Preference Estimation via Opponent Modeling in Multi-Agent Negotiation

Anonymous ACL submission

Abstract

Automated negotiation is essential for facilitating decision-making in societies with diverse stakeholders. While opponent modeling is critical, particularly in multi-party, multi-issue settings, conventional numerical-only methods fail to incorporate the qualitative context embedded in natural language, leading to unstable predictions. Although Large Language Models (LLMs) enable context-aware reasoning, they often struggle to maintain inferential consistency during prolonged interactions. To bridge this gap, we propose a novel preference estimation method integrating natural language information into a structured Bayesian opponent modeling framework. Our approach leverages LLMs to extract qualitative cues from utterances and converts them into probabilistic formats for dynamic belief tracking. Experimental results using a multi-party benchmark demonstrate that our framework achieves high agreement rates among all participants and superior preference estimation accuracy. These findings suggest that the integration of structured probabilistic reasoning and natural language understanding can facilitate robust consensus-building in ambiguous social interactions.

1 Introduction

In modern society, automated negotiation is a pivotal technology for conflict resolution and efficient consensus-building among diverse stakeholders (Memon et al., 2025; Bagga et al., 2021). Historically, the field has matured through integrated development environments like GENIUS (Lin et al., 2014) and international competitions such as ANAC (Baarslag et al., 2012). A significant milestone was the BOA architecture (Baarslag et al., 2014), which standardized negotiating agents into three decoupled components: the Bidding strategy, Opponent model, and Acceptance strategy. Within multi-party, multi-issue settings, opponent modeling remains essential for strategic decision-making (Baarslag et al., 2016). Traditionally, these models have evolved through Bayesian learning (Zeng and Sycara, 1998; Hindriks and Tykhonov, 2008) and reinforcement learning (He et al., 2016), primarily estimating utility functions from numerical proposal histories.

However, numerical-only methods struggle to capture qualitative contexts, leading to unstable estimation under high information uncertainty (Baarslag et al., 2016).

To address these limitations, integrating Large Language Models (LLMs) into negotiation and decision-making frameworks has gained traction (Abdelnabi et al., 2024; Fu et al., 2023). LLMs possess sophisticated capabilities for context understanding and Theory of Mind (ToM) (Kosinski, 2023; Chan et al., 2024), enabling the extraction of qualitative preference signals typically lost in conventional models. Nevertheless, directly applying reasoning techniques like CoT (Wei et al., 2022), ToT (Yao et al., 2023), or MAD (Liang et al., 2024) to LLM-based agents reveals new challenges: a lack of strategic consistency during prolonged negotiations (Chan et al., 2024), fragile generalization across different problem settings (Zhao et al., 2025), and an exponential increase in inference complexity as information grows (Abdelnabi et al., 2024). In addition, prior work on natural language negotiation using LLMs (Chen et al., 2024; Chan et al., 2024) has largely focused on intent inference in static or short-horizon evaluation settings, where strategic dynamics are limited. Such approaches often lack a formal mechanism for belief updating over time, thereby hampering stable preference tracking in dynamic negotiation scenarios.

To address these challenges, we propose a novel preference estimation method that integrates natural language signals from dialogue into a structured Bayesian framework. Our approach utilizes LLMs to extract qualitative cues and subsequently converts these cues into a format compatible with

probabilistic models for dynamic belief tracking.

Our main contributions are summarized as follows. First, we proposed an integrated framework that complements qualitative intent extraction via LLMs with quantitative preference estimation through Bayesian inference. Second, we demonstrated that the proposed method achieves superior preference estimation performance in complex multi-party scenarios compared to baselines relying solely on numerical data or direct LLM inference. Finally, we revealed that our framework improves agreement success rates even under high uncertainty, thereby facilitating more effective autonomous negotiation.

2 Problem Formulation

We adopt the Scorable Negotiation framework Abdelnabi et al. (2024). Let $P = \{p_1, \ldots, p_N\}$ denote the set of parties and $I = \{i_1, \ldots, i_M\}$ the set of issues, where each issue i_m has a finite option set $O_m = \{o_{m,1}, \ldots, o_{m,K_m}\}$. A deal d_t proposed at round t is defined as a combination of options $d_t = \{o_1^*, \ldots, o_M^*\}$, where each $o_m^* \in O_m$ is selected for the corresponding issue i_m.

Each party p_n holds a private score function $s_{n,m}(o_m^*)$ and the utility of a deal d_t for party p_n is defined as the sum of these scores:

$$U_n(d_t) = \sum_{m=1}^{M} s_{n,m}(o_m^*) \tag{1}$$

Upon reaching an agreement, each party receives the utility $U_n(d_t)$. Otherwise, each party receives their Best Alternative to a Negotiated Agreement (BATNA), represented by a private reservation threshold τ_{p_n}.

The left panel of Figure 1 illustrates the flow of the negotiation. The negotiation lasts for up to 24 rounds. In each round t, a designated party proposes a deal d_t and a natural language utterance u_t, without revealing score functions; parties infer others' preferences from the history of d_t and u_t. The success of the negotiation is determined by the deal d_t in the final round. An agreement is reached if and only if at least a minimum required number of parties, including all veto holders, satisfy $U_n(d_t) > \tau_{p_n}$.

3 Bayesian Preference Estimation Method

In this chapter, we describe our method for explicitly estimating opponents' preferences at each negotiation round. Our approach builds on the Bayesian opponent modeling framework established by Hindriks and Tykhonov (2008), extending it to integrate natural language information using a LLM. The right panel of Figure 1 illustrates the specific mechanism of our proposed opponent modeling.

3.1 Model Representation and Hypothesis Space

First, we define the representation of the opponent's strategy and the space of possible preferences.

Hypothesis Space

To estimate the opponent's score function, the agent maintains a finite set of candidate hypotheses $H = \{h_1, \ldots, h_K\}$. Each hypothesis $h_k \in H$ represents a specific combination of a weight vector $\mathbf{w}^{(k)} = [w_1^{(k)}, \ldots, w_M^{(k)}]$, denoting the relative importance of each issue, and a vector of evaluation functions $\mathbf{v}^{(k)} = [v_1^{(k)}, \ldots, v_M^{(k)}]$, representing the preference shapes for each issue.

Estimated Utility Function

Under a given hypothesis h_k, the estimated utility $U(d_t; h_k)$ of a deal d_t is modeled as an additive utility function. It is calculated as the weighted sum of the evaluation functions:

$$U(d_t; h_k) = \sum_{m=1}^{M} w_m^{(k)} \cdot v_m^{(k)}(o_m^*) \tag{2}$$

Likelihood Based on Numerical Offers

Assuming the opponent follows a concession-based strategy, we define the likelihood $P(d_t \mid h_k)$ of observing a deal d_t under hypothesis h_k. This is based on the proximity between the estimated utility $U(d_t; h_k)$ and $u'(t)$:

$$P(d_t \mid h_k) \propto \exp\left(-\frac{(U(d_t; h_k) - u'(t))^2}{2\sigma^2}\right) \tag{3}$$

Here, $u'(t)$ represents the opponent's assumed target utility at round t. This value is modeled based on the premise that the agent gradually lowers its aspiration level in accordance with its concession strategy as the negotiation progresses.

3.2 Linguistic Likelihood Estimation via LLM

We describe how linguistic utterances are converted into probabilities over opponent preferences.

Signal Extraction via LLM

We employ an LLM to parse an utterance u_t into structured signals z_t. Each signal z_t is defined as a tuple consisting of two attributes: (1) *Target* (the issues or options being predicted by the signal. Specifically, a target can be formed in four

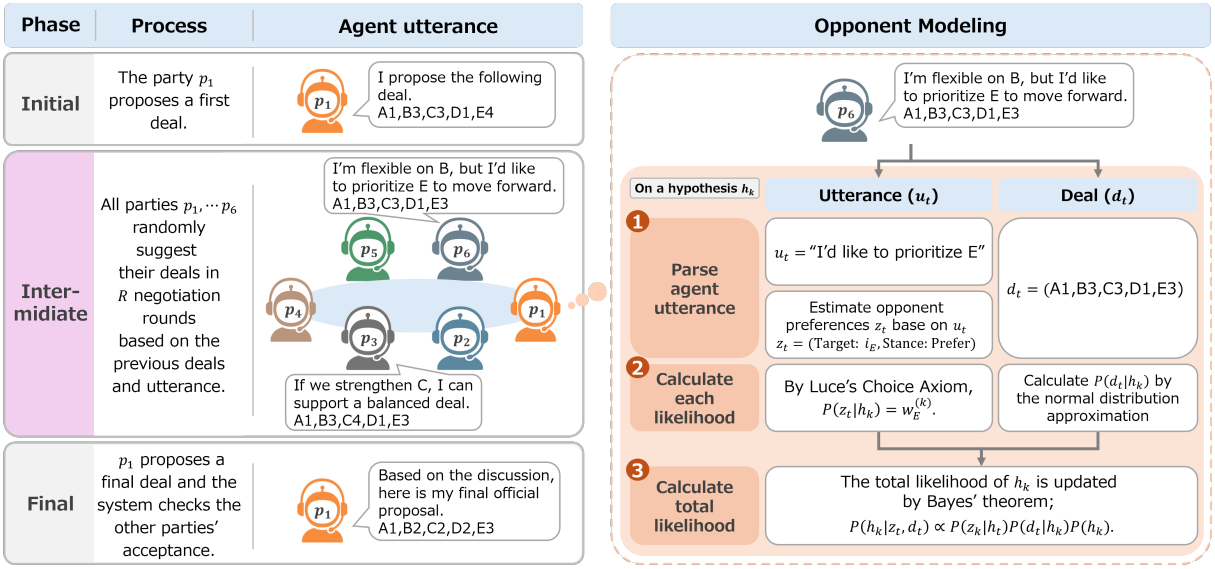

Figure 1: Overview of the negotiation flow (left) and the opponent modeling process by agent p_1 using the proposed framework (right).

possible ways: (i) a single issue, (ii) a comparison between two issues, (iii) a single option, or (iv) a comparison between two options), and (2) *Stance* (the attitude toward the target, such as "Prefer" or "Oppose"). This allows the agent to extract qualitative information—for example, "Issue i_1 is important" or "Option $o_{1,1}$ is preferable to $o_{1,2}$"—into a format suitable for probabilistic computation.

Likelihood Calculation based on Luce's Axiom To quantify the likelihood $P(z_t \mid h_k)$, we apply Luce's Choice Axiom. For instance, the probability of observing a signal that indicates a preference for issue i_x is defined as:

$$P(z_t \in \mathcal{Z}_{i_x,\text{pref}} \mid h_k) = \frac{w_x^{(k)}}{\sum_{m=1}^{M} w_m^{(k)}} \quad (4)$$

where $\mathcal{Z}_{i_x,\text{pref}}$ denotes the set of signals representing a "Prefer" stance toward issue i_x. Similarly, likelihoods for comparison or opposition are calculated based on the relative ratios of components within $\mathbf{w}^{(k)}$ and $\mathbf{v}^{(k)}$ for each hypothesis h_k.

3.3 Preference Update via Multimodal Observations

We now integrate numerical offers and linguistic signals into a unified Bayesian update rule. Assuming conditional independence between the numerical offer d_t and the linguistic signal z_t given a hypothesis h_k, the posterior distribution over hypotheses is updated as follows:

$$P(h_k|d_t, z_t) = \frac{P(d_t, z_t|h_k)P(h_k)}{P(d_t, z_t)} \propto P(d_t|h_k)P(z_t|h_k)P(h_k) \quad (5)$$

Here, $P(d_t|h_k)$ corresponds to the likelihood calculated from the numerical offer, and $P(h_k)$ represents the prior probability of the hypothesis. Our rule weights these terms by the linguistic likelihood $P(z_t|h_k)$, thereby incorporating consistency with the opponent's stated preferences into the estimation process.

4 Experiments

4.1 Experimental Setup

We evaluate our method using the multi-agent negotiation environment proposed by Abdelnabi et al. (2024).

Negotiation Scenario

The scenario involves a negotiation on the construction of a sports facility with $N = 6$ stakeholders, including two veto holders (p_1, p_2) and $M = 5$ issues. The scenario is characterized by relatively diverse preferences among the parties, making consensus building highly challenging; among all 720 possible combinations in the deal space, only 2.9% (21 deals) satisfy the reservation thresholds τ_{p_n} for at least five parties including veto holders, and only 0.4% (3 deals) satisfy the thresholds for all six parties.

Compared Methods

Proposed: The method described in Chapter 4.

We evaluate two configurations: *p1*, where only the leader p_1 performs estimation, and *all*, where all agents perform mutual estimation.

Base-LLM: The original implementation where agents negotiate based on prompts without explicit preference estimation.

Base-OM (Baseline-Opponent Modeling): A conventional Bayesian approach that utilizes only the history of deals d_t for estimation.

LLM-PE (LLM Preference Estimation): A method where an LLM directly infers the numerical values of the opponents' score functions s without a structured Bayesian framework.

For all agents, GPT-4.1 was employed as the underlying model.

4.2 Evaluation Metrics

To assess the negotiation outcomes, we calculate the mean values of the following three metrics across 500 independent negotiation trials for each method:

FAR (Full Agreement Rate): The ratio of trials where all six parties reached a consensus.

PAR (Partial Agreement Rate): The ratio of trials where an agreement was reached by at least five parties including veto holders in the final round d_t.

LAR (Latent Agreement Rate): The ratio of trials where at least one valid deal was proposed during the 24-round process.

Furthermore, we measure the Mean Squared Error (MSE) between the estimated score functions and the true score functions s to evaluate estimation accuracy.

4.3 Results and Discussion

In this section, we evaluate the negotiation outcomes and preference estimation accuracy. Table 1 and Table 2 present these results.

Table 1: Negotiation Outcomes

Method	FAR	PAR	LAR
Proposed (p1)	0.462	0.780	0.958
Proposed (all)	**0.618**	0.890	0.982
Base-LLM	0.368	0.764	0.974
Base-OM (p1)	0.452	0.816	0.968
Base-OM (all)	0.558	**0.918**	**0.990**
LLM-PE (p1)	0.400	0.750	0.968
LLM-PE (all)	0.318	0.690	0.934

Analysis of Negotiation Outcomes

The proposed method (*all*) achieved the strongest overall performance on FAR, while maintaining a competitively high PAR. In particular, it obtained the highest FAR (0.618), indicating a strong capability to identify agreements acceptable to all agents under complex multi-agent interactions. This result highlights the advantage of explicitly modeling preferences using linguistic signals extracted by LLMs.

Compared to the single-estimator setting (*p1*), mutual preference estimation (*all*) further improved both FAR and PAR, demonstrating enhanced strategic coordination among agents.

Table 2: Estimation Error (MSE) by p1

Method	May	Cit	Uni	DoT	Env	Avg
Proposed	158.6	216.7	119.6	99.1	200.8	**158.9**
Base-OM	111.9	232.1	154.8	120.2	323.7	188.5
LLM-PE	166.6	237.7	184.9	96.2	128.8	162.8

Analysis of Estimation Accuracy

The proposed method achieved a lower MSE (158.996) than Base-OM approach (188.592), showing that natural language information improves preference estimation. Although its average error was slightly lower (i.e., better) than that of LLM-PE, the proposed method showed more balanced accuracy across agents, enabling less biased preference prediction. This likely helped the agents propose deals d_t that better satisfy the complex multi-party constraints needed for consensus.

5 Conclusion

In this work, we proposed a novel Bayesian preference estimation framework that integrates numerical proposals with qualitative natural language signals extracted from dialogue. Our experiments in a multi-agent, multi-issue negotiation setting demonstrated that the proposed method achieves consistently higher FAR than baselines relying solely on numerical data or direct LLM inference, while also maintaining a high PAR. These findings highlight the potential of combining the linguistic intelligence of LLMs with mathematically rigorous Bayesian inference to facilitate effective conflict resolution.

Limitations

While our framework significantly improves multi-party negotiation, there remain several avenues for future enhancement. First, although we validated our method using a complex benchmark, further research is required to verify its generalizability across more diverse utility structures and larger

agent populations. Second, while we achieve high accuracy by assuming sincere dialogue, robustness could be improved by incorporating mechanisms to account for strategic behaviors like deception or bluffing. Third, while we successfully focused on learning preference shapes, integrating the inference of opponents' reservation values would allow for more sophisticated coordination, especially in ambiguous settings where agreement zones are difficult to identify.

References

Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2024. Cooperation, competition, and maliciousness: Llm-stakeholders interactive negotiation. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc.

Tim Baarslag, {Mark J.C.} Hendrikx, {Koen V.} Hindriks, and {Catholijn M.} Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, 30(5):849–898.

Tim Baarslag, Koen Hindriks, Mark Hendrikx, Alex Dirkzwager, and Catholijn Jonker. 2014. *Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies*, volume 535, pages 61–83. Springer Japan.

Tim Baarslag, Koen Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. 2012. The first automated negotiating agents competition (anac 2010). In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, 383, pages 113–135. Springer, Berlin, Heidelberg.

Pallavi Bagga, Nicola Paoletti, Bedour Alrayes, and Kostas Stathis. 2021. Anegma: an automated negotiation model for e-markets. *Autonomous Agents and Multi-Agent Systems*, 35(2).

Chunkit Chan, Cheng Jiayang, Yauwai Yim, Zheye Deng, Wei Fan, Haoran Li, Xin Liu, Hongming Zhang, Weiqi Wang, and Yangqiu Song. 2024. NegotiationToM: A benchmark for stress-testing machine theory of mind on negotiation surrounding. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4211–4241, Miami, Florida, USA. Association for Computational Linguistics.

Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, and Minlie Huang. 2024. Tombench: Benchmarking theory of mind in large language models. *Preprint*, arXiv:2402.15052.

Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. Improving language model negotiation with self-play and in-context learning from ai feedback. *Preprint*, arXiv:2305.10142.

He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé. 2016. Opponent modeling in deep reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, page 1804–1813. JMLR.org.

Koen Hindriks and Dmytro Tykhonov. 2008. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '08, page 331–338, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Michal Kosinski. 2023. Theory of mind may have spontaneously emerged in large language models. *ArXiv*, abs/2302.02083.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.

Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. 2014. Genius: An integrated environment for supporting the design of generic automated negotiators. *Comput. Intell.*, 30(1):48–70.

Mashal Afzal Memon, Gian Luca Scoccia, and Marco Autili. 2025. A systematic mapping study on automated negotiation for autonomous intelligent systems. *Automated Software Engg.*, 32(2).

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA. Curran Associates Inc.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: deliberate problem solving with large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA. Curran Associates Inc.

Dajun Zeng and Katia Sycara. 1998. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48(1):125–141.

Guangxiang Zhao, Saier Hu, Xiaoqi Jian, Wu Jinzhu, Yuhan Wu, Lin Sun, and Xiangzheng Zhang. 2025.

Large language models badly generalize across option length, problem types, and irrelevant noun replacements. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26825–26834, Suzhou, China. Association for Computational Linguistics.

A Negotiation Scenario Details

A.1 Parties and Roles

Table 3 describes the six stakeholders involved in the Harbour Sport Park negotiation and their strategic characteristics.

Table 3: Roles and Characteristics of Negotiation Parties

Party	Characteristics
SportCo	Proposer and facilitator of the project. Holds veto power over the final decision.
Dept. of Tourism (DoT)	Provider of federal funding. Holds veto power over the final decision.
Environmental League	Prioritizes ecological preservation above all else.
Local Labour Union (LLU)	Advocates for union priority in employment rules.
Other Cities	Demands increased compensation for neighboring municipalities.
Mayor	Participates as the head of the host city.

A.2 Issues and Options

The negotiation consists of five issues, each with three to five options. Table 4 provides the detailed definitions for each option.

Table 4: Definitions of Negotiation Issues and Options

Issue	Option 1	Option 2	Option 3	Option 4	Option 5
A: Infrastructure	Water-based	Amphibious	Land-based	-	-
B: Ecology	Accept damage	Balanced	Max effort	-	-
C: Employment	Union priority	2:1 Ratio	1:1 Ratio	No priority	-
D: Fed. Funding	$3B	$2B	$1B	None	-
E: Compensation	$600M	$450M	$300M	$150M	None

A.3 Preference Profiles

Table 5 presents the private score functions and reservation thresholds (τ) for each party. Under these high thresholds, the negotiation is highly challenging, with only 0.4% of all possible deals satisfying the conditions for a full agreement.

Table 5: Agent Score Functions and Reservation Thresholds (τ)

Party	Threshold (τ)	Issue A	Issue B	Issue C	Issue D	Issue E
SportCo (Veto)	53	[14, 8, 0]	[11, 7, 0]	[0, 5, 10, 17]	[35, 29, 20, 0]	[0, 5, 10, 15, 23]
DoT (Veto)	70	[0, 11, 5]	[0, 20, 25]	[0, 2, 4, 9]	[10, 26, 40, 0]	[4, 8, 15, 12, 0]
Env. League	45	[0, 22, 45]	[0, 25, 55]	[0, 0, 0, 0]	[0, 0, 0, 0]	[0, 0, 0, 0, 0]
LLU	50	[15, 20, 0]	[0, 0, 0]	[42, 35, 25, 0]	[30, 20, 10, 0]	[2, 4, 6, 8, 0]
Other Cities	50	[0, 4, 10]	[0, 0, 0]	[12, 8, 6, 0]	[0, 8, 13, 18]	[60, 45, 30, 15, 0]
Mayor	55	[14, 8, 0]	[12, 8, 0]	[24, 18, 12, 0]	[40, 30, 23, 0]	[0, 2, 4, 7, 10]

B Experimental Setup and Computational Resources

B.1 Implementation and Infrastructure

The proposed framework was implemented by extending the open-source negotiation environment provided by Abdelnabi et al. (2024), which is distributed under the MIT License. All experiments were conducted within a Docker container on an Amazon Web Services (AWS) instance, accessed via SSH from a macOS workstation. The inference process utilized the `gpt4.1-2025-04-14` model via the OpenAI API. Total execution time for 500 trials of the proposed method was approximately several hours.

B.2 Bayesian Hypothesis Space and Parameters

Following the methodology of Hindriks and Tykhonov (2008), the hypothesis space H is defined as follows:

- **Issue Weights (w):** We consider all possible permutations of issue rankings. For $M = 5$ issues, this results in $5! = 120$ weight hypotheses.
- **Evaluation Functions (v):** Preference shapes are modeled as linear functions. Given the issues' options (A, B: 3; C, D: 4; E: 5), the total combination of evaluation functions is $3 \times 3 \times 4 \times 4 \times 5 = 720$ patterns.
- **Likelihood Parameters:** The standard deviation for numerical offer likelihood calculation was set to $\sigma = 1.0$.
- **LLM Settings:** The sampling temperature for the LLM was set to 0 to ensure deterministic reasoning, with `max_tokens` maintained at the default configuration of the `gpt-4.1` model.

C Societal Impact

The primary goal of this research is to enhance the transparency and efficiency of consensus-building. By integrating LLMs with structured Bayesian inference, we provide an interpretable framework where an agent's internal belief state can be inspected. This facilitates smoother decision-making in multi-party organizational contexts. We acknowledge potential risks: while our framework assumes truthful signaling, future iterations could be misused for deceptive strategy optimization. We recommend human oversight in high-stakes deployments.

D Use of AI Assistants

AI assistants were utilized in several stages of this research. Specifically, Large Language Models were employed to assist in (1) literature research and information synthesis, (2) debugging and optimizing parts of the experimental code implementation, and (3) refining and proofreading the manuscript to improve clarity and grammatical accuracy. All final decisions, scientific interpretations, and content verifications were performed by the authors.

E Prompt for Signal Extraction

Figure 2 illustrates the detailed prompt used to extract qualitative signals z_t from negotiation dialogues.

You are an expert in negotiation analysis. Analyze the following chat history and extract opponent modeling signals for each agent.

Negotiation Rules:
{negotiation_rule}

Chat history:
{chat_history}

Task:
For each agent appearing in the chat history, extract structured opponent modeling signals. Identify behavioral signals that indicate their preferences.

Important instructions:

1. Be sure to extract at least one signal for each agent.
2. Include signals that can be inferred by comprehensively considering the chat history and negotiation rules, even if the agent did not directly mention them in their statement. Do not limit signal extraction only to the options proposed in the deal; also extract signals regarding issue preferences and comparisons of preferences between two issues or options.
3. Extract signals in chronological order as they appear in the chat history. Process the conversation from beginning to end, and add signals to the array in the order you encounter them.
4. Classify each signal using the following information:

- **entity**: The type of reference ("issue" or "option")
    - "issue": Refers to the name of an issue (e.g., A, B)
    - "option": Refers to a specific choice within an issue (e.g., A1, B1)
- **signal_type**: The type of signal ("point" or "comparison")
    - "point": A direct preference toward a specific target ("A", "A1")
    - "comparison": A preference comparison between two targets ("A, B", "A1, B1")
- **target**: The specific target (e.g., "A", "A1", "A, B", "A1, B1")
- **stance**: The agent's position toward the target ("prefer" or "oppose"). When the agent gives importance to the target, use "prefer". When the agent devalues or rejects the target, use "oppose".

When extracting signals, pay particular attention to the stated reasons why the agent proposed a specific deal. Accurate capture of each party's preferences by considering their public answers and the flow of proposed deals is required. While the agent is likely to prefer the option they proposed, consider that they may have compromised to accommodate other parties; thus, an even more preferred option might exist.

Using the provided function schema, create a structured response with agent names as keys. Do not return empty results. Always extract at least some signals.

Figure 2: Prompt for Qualitative Signal Extraction