrium is what the solver computes — type = hole cards, strategy = what you do with each hand, and the equilibrium is a fixed point where both players are best-responding to each other's type-contingent strategies given the common prior.

Formal Definition

Let $G^B = (N, \{\Theta_i\}, \{S_i\}, p, \{u_i\})$ be a Bayesian game with common prior p.

Type-contingent strategy. A pure strategy is $s_i : \Theta_i \to S_i$. A behavioral (mixed) type-contingent strategy is $\sigma_i : \Theta_i \to \Delta(S_i)$.

Interim expected payoff. Given opponent strategy profile σ_{-i} and own type θ_i, player i's interim expected payoff from action $a \in S_i$ is

$$U_i(a \mid \theta_i, \sigma_{-i}) = \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) \sum_{a_{-i}} \sigma_{-i}(\theta_{-i})(a_{-i}) \cdot u_i(a, a_{-i}, \theta_i, \theta_{-i}).$$

Bayesian Nash equilibrium. A strategy profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a **Bayesian Nash equilibrium** if, for every player i and every type $\theta_i \in \Theta_i$ with $p(\theta_i) > 0$,

$$\sigma_i^*(\theta_i) \in \arg\max_{\sigma_i \in \Delta(S_i)} U_i(\sigma_i \mid \theta_i, \sigma_{-i}^*).$$

Equivalently (by linearity): for every type θ_i and every action a in the support of $\sigma_i^*(\theta_i)$, a is a best response to σ_{-i}^* at θ_i.

Ex ante characterization. A profile is BNE iff each player's ex ante expected payoff (averaged over all types) is maximized by their strategy, holding opponents' strategies fixed. Under common prior, ex ante and interim optimality are equivalent at types in the support of p.

Theorem (existence). Every finite Bayesian game has a Bayesian Nash equilibrium. (Proof: treat the Bayesian game as a Nash equilibrium problem in a suitably defined "agent-normal-form" game where each type of each player is a separate pseudo-agent; apply Nash's existence theorem to the finite game on strategy profiles $\prod_{i,\theta_i} \Delta(S_i)$.)

Intuition

Think of each type of each player as a "separate pseudo-agent" with its own best-response problem. The common prior ties these agents together through the interim belief computation: when type θ_i of player i decides what to do, it computes an expected payoff by averaging over the opponent's types weighted by $p(\theta_{-i} \mid \theta_i)$.

A Bayesian Nash equilibrium is a profile where *every one* of these pseudo-agents is best-responding to the pseudo-agents of the other players. Equivalently: every type of every player is playing optimally given their beliefs, and no type wants to deviate.

Why "pseudo-agent"? Because the players are the real agents, but the action plan must account for every contingency (every type). Given the common prior, the whole profile can be computed all at once as a joint fixed point.

The notion of equilibrium here is *interim* — checked type by type. This is important because a player of type θ_i is not averaging over her own types (she knows her type); she is averaging only over opponents' types. The ex ante formulation (averaging over everyone's types) is equivalent under common prior but is a different computation — it ignores the conditional structure.

One conceptual subtlety: **BNE is just Nash equilibrium of the agent-normal-form game.** The Bayesian game, after the Harsanyi transformation, is a game of imperfect information with a chance node at the start. Its Nash equilibrium is the BNE. So BNE is not a new solution concept — it is Nash equilibrium applied to the right object.

Structural Role

Bayesian Nash equilibrium is the solution concept for static incomplete-information games:

- **2.1 Nash ⇒ 4.3 BNE.** BNE is Nash on the agent-normal-form game.
- **4.1, 4.2 ⇒ 4.3.** Types and type spaces are the prerequisites; BNE is the equilibrium on them.
- **4.3 ⇒ 4.4 PBE.** For dynamic games, PBE extends BNE with belief consistency at information sets.
- **4.3 ⇒ 4.5 Signaling games.** Signaling equilibria are refinements of BNE.
- **4.3 ⇒ 6.1 Mechanism design.** The designer chooses rules such that BNE yields desired outcomes.
- **4.3 ⇒ 6.3 Auction theory.** Bidders play BNE of the induced Bayesian auction game.
- **4.3 ⇒ 9.3 CFR.** CFR's convergence theorem is to an ϵ-BNE (or equivalent Nash) in the extensive-form representation.

Mathematical Development

Equivalence with Nash in the agent-normal-form game

Given Bayesian game G^B, define the **agent-normal-form game** G^{ANF}:

- **Players:** one per (player, type) pair (i, θ_i) — call these "agents."
- **Strategies:** each agent (i, θ_i) picks $\sigma_i(\theta_i) \in \Delta(S_i)$.
- **Payoffs:** agent (i, θ_i)'s payoff is $U_i(\sigma_i(\theta_i) \mid \theta_i, \sigma_{-i}) \cdot p(\theta_i)$ (weighting by prior probability of being this type).

Then the Nash equilibria of G^{ANF} correspond exactly to the Bayesian Nash equilibria of G^B. Nash's existence theorem applied to G^{ANF} gives BNE existence for finite Bayesian games.

Purification (Harsanyi, 1973)

Harsanyi showed that almost every mixed-strategy Nash equilibrium of a complete-information game is the *limit* of pure-strategy Bayesian Nash equilibria of a perturbed incomplete-information game. Specifically, add a small independent payoff perturbation $\epsilon\theta_i$ to each player's payoffs, with θ_i drawn from a smooth distribution. The resulting Bayesian game has a pure-strategy BNE (each type plays a pure action), and as $\epsilon \to 0$, the distribution of actions induced by the pure BNE converges to the mixed strategy Nash of the original game.

This is Harsanyi's **purification theorem**. It provides an epistemic justification for mixed strategies: the randomness is not arbitrary — it reflects hidden private information that would deterministically select an action at each type.

Efficiency and information aggregation

In Bayesian games, equilibria are generically *not* ex-post efficient. Even if the complete-information game would have a Pareto-optimal equilibrium, under incomplete information the interim rationality constraints rule out many ex-post-efficient profiles. This motivates mechanism design (Module 6): by carefully structuring the game, the designer can enforce efficient outcomes as BNE.

Myerson–Satterthwaite (1983). In a bilateral trade setting with private values, no BNE of any trading mechanism can achieve ex-post efficiency while maintaining budget balance and voluntary participation. This is a negative result: information asymmetry *forces* inefficiency in equilibrium.

Beauty-contest / level-k refinements

BNE assumes common knowledge of rationality and common prior. Empirically, subjects often play at "level-k" — optimal response to level-$k-1$ opponents. Cognitive hierarchy models (Camerer–Ho–Chong, 2004) model this as a distribution over levels. For $k \to \infty$, level-k thinking converges to BNE. For finite k, it gives a bounded-rationality refinement.

Worked Example

The Bayesian coordination game (completed)

Two players choose $\{A, B\}$. Player 1 is type-less (one type). Player 2 has $\theta_2 \in \{H, L\}$ with $p(H) = p(L) = 1/2$.

Payoffs: - Player 1 (independent of type): $u_1(A, A) = 2, u_1(B, B) = 1, u_1(\text{mismatch}) = 0$. - Player 2 type H: $u_2(A, A) = 1, u_2(B, B) = 0, u_2(\text{mismatch}) = 0$ — prefers coordinating on A. - Player 2 type L: $u_2(A, A) = 0, u_2(B, B) = 1, u_2(\text{mismatch}) = 0$ — prefers coordinating on B.

Guess: Player 2's strategy: type H plays A, type L plays B.

Player 1's expected payoff from A: $0.5 \cdot 2 + 0.5 \cdot 0 = 1$ (from type H match, type L mismatch).

Player 1's expected payoff from B: $0.5 \cdot 0 + 0.5 \cdot 1 = 0.5$ (from type H mismatch, type L match).

Player 1 plays A.

Player 2 best-response check given Player 1 plays A: - Type H gets 1 from A (coordination), 0 from B (mismatch). Plays A. ✓ - Type L gets 0 from A (mismatch), 0 from B (mismatch). Indifferent; any action is a best response.

So a BNE exists: $(A, (A, \text{anything}))$. The L-type's action doesn't matter because it doesn't affect Player 1's best response. There are multiple BNE parameterized by L-type's action.

Non-trivial case. Change Player 1's payoffs to $u_1(A, A) = u_1(B, B) = 2, u_1(\text{mismatch}) = 0$ and Player 2's weight $p(H) = 0.4, p(L) = 0.6$. Then Player 1 playing A yields $0.4 \cdot 2 + 0.6 \cdot 0 = 0.8$; playing B yields $0.4 \cdot 0 + 0.6 \cdot 2 = 1.2$. Player 1 plays B. Player 2 type H best response to Player 1 = B: type H gets 0 from A (mismatch) and 0 from B (mismatch — type H wants A). Type H is indifferent. For an equilibrium, type H plays some action, both yield 0. Check if there's an equilibrium where type H plays A: Player 1's expected payoff from A is $0.4 \cdot 2 + 0.6 \cdot 0 = 0.8$ (vs 1.2 from B); Player 1 still plays B. So (B, (A, B)) is a BNE with Player 2 effectively miscoordinating in type H and coordinating in type L.

First-price auction (two bidders, uniform values)

Bidders $i \in \{1, 2\}$ with private values $v_i \sim U[0, 1]$. Each bids $b_i \geq 0$; highest bid wins and pays bid.

Guess: symmetric linear strategy $b_i(v_i) = \alpha v_i$ for some $\alpha \in (0, 1)$.

Given opponent plays $b_j(v_j) = \alpha v_j$, bidder 1 of type v_1 bids b_1 to maximize:

$$U_1(b_1 \mid v_1) = (v_1 - b_1) \cdot \Pr[b_1 > \alpha v_2] = (v_1 - b_1) \cdot \frac{b_1}{\alpha}.$$

FOC: $\frac{\partial}{\partial b_1}\left[(v_1 - b_1)\frac{b_1}{\alpha}\right] = \frac{v_1 - 2b_1}{\alpha} = 0 \Rightarrow b_1 = v_1/2$.

So $\alpha = 1/2$: the symmetric BNE is $b_i(v_i) = v_i/2$. Each bidder shades their bid by half to trade off winning probability against surplus.

General symmetric BNE with n bidders: $b_i(v_i) = \frac{n-1}{n} v_i$. More bidders $\Rightarrow$ less shading.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Type-contingent strategy $\sigma_i(\theta_i)$	Your range strategy: what you do with each hole-card combination at each info set
Interim best response	Your best action given your specific hand and your belief about opponent's range
Interim belief $p(\theta_{-i} \mid \theta_i)$	Your posterior over opponent's hole cards given observed betting
Bayesian Nash equilibrium	**The GTO solution of the hand** — both players' range strategies are mutual best responses
Agent-normal-form game	"Each hand is a separate pseudo-agent choosing its action" — the solver's internal representation
Purification	Mixed strategies in poker (e.g., 50/50 bet/check with KQ) reflect implicit card-removal randomness

Concrete situation

The flop spot, formally. BTN c-bet flop decision on $Kd9s4c$. BTN's range has ~200 combos; BB's range has ~250 combos. BTN's strategy is a function $\sigma_{BTN} : \Theta_{BTN} \to \Delta(\{\text{bet } 33\%, \text{bet } 75\%, \text{check}\})$. BB's strategy, after observing BTN's action, is similar.

At a Bayesian Nash equilibrium: - For every combo θ_{BTN}, BTN's action at θ_{BTN} is a best response given BB's strategy. Concretely: for each combo, EV of checking must equal EV of betting 33% must equal EV of betting 75%, if all three are in the solver's output with positive probability. If only betting 33% is positive, then betting 33% must have strictly higher EV than check and 75%. - For every combo θ_{BB}, BB's action after BTN's chosen sizing is a best response given BTN's strategy.

This is the formal Bayesian Nash equilibrium condition. The solver's output is a strategy profile satisfying this at every info set, every combo, simultaneously.

CFR computes an ε-BNE. Counterfactual regret minimization iterates on regret values at each (player, info set, action) triple and converges to a strategy profile such that no combo of any player has more than ϵ counterfactual regret. This is equivalent to an ϵ-Bayesian Nash equilibrium of the poker game — the smaller ϵ, the closer to exact BNE. "GTO" in practice is ϵ-BNE with ϵ small enough (e.g., 0.1% of the pot) that exploitation gains are negligible.

Solver mixing is purification in disguise. When the solver outputs "with $A\spadesuit K\spadesuit$ on $Kd9s4c$, bet 33%: 62%, check: 38%," this is a mixed strategy. Harsanyi's purification says this mixing is the limit of pure strategies in a perturbed Bayesian game where "nearby" combos have slightly different payoffs and deterministically choose different actions. In poker, the purification is essentially achieved by card removal: $A\spadesuit K\spadesuit$ is close to $A\heartsuit K\spadesuit$ but not identical; the solver handles them separately, and the aggregate frequency of betting across all "AKs-like" combos is ~62%. The mixing you see in the output is the composition of pure strategies over these finely distinguished combos.

What this understanding buys you

GTO = Bayesian Nash equilibrium, precisely. When a poker player says "GTO says bet 33% with 62% frequency," they are reporting the strategy component of a Bayesian Nash equilibrium at a specific information set. This is not a heuristic, not an approximation, not a suggestion — it is the formal solution concept of game theory, instantiated in poker. The entire apparatus of BNE (existence, best response, interim optimality) applies.

Every combo is a best response. The key property of BNE at your info set is that *every hand in your range* is playing a best response. If the solver has you mixing with AKs, both actions (bet and check) must yield the same EV for AKs — otherwise you should only play the better one. This is the indifference principle (from 2.3) applied to type-contingent strategies. Coaches who say "AKs is indifferent between betting and checking" are stating the BNE indifference condition.

Opponent mixing is their purification, and you exploit deviations. If your opponent plays a strategy that is not BNE — say, they always bet 33% with AKs (100% instead of 62%) — then some *other* hand in their range is no longer indifferent, and you can exploit. The exploit target is the combo whose EV calculus is broken by their deviation. This is how solver-aware humans beat non-solver opponents: detect deviations from BNE, then play the best response to the deviated strategy (not to BNE).

Interim vs ex ante matters for study. When you study a spot, you look at your strategy conditional on your hand (interim). When you look at "range performance" or "how often this range wins showdown," you are doing ex ante analysis. BNE is defined at the interim level — each hand must be best-responding. Ex ante equivalence holds under common prior. If you find an ex post analysis that differs from the interim conclusion, it is because you are implicitly averaging over something that shouldn't be averaged. Always reason at the interim level when asking "should I play this way with this specific hand."

Cross-Links

- **2.1 Nash equilibrium** — the base concept extended here.
- **4.1 Bayesian games** — the underlying framework.
- **4.2 Types and type spaces** — the strategy domain.
- **4.4 Perfect Bayesian equilibrium** — dynamic refinement.
- **4.5 Signaling games** — canonical Bayesian game with receiver best response.
- **6.1, 6.3 Mechanism design, auctions** — BNE underlies both.
- **9.3 CFR** — algorithmic computation of ϵ-BNE.

Structural Compression

Invariant: A Bayesian Nash equilibrium is a type-contingent strategy profile in which every type of every player is playing a best response to the profile at their interim belief; equivalently, it is a Nash equilibrium of the agent-normal-form game where each (player, type) pair is a separate pseudo-agent weighted by its prior probability.

Mechanism: Interim expected payoff computation via Bayes' rule on the common prior, followed by type-by-type best response optimization, with the joint fixed point constituting the equilibrium.

Cross-System Echo: Bayesian Nash equilibrium is the game-theoretic counterpart of a Bayes-optimal decision rule in statistical decision theory: in both, an agent conditions on private information to compute a posterior and acts optimally against that posterior. The difference is that in BNE the "posterior" is about opponents' types rather than nature, and the opponents are themselves acting optimally against their own posteriors — a fixed point, not a one-sided optimization.

Exercises

1. Verify that the symmetric linear BNE of a first-price auction with two bidders and uniform values is $b_i(v_i) = v_i/2$.
2. Show that every finite Bayesian game has a BNE by applying Nash's existence theorem to the agent-normal-form game.
3. Construct a Bayesian coordination game where the unique BNE has both players miscoordinating in expectation.
4. In a Bayesian prisoner's dilemma where each player's cooperation cost is private (two types, H and L, with independent draws), compute the BNE. Does it depend on the distribution over types?
5. Show that BNE is equivalent to ex-ante optimality under common prior. Where does the argument break without common prior?
6. Formulate heads-up NLHE flop spot as a Bayesian game. Specify strategies, interim beliefs (accounting for blockers), and write the BNE indifference condition for one specific combo.
7. **Argue, structurally,** why Harsanyi's purification theorem provides an epistemic justification for mixed strategies. What does it say about the "arbitrariness" of mixing?

Game Theory · Module 4 — Information And Beliefs · Node 4.3
Concepts: bayesian-nash-equilibrium, incomplete-information, expectation-value, conditional-probability
Prerequisites: 2.1, 4.1, 4.2
Enables: 4.4, 4.5, 4.6, 6.1, 6.3

hbar.university — own.your.cognition