Types and Type Spaces

Game Theory · Module 4 — Information And Beliefs

Node 4.2

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Information and Beliefs **Node:** 4.2

Node 4.1 introduced the Bayesian game as a tuple with a set of types Θ_i for each player. This node zooms in on the type object itself: what a type encodes, how it relates to hierarchies of beliefs about beliefs, and when two different-looking type spaces are strategically equivalent. The key insight, due to Harsanyi and formalized by Mertens–Zamir (1985), is that any coherent hierarchy of beliefs collapses into a point in a **universal type space**. For applied work — including poker — one almost always uses a finite type space that directly encodes payoff-relevant private information. This node makes precise what "a type" contains and sets up Bayesian Nash equilibrium (4.3) as equilibrium over type-contingent strategies.

Formal Definition

Let N be the set of players, and for each player i let Θ_i be a measurable set of **types**. The product $\Theta = \prod_{i \in N} \Theta_i$ is the set of **type profiles**.

Payoff-parameter type. The simplest type encodes player i's private payoff parameter: a type θ_i indexes a payoff function $u_i(\cdot, \theta_i)$. Opponents do not observe θ_i.

Belief-hierarchy type. A richer type also encodes player i's beliefs about others. Define the hierarchy recursively:

- **Level 0:** X_i^0 = the payoff-relevant state for player i (their payoff parameter, private signal).
- **Level 1:** $X_i^1 = X_i^0 \times \Delta(X_{-i}^0)$ — own parameter plus a belief about others' parameters.
- **Level $k+1$:** $X_i^{k+1} = X_i^k \times \Delta(X_{-i}^k)$ — own parameter, belief about others' levels up to k.

A **belief hierarchy** for player i is an element $(\mu_i^1, \mu_i^2, \ldots)$ with $\mu_i^k \in \Delta(X_{-i}^{k-1})$, subject to **coherence**: each μ_i^{k+1} has marginal on X_{-i}^{k-1} equal to μ_i^k.

Type space. A type space is a tuple $(\Theta_i, \hat{\theta}_i, \pi_i)_{i \in N}$ where each type $\theta_i \in \Theta_i$ is associated with:

- $\hat{\theta}_i(\theta_i) \in X_i^0$: the payoff parameter implied by type θ_i.
- $\pi_i(\theta_i) \in \Delta(\Theta_{-i})$: an interim belief about the other players' types.

Iterating these two maps generates a belief hierarchy from θ_i. The type space is **universal** if every coherent belief hierarchy is represented by exactly one type.

Theorem (Mertens–Zamir, 1985). The universal type space exists and is unique up to isomorphism. Every type space embeds into it.

Common prior. A type space has the **common prior property** if there exists a single distribution $p \in \Delta(\Theta)$ such that each $\pi_i(\theta_i)$ equals the conditional $p(\theta_{-i} \mid \theta_i)$ derived from p.

Intuition

Two distinct questions live inside "what player 2 doesn't know about player 1":

1. **What is player 1's payoff parameter?** (their valuation, their hand, their cost.)
2. **What does player 1 believe about player 2's payoff parameter?** And what does player 1 believe player 2 believes player 1 believes? And so on.

Question 1 is simple: player 1 has a parameter, player 2 has a distribution over it. Question 2 is the Russian-doll one: beliefs about beliefs about beliefs, stretching infinitely. Harsanyi's genius was to notice that *you can fold the whole tower into a single "type" object* — a point in a sufficiently rich space — so that knowing the type tells you both the payoff parameter and the beliefs at every level.

Why does this work? Because the beliefs at level $k+1$ are distributions over level-k beliefs, each belief at every level is eventually a distribution over the original payoff parameters. Coherence (the marginals must match) forces the hierarchy to be consistent. Under coherence, the hierarchy is fully determined by a single probability measure, and that measure is what the "type" represents.

The common prior assumption is a separate and strong restriction. Without it, each player can have an arbitrary interim belief $\pi_i(\theta_i)$; with it, all interim beliefs are consistent with a single joint distribution that players would have agreed on before observing their types. Aumann's agreement theorem (1976) and no-speculative-trade rely on this property.

For applied work, you almost never write down the belief hierarchy explicitly. Instead, you pick a finite type space where each θ_i is literally a payoff parameter (a hole card combination, a valuation, a cost), and the hierarchy is generated automatically from the common prior. This is fine: the Mertens–Zamir theorem guarantees that what you lose is only redundant representation, not strategic content.

Structural Role

Types and type spaces are the bridge between Bayesian games as tuples (4.1) and Bayesian Nash equilibrium (4.3):

- **4.1 ⇒ 4.2.** Bayesian games are defined; this node unpacks the type object.
- **4.2 ⇒ 4.3.** Bayesian Nash equilibrium is equilibrium over type-contingent strategies.
- **4.2 ⇒ 4.4.** Perfect Bayesian equilibrium extends to sequential games with belief updating.
- **4.2 ⇒ 4.5.** Signaling games use types where the sender's type is private and the receiver updates beliefs from signals.
- **4.2 ⇒ 4.6.** Sequential equilibrium imposes consistency on beliefs at off-path information sets.
- **4.2 ⇒ 6.1, 6.3.** Mechanism design and auction theory rest on a type space for private valuations.

The universal type space result is mostly of theoretical importance — it guarantees the foundational framework is well-posed. Applied game theory, including every poker solver, uses small finite type spaces.

Mathematical Development

Coherence of belief hierarchies

A belief hierarchy $(\mu_i^1, \mu_i^2, \ldots)$ is **coherent** if for every k,

$$\mathrm{marg}_{X_{-i}^{k-1}} \mu_i^{k+1} = \mu_i^k.$$

Incoherent hierarchies are rejected: they would correspond to a player believing contradictory things at different levels. Under coherence, the sequence of beliefs converges to a single probability measure on a suitable limit space, by Kolmogorov's extension theorem.

Brandenburger–Dekel (1993). The set of coherent hierarchies, with an additional "common knowledge of coherence" requirement, is homeomorphic to a compact space — the universal type space. Every type space maps into it continuously.

Finite type spaces vs universal type space

A **finite type space** uses Θ_i = finite set (e.g., $\{H, L\}$ for "high-cost" vs "low-cost" player), with a specified common prior $p \in \Delta(\Theta)$. The induced hierarchy is derived by iterating π_i.

The universal type space is infinite-dimensional and useful for existence theorems and foundational work. In applied settings, one argues "pick a sufficiently rich finite type space for the payoff-relevant parameters, assume common prior, and compute." This gives all the strategic content of the universal space for the problem at hand.

When finite type spaces suffice. If the payoff parameter takes finitely many values and players know the distribution, a finite type space with types indexing parameter values and common prior equal to the parameter distribution is strategically equivalent to the full universal type space (for this problem).

Common prior property and its strength

The common prior assumption is logically separate from the definition of types. A type space without common prior has types that specify interim beliefs directly; there need not be any joint prior consistent with all of them.

Why common prior is strong. Aumann argued (1987) that rational agents with different priors are behaving as if they had different information, and there is no way to distinguish prior from posterior without referring to some earlier state. He therefore advocated common prior as a modeling default. Critics argue that prior heterogeneity is real and economically important (Morris 1995; asset pricing with disagreement).

In game theory practice, most models assume common prior. Departures are flagged as "belief heterogeneity" or "non-common-prior types."

Type space morphisms and equivalence

Two type spaces are **strategically equivalent** if they induce the same set of Bayesian Nash equilibrium outcome distributions for any game played on them. Equivalence classes of type spaces embed uniquely into the universal type space. This is the sense in which "the universal type space is universal": it represents every equivalence class exactly once.

Different authors use different finite type spaces for the same problem (e.g., different poker solver ranges); as long as the payoff-parameter structure and common prior agree, the equilibria are the same.

Interim, ex ante, and ex post

Three vantage points on the Bayesian game:

- **Ex ante (before types are drawn).** Expected payoffs averaged over own type and opponents' types, under the common prior.
- **Interim (after own type is observed but before play).** Expected payoffs averaged over opponents' types only, conditional on own type.
- **Ex post (after play is realized).** The actual realized payoffs.

Bayesian Nash equilibrium is typically defined at the interim stage: each type of each player must be playing a best response given the interim belief. This is equivalent to ex ante optimality when the common prior is used.

Worked Example

A two-player Bayesian game with finite types

Player 1 has type $\theta_1 \in \{H, L\}$; Player 2 has type $\theta_2 \in \{H, L\}$. Common prior:

$\theta_1 \setminus \theta_2$	H	L
H	0.3	0.1
L	0.1	0.5

Marginal distributions. $p(\theta_1 = H) = 0.4, p(\theta_1 = L) = 0.6$; similarly for θ_2.

Interim beliefs.

$$p(\theta_2 = H \mid \theta_1 = H) = \frac{0.3}{0.4} = 0.75, \quad p(\theta_2 = L \mid \theta_1 = H) = 0.25.$$

$$p(\theta_2 = H \mid \theta_1 = L) = \frac{0.1}{0.6} = 0.167, \quad p(\theta_2 = L \mid \theta_1 = L) = 0.833.$$

The types are **positively correlated**: knowing you are H raises your probability that the opponent is H from 0.4 to 0.75.

Belief hierarchy (level 2). Given $\theta_1 = H$, Player 1's belief about Player 2's belief about θ_1 is derived by asking: *given Player 2's type is H (prob 0.75 from Player 1's view), what does Player 2 believe about θ_1?* Player 2 of type H thinks $p(\theta_1 = H \mid \theta_2 = H) = 0.3/0.4 = 0.75$, and similarly Player 2 of type L thinks $p(\theta_1 = H \mid \theta_2 = L) = 0.1/0.6 = 0.167$. So Player 1's level-2 belief is a mixture $0.75 \cdot$ (Player 2 of type H's belief) $+$ $0.25 \cdot$ (Player 2 of type L's belief), all derived automatically from the common prior via iterated Bayes' rule.

The entire infinite hierarchy is generated from the 4-entry common prior table. This is the compactness of types: infinite hierarchy, finite representation.

Non-common-prior types

Suppose instead Player 1's interim belief $\pi_1(H \mid \theta_1 = H) = 0.9$ and Player 2's interim belief $\pi_2(H \mid \theta_2 = H) = 0.5$. Can these come from a common prior?

Attempted common prior: let $p(H, H) = a$, $p(H, L) = b$, $p(L, H) = c$, $p(L, L) = d$. Then $0.9 = a/(a + b) \Rightarrow a = 9b$, and $0.5 = a/(a + c) \Rightarrow a = c$. So $a = 9b = c$, and normalization $a + b + c + d = 1 \Rightarrow 9b + b + 9b + d = 1 \Rightarrow 19b + d = 1$. This is one equation in two unknowns — there exist such (b, d) satisfying it (e.g., $b = 0.05, d = 0.05, a = 0.45, c = 0.45$).

So these interim beliefs are compatible with a common prior. This illustrates that common prior is not automatic from interim beliefs, but is a testable constraint — and when satisfied, the common prior may still be non-unique.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Type θ_i	Your specific hole cards — the unique realization of your private information

Formal object	Poker instantiation
Type space Θ_i	The set of hole-card combinations in your range — e.g., all hands BTN opens preflop
Payoff parameter	How strong your hand is against the opponent's range — your equity
Common prior	The deck is uniform, and both players agree on it (standard modeling assumption)
Interim belief $\pi_i(\theta_{-i} \mid \theta_i)$	Your posterior over the opponent's hole cards given everything observed so far, including the cards you hold
Belief hierarchy	"I think they have AK, and I think they think I have a bluff, and I think they think I think they're bluffing, . . . "
Finite vs universal type space	Solvers use finite type spaces (bucketed ranges); universal type space is theoretical

Concrete situation

Range interpretation as type space. When you load a solver and see a preflop range chart like "BTN opens 22+, A2s+, A5o+, K9s+, KTo+, Q9s+, QTo+, J9s+, JTo, T9s, 98s, 87s, 76s, 65s," you are specifying a type space for the BTN player. Each element of the range is a possible type. The common prior is the uniform distribution over the deal, and the BTN's "type" is drawn from this range by Bayes' conditioning on the observation that BTN opened (since some hands fold and some open).

Card removal and correlation. Suppose BTN opens and BB calls with range R_{BB}. You (BTN) hold $A♠K♠$. What is your interim belief about BB's hand?

Two effects: 1. **Range filtering.** BB's hand is in R_{BB}. 2. **Card removal.** BB's hand cannot contain $A♠$ or $K♠$.

The interim belief is $p(\theta_{BB} \mid \theta_{BTN} = A♠K♠)$ = uniform over R_{BB} minus combinations containing the blockers, renormalized. This is the formal object the solver uses internally — Bayes' rule on hole cards conditional on observed actions and own cards.

Belief hierarchy in action. Consider "level reasoning" in poker: - Level 0: "I have AK, I'm strong." - Level 1: "My opponent has range R, I should play best response to R." - Level 2: "My opponent is thinking about my range R', they will play best response to R'." - Level 3: "My opponent thinks I think they have R, so they will deviate from R by bluffing more." - . . .

Each level is a belief about beliefs. In equilibrium (Bayesian Nash), all levels are consistent: your strategy is best-response to their strategy given your belief, and their strategy is best-response to yours given their belief, and these beliefs are the common-prior posteriors. The universal type space is the formal setting in which this consistency is well-defined.

Finite solver type space = universal in practice. For a specific spot (specific board, specific preflop history), the solver uses a finite type space per player (the preflop calling range, which has ~100–500 combos). This is strategically equivalent to the universal type space for the relevant payoff-relevant information. The Mertens–Zamir theorem guarantees nothing is lost.

What this understanding buys you

Type = hole cards is exact, not metaphor. When you read "in a Bayesian game, player i's type encodes their private information," and you're a poker player, the correct interpretation is literal: your type is $A♠K♠$, period. Every other object in the Bayesian game (type space, common prior, interim belief, type-contingent strategy) has a precise poker counterpart. The mathematics is the game.

Common prior is what "agreed ranges" means. GTO play assumes both players agree on what each player's preflop range is. This is a common prior on the joint type space. If you and your opponent disagree about what "BTN opens" or "BB defends" means, you are playing a non-common-prior game, and exploitation becomes possible (or impossible, depending on direction). This is the formal content of "the solver assumes equilibrium ranges."

Level reasoning collapses to equilibrium fixed point. "Level 1, level 2, level 3 thinking" is informal; the formal object is the Bayesian Nash equilibrium, which is the fixed point of all levels. Your interim strategy and opponent's interim strategy are jointly best-responses given the common prior. Solving for level-k thinking for finite k gives bounded rationality models (cognitive hierarchy, level-k models); as $k \to \infty$ and the beliefs become consistent, you converge to the Bayesian Nash equilibrium. This is what every solver computes.

Card removal is Bayesian conditioning, not a side effect. Poker players often treat "card removal" as a separate consideration ("don't forget blockers!"), but formally it is just the computation of $\pi_i(\theta_{-i} \mid \theta_i)$ under the common prior of a uniform deck. The blockers are automatically in the conditional distribution. When a solver reports your optimal frequencies, card-removal effects are already absorbed.

Cross-Links

- **4.1 Bayesian games** — the containing framework.
- **4.3 Bayesian Nash equilibrium** — the solution concept defined on type-contingent strategies.
- **4.4 Perfect Bayesian equilibrium** — dynamic version with belief updating on information sets.
- **4.5 Signaling games** — canonical two-stage Bayesian game with sender/receiver types.
- **4.6 Sequential equilibrium** — consistency of off-path beliefs.
- **6.1 Mechanism design** — design a game for a given type space.
- **6.3 Auction theory** — private valuations as types drawn from a prior.

Structural Compression

Invariant: A type encodes all payoff-relevant private information plus all levels of belief about other players' information; under coherence, an infinite hierarchy of beliefs collapses to a point in a universal type space, and any finite type space with a common prior is strategically equivalent to its embedding in that universal space.

Mechanism: Iterated belief-about-belief construction modulo coherence constraints; Bayesian conditioning on types at each level; common prior as a consistency condition tying all interim beliefs to a single joint distribution.

Cross-System Echo: Type spaces generalize the "state" variable in stochastic processes and POMDPs: a type is a sufficient statistic for a player's private information and all higher-order beliefs, just as a POMDP belief state is a sufficient statistic for the agent's knowledge about the hidden environment. The universal type space result parallels the Kolmogorov extension theorem for stochastic processes.

Exercises

1. For a two-player Bayesian game with $|\Theta_1| = |\Theta_2| = 2$, enumerate all coherent belief hierarchies at levels 1 and 2 and count them.

2. Show that under common prior, the level-k belief of player i can be derived from the common prior by k iterations of Bayes' rule.

3. Construct a two-player example where players have coherent interim beliefs but no common prior is compatible with them.
4. For the 2×2 common prior in the worked example, compute the level-2 belief of Player 1 of type H about Player 2's belief about θ_1.
5. Define strategic equivalence of type spaces formally and verify it for a simple example with two different finite type spaces encoding the same payoff structure.
6. In a poker context, compute the interim belief of BTN holding $A\spadesuit K\spadesuit$ about BB's hand, given BB's preflop calling range is $\{AK, AQ, AJ, KQ, KJ, QJ\}$ (all suits). Account for card removal.
7. **Argue, structurally,** why the universal type space result is necessary for the foundations of Bayesian game theory, even though applied work always uses finite type spaces. What would go wrong foundationally without it?

Game Theory · Module 4 — Information And Beliefs · Node 4.2
Concepts: bayesian-updating, prior-posterior, incomplete-information
Prerequisites: 4.1
Enables: 4.3, 4.4, 4.5, 4.6

hbar.university — own.your.cognition