Beyond Common Knowledge: A Belief-Reasoning Framework for Non-Equilibrium Games and Its Graphical Representation (Extended Abstract)

Yiming Bai[1], **Sheng Zhong**[2]
[1]Independent Researcher
[2]Sichuan University
baiyiming2013@gmail.com, zhsh@scu.edu.cn

Abstract

Classical Game Theory heavily relies on the "Common Knowledge" assumption, which mandates that all players know the game structure and others' rationality ad infinitum. However, in asymmetric information scenarios, this assumption often collapses. This paper proposes a **Belief-Reasoning Knowledge Model** that decouples "States" from "Observable Features." We introduce the **Belief Forest**, a graphical data structure to visualize and compress nested belief chains. We demonstrate the model's efficacy through the "Hat Game" and a specialized **Full-Belief Static Game** analysis of the "Huarong Path" (Huarongdao) problem. We show that agents can achieve belief-driven strategies by assigning probability weights to belief trees, even in the absence of common knowledge.

1 Introduction

The concept of Nash Equilibrium (Nash 1951; Fudenberg and Tirole 1991) provides a powerful tool for predicting outcomes in strategic interactions. However, its practical application is often constrained by the **Common Knowledge** assumption, which mandates that the game structure and player rationality are known by all players ad infinitum (Aumann 1976; Binmore and Brandenburger 1989). In many real-world scenarios—ranging from the tactical maneuvers in the classic *Huarong Path* to decentralized multi-agent systems (MAS)—information is inherently asymmetric and non-common.

Traditional epistemic models represent knowledge as partitions over a universal state space Ω (Hintikka 1962; Halpern 1986). While theoretically elegant, this approach faces two critical challenges: (1) **Ambiguity of States:** A "state" $\omega \in \Omega$ must capture every possible detail of the universe, making Ω exponentially large and conceptually elusive. (2) **The Inference Gap:** There is no formal mechanism to describe how an agent transitions from a physical observation to a refined set of possible worlds.

Unlike Kripke/partition models, which define epistemic accessibility directly over a global state space, our framework inserts an explicit feature-observation layer and models knowledge construction as inference from observations to possible states. Specifically, this paper proposes a **Belief-Reasoning Knowledge Model**. By decoupling abstract "States" from observable "Features," we allow agents to construct subjective knowledge through an active inference process. This framework not only provides a computational data structure for high-order nested beliefs—the **Belief Forest**—but also enables the derivation of belief-driven strategies in non-equilibrium environments.

2 The Model: Feature-Based Inference

We move away from the assumption that agents directly perceive partitions of Ω. Instead, we model the reasoning process as a mapping from perceived features to state sets.

2.1 Decoupling Features and States

Let $C = \{c_1, c_2, \ldots, c_n\}$ be the **Universal Feature Set**, where each c_k represents an atomic, observable property of the environment.

Definition 1 (Feature Description Mapping): Each state $\omega \in \Omega$ is characterized by a specific set of features. We define the mapping $\mathcal{F} : \Omega \to 2^C$ such that:

$$\mathcal{F}(\omega) = \{c \in C \mid \omega \text{ possesses feature } c\} \tag{1}$$

In this framework, a state ω is no longer a primitive black box but is uniquely identified by its "feature fingerprint" $\mathcal{F}(\omega)$.

2.2 Inference Function and Uncertainty Collapse

Agents construct knowledge by filtering the state space through their observations.

Definition 2 (Observation Set): For each agent i, the observation set $I_i \subseteq C$ represents the subset of features perceived by agent i.

Definition 3 (Impossible Feature Set): For each state $\omega \in \Omega$, we define $\bar{C}(\omega) \subseteq C$ as the set of features that are *impossible* (contradictory) under ω. This induces a three-way partition of C: features that are positively present $\mathcal{F}(\omega)$, features that are impossible $\bar{C}(\omega)$, and features whose status is unknown under ω.

Definition 4 (Inference Function): We define the mapping $P : 2^C \to 2^\Omega$ such that for any observation set I:

$$P(I) = \{\omega \in \Omega \mid I \cap \bar{C}(\omega) = \varnothing\} \tag{2}$$

Equation (2) formalizes the reasoning process: an agent eliminates all states where any observed feature is *impossible*, rather than requiring all observed features to be positively confirmed. This formulation reflects partial observability, where agents can often rule out contradictions without fully confirming feature presence. This three-way partition (present / unknown / impossible) is more expressive than a binary (present / absent) model, as it preserves epistemic uncertainty about features whose status is indeterminate. The impossible-feature formulation is strictly more general, allowing features whose status is neither confirmed nor ruled out. Note that when every feature is definitively present or impossible under each state (i.e., $\mathcal{F}(\omega) \cup \bar{C}(\omega) = C$ for all ω), Definition 4 reduces to the stricter condition $I \subseteq \mathcal{F}(\omega)$; the three-way formulation is thus a natural generalization.

2.3 Properties of the Knowledge Model

A critical advantage of this formulation is the formalization of knowledge growth as a monotonic process.

Theorem 1 (Knowledge Monotonicity): *For any two observation sets I_1 and I_2, if $I_1 \subseteq I_2$, then $P(I_2) \subseteq P(I_1)$.*

Proof: According to Definition 4, $\omega \in P(I_2)$ implies $I_2 \cap \bar{C}(\omega) = \varnothing$. Since $I_1 \subseteq I_2$, we have $I_1 \cap \bar{C}(\omega) \subseteq I_2 \cap \bar{C}(\omega) = \varnothing$, which implies $\omega \in P(I_1)$. Thus, $P(I_2) \subseteq P(I_1)$. □

This theorem ensures monotone uncertainty reduction during inference. When $|P(I)| = 1$, the agent has achieved perfect knowledge regarding the current state. This property ensures that belief refinement can be implemented as incremental filtering, which is crucial for computational tractability.

2.4 Higher-Order Observations

In multi-agent systems, agents reason about the perceptions of others. We define $I_{ki} = \{c \in I_k \mid k \text{ believes } i \text{ can observe } c\}$: the subset of k's observations that k attributes to i. This definition is informal and serves to motivate higher-order observation structures. We do not formalize the belief semantics here; it serves as a notational device for higher-order observation. It extends recursively to $I_{i_n \cdots i_1}$, and adjacent identical subscripts collapse ($I_{\ldots ii \ldots} = I_{\ldots i \ldots}$), eliminating redundant self-reflection in belief chains. This structure relates to probabilistic common knowledge (Brandenburger and Dekel 1987; Monderer and Samet 1989) and transforms infinite-dimensional epistemic reasoning into finite combinatorial operations over C, providing the foundational data for the Belief Forest.

3 The Belief Forest: Structure and Generation

To handle the recursive nature of higher-order beliefs without triggering a state-space explosion, we map the inference results into a graphical structure. Unlike Kripke structures, which expand exponentially with nested accessibility relations, the Belief Forest enables compression via structural equivalence. This section defines the **Belief Forest** and the core algorithm used to compress redundant reasoning paths.

3.1 Formal Definitions

Definition 5 (Belief Chain): A belief chain L is a sequence $L = (i_n, i_{n-1}, \ldots, i_1, E)$, representing the nested proposition: "Agent i_n believes that agent i_{n-1} believes ... that agent i_1 knows the event set E" (Fagin et al. 2003). Here, $E \subseteq \Omega$ is derived from the inference function $P(I_{i_n \ldots i_1})$.

Definition 6 (Belief Tree and Forest): A Belief Tree T_i for agent i is a rooted in-tree (i.e., a directed tree where all edges point toward the root) formed by the union of all belief chains originating from i. A Belief Forest is the collection of all players' belief trees within a specific reference frame.

3.2 The Node Merging Algorithm

The primary innovation in our graphical representation is the **Equivalence Node Merging Rule**, which ensures computational tractability by collapsing redundant logical paths.

Rule 1 (Structural Compression): Two nodes v_1 and v_2 in a belief tree are merged if they satisfy the following conditions:

1. They exist at the same depth of the belief chain (same recursion order).
2. They represent the perspective of the same agent.
3. Their associated possible state sets are identical: $P(I_{v_1}) = P(I_{v_2})$.

This merge preserves all subsequent reasoning paths: since both nodes share the same agent, depth, and possible state set, the downstream chains are logically identical, and the merged subtree retains full strategic equivalence.

3.3 Case Study: The Hat Game Evolution

We demonstrate the model through the 3-agent Hat Game. Three agents $\{A, B, C\}$ each wear a white hat (W). They can see others' colors but not their own. The state space is $\Omega = \{WWW, WWB, \ldots, BBB\}$.

Phase 1: Initial Parallel Trees Initially, C observes features $I_C = \{c_A, c_B\}$, where c_A and c_B denote "A is wearing White" and "B is wearing White." Each observation independently generates a belief tree, producing two disjoint trees in C's mental model (Figure 1).

Phase 2: Merging via Rule 1 Before any announcement, both trees from Phase 1 already share the same root agent C and the same $P(I_C)$—their possible state sets coincide, enabling structural merging by Rule 1 into a single tree (Figure 2).

Phase 3: Symmetric Common Knowledge A public announcement "at least one hat is white" (c_{pub}) is added to every agent's observation set. In C's mental model, this restructures the belief tree into the fully symmetric form shown in Figure 3. While each agent already privately knew this fact, the announcement elevates it to *common knowledge*, restructuring higher-order beliefs. A and B construct identical trees in their own reasoning. The three phases illustrate a general pattern: private observation generates belief trees that remain subjective (Phase 1–2), while public announcement restructures them into a symmetric common-knowledge form (Phase 3).

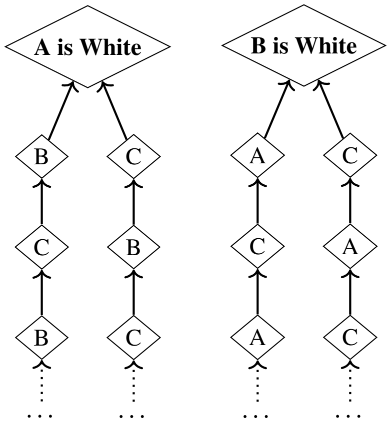

Figure 1: C's Initial Belief Forest. Observing A and B are white creates two parallel, disconnected belief trees.

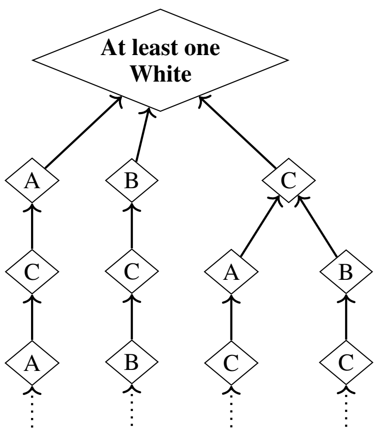

Figure 2: C's Belief Tree after merging Phase 1 trees via Rule 1.

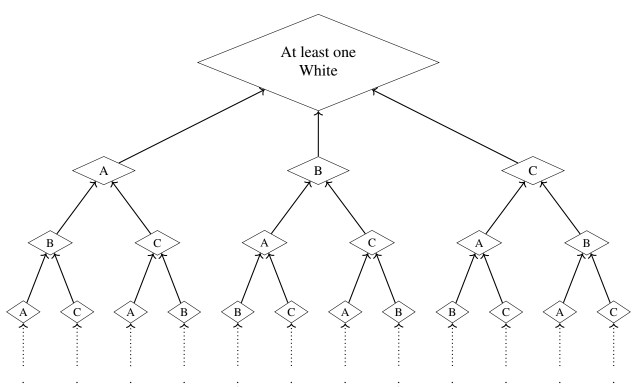

Figure 3: C's Belief Tree after the public announcement. Under common knowledge, A and B hold identical trees.

Cao / Liang	Ambush Huarong (s_H)	Ambush Main Road (s_M)
Take Huarong (s_H)	(-1, 1)	(1, -1)
Take Main Road (s_M)	(1, -1)	(-1, 1)

Table 1. Payoff Matrix for the Huarong Path Game

4 Application: Full-Belief Static Games

The Belief Forest framework allows for the formalization of strategic decisions under non-common knowledge.

Definition 7 (Full-Belief Static Game): A full-belief static game is a tuple $G = (N,\ \{S_i\},\ \{u_i\},\ \{B_i\})$, where N is the player set, S_i is player i's strategy set, u_i is the payoff function, and $B_i = (T_i, \theta_i)$ is player i's **belief structure**—consisting of the belief tree T_i and a weight function $\theta_i : V(T_i) \to [0, 1]$ with $\sum_v \theta_i(v) = 1$. Each node v corresponds to a reasoning path from v to the event root, and $\theta_i(v)$ quantifies i's confidence in that path; weights concentrate on nodes representing i's distinct beliefs. The weights θ_i are *subjective parameters*, analogous to priors in Bayesian analysis; they may be elicited from domain knowledge, calibrated empirically, or derived from historical data.

4.1 Decision-Making via Belief Weights

Definition 8 (Subjective Expected Utility): The expected utility for strategy $s_i \in S_i$ is:

$$U_i(s_i) = \sum_{v \in T_i} \theta_i(v) \cdot u_i(s_i, s_{-i}(v)) \tag{3}$$

where $s_{-i}(v)$ is the anticipated strategy of opponents at node v. Under the **belief-proportional decision rule**, player i's mixed strategy (Aumann 1974) directly mirrors the probability distribution over opponent actions induced by θ_i: the probability of choosing action a equals i's aggregate confidence that a is safe. We do not assume that agents compute pure best responses to U_i. Instead, the belief-proportional rule directly maps subjective belief weights to mixed strategies, reflecting agents whose actions are driven by graded confidence rather than strict optimization. This rule is consistent with probability-matching and bounded-rationality behaviors observed in experimental settings; we adopt it as the simplest linear mapping from beliefs to actions.

4.2 The Huarong Path Case Study

We reconstruct the asymmetric reasoning between Cao and Liang. The presence of smoke (c_1) serves as the primary observable feature. Table 1 shows the zero-sum payoff structure.

Belief Chain Construction As shown in Figure 4, Cao's decision process is a chain of nested beliefs. Cao observes smoke and reasons: "Liang knows I see the smoke." This extends the chain to higher orders.

Derivation of Belief-Driven Strategy Cao's weights $\{\theta_k\}$ (Figure 4) attach to the Liang nodes at odd depths—Cao's distinct beliefs—with $\sum_k \theta_k = 1$. The intervening Cao nodes at even depths collapse via the adjacent-subscript rule of Section 2.4 ($I_{\dots ii \dots} = I_{\dots i \dots}$) and carry no independent weight. Parity determines the predicted ambush: θ_1

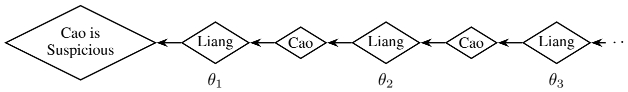

Figure 4: Cao's belief chain. θ_k is Cao's weight on the $(2k-1)$-th order belief; odd-order beliefs predict Liang at the Main Road, even-order predict Liang at Huarong. "Suspicious" is an informal interpretation of the observed signal (smoke).

("Liang knows I am suspicious") and all odd orders predict Liang at the Main Road; θ_2 ("Liang counter-bluffs: knows I know he knows I am suspicious") and all even orders predict Liang at Huarong. Let P denote Cao's subjective probability that Liang ambushes Huarong:

$$P = \theta_2 + \theta_4 + \theta_6 + \cdots \tag{4}$$

For the two-level case $\theta_1 + \theta_2 = 1$, this simplifies to $P = \theta_2 = 1 - \theta_1$.

By the belief-proportional decision rule (Def. 8), Cao's probability x of choosing Huarong equals his aggregate confidence that Huarong is safe:

$$x = 1 - P = \theta_1 + \theta_3 + \cdots \tag{5}$$

Expanding Cao's expected utility from the payoff matrix (Table 1):

$$\begin{aligned} U_{\text{Cao}}(x) &= x\big[P \cdot (-1) + (1-P) \cdot 1\big] \\ &\quad + (1-x)\big[P \cdot 1 + (1-P) \cdot (-1)\big] \\ &= (2x - 1)(1 - 2P) \end{aligned} \tag{6}$$

Substituting $x = 1 - P$:

$$U_{\text{Cao}} = (1 - 2P)^2 = (2P - 1)^2 \tag{7}$$

Since $(2P-1)^2 \geq 0$ for all $P \in [0, 1]$, the believing agent always achieves non-negative expected utility, with equality *only* at $P = 0.5$. This has two critical implications:

Case 1: Deterministic first-order belief ($\theta_1 = 1$). All weight on the first-order chain gives $P = 0$ and $x = 1$: Cao takes Huarong with certainty, achieving $U = 1$. This matches the *Romance of the Three Kingdoms* narrative, where Cao Cao, fully convinced Liang is one step ahead (smoke bluff, real ambush on the Main Road), chose the shorter path.

Case 2: Common knowledge ($\theta_1 = \theta_2 = 0.5$). When neither player possesses private information, beliefs become symmetric: $P = 0.5$, $x = 0.5$, and $U = 0$. This is exactly the Nash equilibrium of the matching-pennies game, where both players adopt uniform mixed strategies. **Nash equilibrium thus emerges as a special case of our model—the degenerate case where beliefs collapse to common knowledge.** We demonstrate this here for the symmetric zero-sum (matching-pennies) structure; generalization to broader game classes depends on the alignment between belief weights and best-response structure.

We define the *belief rent* as $R(P) = U(P) - U(0.5) = (2P - 1)^2$: the expected-utility surplus that an agent earns, relative to the Nash baseline, by holding a private belief $P \neq 0.5$. Any deviation from Nash yields strictly positive belief rent, quantifying the strategic advantage of asymmetric information. By symmetry, Liang's analysis from his own perspective yields the same formula with the payoff roles reversed.

4.3 Relationship to Harsanyi's Framework

Harsanyi's incomplete information game (Harsanyi 1967–1968) models private information through a type space: each player i has a type $t_i \in T_i$ drawn from a common prior p_i. The pair (t_i, p_i) captures all private information, but the structure is inherently *single-layered*—players hold beliefs about types, not about others' beliefs about types.

Our Belief Forest generalizes this along two dimensions:

(1) Multi-layered depth. The observation structure $I_{ijk\ldots}$ and belief forest $\mathcal{F}_i$ encode beliefs of arbitrary depth. Harsanyi's (t_i, p_i) is the special case where the tree has depth 1: in simple two-player binary settings, type t_i maps to a single observation I_i, and prior p_i maps to the first-level weight θ_1. Higher-order beliefs ("i believes j believes k knows…") extend naturally beyond what (t_i, p_i) can express.

(2) Subjective heterogeneity. Harsanyi requires a *common prior* over types (the consistency condition). In our framework, each player constructs a subjective belief forest from private observations, without assuming a shared prior. This accommodates genuinely asymmetric reasoning where players disagree not only about states but about the structure of each other's beliefs.

When the forest is truncated to depth 1 and a common prior is imposed on θ, our model recovers the standard Bayesian-game formulation, establishing Harsanyi's framework as a special case.

5 Conclusion

This paper replaces the rigid common knowledge assumption with a dynamic, feature-based inference process. The Belief Forest and its associated merging rules provide a computationally efficient way to represent and solve games with asymmetric information. The Huarong Path derivation reveals a key result: $U = (2P - 1)^2 \geq 0$, with Nash equilibrium ($P = 0.5$, $U = 0$) emerging as the degenerate special case where beliefs collapse to common knowledge. This illustrates how classical equilibrium behavior can emerge as a boundary case under specific belief-alignment conditions.

Several directions remain open: formal complexity analysis of belief forest merging, extension to dynamic games with sequential belief revision, and relaxation of the belief-proportional rule to accommodate bounded-rationality models such as quantal response equilibrium.

References

Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1): 67–96.

Aumann, R. J. 1976. Agreeing to disagree. *The Annals of Statistics*, 4(6): 1236–1239.

Binmore, K.; and Brandenburger, A. 1989. Common knowledge and game theory. In *Essays on the Foundations of Game Theory*. Blackwell, Oxford.

Brandenburger, A.; and Dekel, E. 1987. Common knowledge with probability. *Journal of Mathematical Economics*, 16(3): 237–245.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 2003. *Reasoning about Knowledge*. MIT Press.

Fudenberg, D.; and Tirole, J. 1991. *Game Theory*. MIT Press.

Halpern, J. Y. 1986. Reasoning about knowledge: An overview. In *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference*, 1–17. Morgan Kaufmann.

Harsanyi, J. C. 1967–1968. Games with incomplete information played by "Bayesian" players, I–III. *Management Science*, 14(3): 159–182.

Hintikka, J. 1962. *Knowledge and belief: An introduction to the logic of the two notions*. Cornell University Press.

Monderer, D.; and Samet, D. 1989. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2): 170–190.

Nash, J. 1951. Non-cooperative games. *Annals of Mathematics*, 54(2): 286–295.